Quality Meshing with Weighted Delaunay Refinement*

Siu-Wing Cheng† Tamal K. Dey‡

September 9, 2003

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. As in previous work, we assume that no input angle is acute. This success is achieved by combining the weight pumping method for sliver exudation and the Delaunay refinement method for boundary conformation.

1 Introduction

In finite element methods a three-dimensional domain is often partitioned with tetrahedra. The quality of their shapes influences the quality of the finite element solution [23]. This motivated the decade long research on generating meshes with guaranteed aspect ratio called *quality meshes* [1, 3, 5, 7, 8, 20, 21, 22]. A considerable literature has built up on the subject, see the surveys and books [2, 12, 15, 24]. 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 [20] 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 because of the axis parallel boxes used in quadtree/octtree 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 [15, 12]. 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 two dimensions [5] which produces an uniform mesh of quality triangles. Ruppert [21], in a pioneering work called Delaunay refinement, showed how circumcenter insertion can be used to produce a quality graded mesh with optimal size.

*This research has been supported by the Research Grant Council, Hong Kong, China, project no. HKUST6088/99E and NSF under grant CCR-9988216.

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

‡Department of Computer and Information Sciences, Ohio State U., Columbus, OH 43210, USA. e-mail: tamaldey@cis.ohio-state.edu

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 three-dimensional domains with boundaries [22]. This refinement eliminates tetrahedra that have large ratio of their circumradius to the shortest edge length. Consequently the resulting mesh satisfies the *radius-edge ratio property*, i.e., all tetrahedra have radius-edge ratio below a threshold. Although most of the poor quality tetrahedra are removed by the radius-edge ratio property, one type of bad tetrahedra called *slivers* are not eliminated, see [9]. 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. [8] proposed the *sliver exudation* method to get rid of the slivers from a Delaunay mesh that already have the radius-edge 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. [13] developed a method to perturb the points so that the unweighted Delaunay tetrahedralization has the radius-edge ratio property and is sliver-free.

Unfortunately, the algorithms of [8, 13] could not handle boundaries and were applied to periodic point sets. In a recent work Edelsbrunner and Guoy [14] experimented with the sliver exudation method which reveals that the technique is very effective in eliminating slivers. After sliver exudation, almost all tetrahedra have angles greater than $5°$ (except that some tetrahedra with angles less than $5°$ survive near the boundary). Thus, boundary handling remains a challenge. Recently, Li and Teng [16] showed 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 [7]. A sliver is destroyed by inserting a random point near its circumcenter. The analysis of this algorithm is a nice example of the confluence of the results developed over the years by Chew [7], Ruppert [21], Shewchuk [22] and Cheng et al. [8].

In this paper we present the first *deterministic* algorithm to construct a Delaunay mesh, with the radius-edge ratio property and without any sliver, of a three-dimensional domain with boundaries. We combine Delaunay refinement with sliver exudation and obtain what we call the *weighted Delaunay refinement.* We show that this technique produces a graded mesh with asymptotically optimal size. Our work can be viewed as an advance along the line of research initiated by Chew and Ruppert [5, 21], carried forward by Dey et al. [9], Shewchuk [22], Cheng et al. [8], Edelsbrunner et al. [13] and Li and Teng [16]. Encouraged by the experimental results of Edelsbrunner and Guoy [14], we believe that weighted Delaunay refinement is imminently practical. Their experiments show that after sliver exudation, relatively few slivers near the domain boundary survive. Thus, we expect that most slivers can be destroyed without adding extra points.

The rest of the paper is organized as follows. Section 2 describes the basic definitions. Section 3 presents our algorithm and its behavior is analyzed in Sections 4, 5 and 6. Section 7 proves the guarantees achieved by our algorithm. We conclude in Section 8.

2 Definitions

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

Quality of tetrahedra. The volumes of tetrahedra in a normalized sense capture their quality. Poor quality tetrahedra have small normalized volume. Let R, L and V be the circumradius, shortest edge length and volume of a tetrahedron τ respectively. We characterize τ by two ratios $\rho(\tau) = R/L$ and $\sigma(\tau) = V/L^3$. If $\rho(\tau)$ exceeds a threshold ρ_0, then we call it *skinny.*

If $\rho(\tau) \leq \rho_0$ and $\sigma(\tau)$ does not exceed a threshold σ_0, then we call τ a sliver. As in previous work [8, 16], we will show that there exist ρ_0 and σ_0 independent of the domain such that our algorithm does not produce any skinny tetrahedron or sliver with respect to ρ_0 and σ_0.

Piecewise linear complex. The domain to be meshed is a bounded volume and its boundary is presented using a *piecewise linear complex* (PLC). A collection $\mathcal{P}$ of vertices, segments and facets in $\mathbb{R}^3$ is called a PLC if (i) all elements on the boundary of an element in $\mathcal{P}$ also belong to $\mathcal{P}$, and (ii) if any two elements intersect, their intersection is a lower dimensional element in $\mathcal{P}$.

Input. We assume that the domain to be meshed is a convex bounded volume, containing a collection of vertices, segments, and facets represented as a PLC. An input angle is the angle between two segments sharing a vertex, a segment and a facet sharing one vertex, or two facets sharing a vertex or a segment. We assume that no input angle is acute. If the input PLC does not bound a convex volume, we enclose it in a large box, mesh the inside of the box, and only keep the tetrahedra covering the original domain. This technique has been used before [12, 21]. We use $\mathcal{P}$ to denote the (possibly extended) input PLC.

Incidence. Two elements in $\mathcal{P}$ are incident if one is in the boundary of the other.

Adjacent elements. We call two elements of $\mathcal{P}$ *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 $\mathcal{P}$ is not adjacent to any other element that is not incident to it.

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

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 *weighted 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 for any point $z \in \widehat{x} \cap \widehat{y}$, $\|x - z\|^2 + \|y - z\|^2 = X^2 + Y^2 = \|x - y\|^2$. That is, $\angle xzy = \pi/2$. So we say that $\widehat{x}$ and $\widehat{y}$ are *orthogonal.* If $\pi(\widehat{x}, \widehat{y})$ is greater (resp. smaller) than 0, then for any point $z \in \widehat{x} \cap \widehat{y}$, $\angle xzy > \pi/2$ (resp. $\angle xzy < \pi/2$) and we say that $\widehat{x}$ and $\widehat{y}$ are *further (resp. closer) than orthogonal* from each other. The *bisector plane* of $\widehat{x}$ and $\widehat{y}$ is the locus of points at equal weighted 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 an edge, a triangle or a tetrahedron. 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 *smallest circumspheres* 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.

For a weighted point set $\mathcal{V}$, 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*

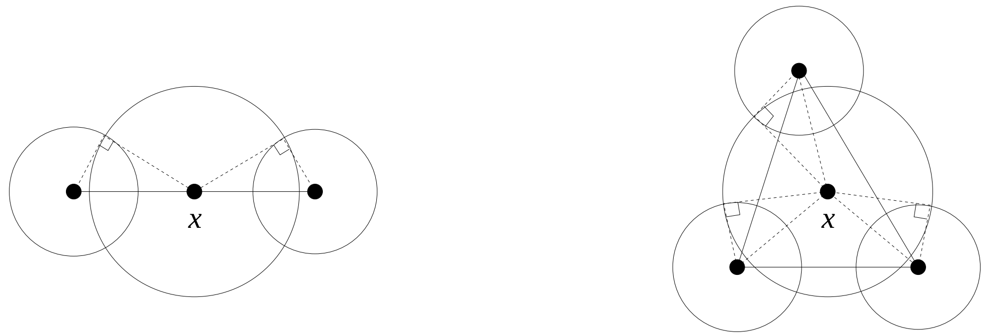

Figure 1: Smallest orthospheres of an edge and a triangle with orthocenter x.

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 [21] and later extended to three dimensions by Shewchuk [22] iteratively inserts circumcenters of 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, 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 because only a finite number of points 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 exudation method of Cheng et al. [8] which eliminates slivers from a Delaunay mesh that already have bounded radius-edge ratio. The key algorithmic tool in sliver exudation is the assignment of weights to unweighted vertices. The weight assignment can be viewed as *pumping* the unweighted vertex to grow it to a sphere (weighted vertex). When an unweighted vertex v is pumped, there is a restriction on the weight to be assigned to v. The weight must be selected from the interval $[0, \omega^2 N(v)^2]$, where $N(v)$ is the Euclidean distance to its nearest vertex and $\omega \in (0, 1/2)$ is a constant. It is shown 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 [8]. This is stated precisely in the following sliver theorem.

Theorem 3.1 (Sliver theorem [8]) *Given a periodic point set $\mathcal{V}$ and a Delaunay triangulation of $\mathcal{V}$ with radius-edge ratio $\leq \rho$, there exists $\rho_0 > 0$ and $\sigma_0 > 0$ and a weight assignment in $[0, \omega^2 N(v)^2]$ for each vertex v in $\mathcal{V}$ such that $\rho(\tau) \leq \rho_0$ and $\sigma(\tau) > \sigma_0$ for each tetrahedron τ in the weighted Delaunay triangulation of $\mathcal{V}$.*

The required weights can be assigned in a deterministic manner for periodic point sets as defined by Cheng et al.[8]. Periodic point sets are infinite points sets without boundaries. So Theorem 3.1 does not give an algorithm for meshing bounded domains. 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.

Our algorithm refines a Delaunay triangulation till it determines that it is safe to pump vertices to remove slivers. In the refinement process it attempts to insert circumcenters of skinny tetrahedra. But, these centers may come close or challenge some boundary elements which are then subdivided. The subdivision process may trigger further subdivision with new vertices. This refinement process is exactly the same as that of Ruppert [21] and Shewchuck [22]. We introduce another stimulus for refinement in preparation for pumping the vertices to remove slivers in a final stage. If a vertex v has a sliver incident to it, we check if the vertex with maximum allowed weight challenges any boundary element. If so, a refinement process is triggered. We will see that even if v does not challenge any boundary element, $\widehat{v}$, the weighted vertex may. At these stages, we only maintain the unweighted Delaunay triangulation of the current vertex set. At the end of the refinement process, when no more boundary element is challenged by weighted or unweighted vertices, we safely pump the vertices to eliminate slivers.

3.2 Subsegments and subfacets

Our algorithm maintains a set of vertices $\mathcal{V}$ which consists of the input vertices initially and it grows as we refine boundary elements and insert circumcenters of skinny tetrahedra.

The vertices in $\mathcal{V}$ on a segment of $\mathcal{P}$ subdivide it into *subsegments*. Let ab be a subsegment. A point p *encroaches* upon ab if p lies inside the smallest circumsphere of ab. Suppose that there is no encroached subsegment. Consider the vertices in $\mathcal{V}$ on a facet of $\mathcal{P}$ in isolation. The two-dimensional Delaunay triangulation of these vertices conforms to the boundary of the facet. Each triangle on the facet is called a *subfacet*. A point p *encroaches* upon a subfacet abc if p lies inside the smallest circumsphere of abc.

Eventually, we will assign weights to vertices in $\mathcal{V}$. This will require us to deal with the weighted versions of subsegments, subfacets, and encroachment. The *weighted-subsegments* are exactly the same as subsegments but possibly with weighted endpoints. A weighted point $\widehat{p}$ encroaches upon a weighted-subsegment ab if $\widehat{p}$ is closer than orthogonal from the smallest orthosphere of ab. See Figure 2 for an illustration. Suppose that there is no encroached weighted-subsegment. Consider the weighted vertices on a facet of $\mathcal{P}$ in isolation. The two-dimensional weighted Delaunay triangulation of these vertices conforms to the boundary of the facet. Each triangle on the facet is called a *weighted-subfacet*. A weighted point $\widehat{p}$ encroaches upon a weighted-subfacet abc if $\widehat{p}$ is closer than orthogonal from the smallest orthosphere of abc.

We will prove several properties for the weighted Delaunay triangulation (Lemmas 3.1–3.5) later. As the weighted case is more general than the unweighted case, these results are applicable for the unweighted Delaunay triangulation as well as when only some of the vertices are weighted.

3.3 Weight assignment

For a vertex u, there are at most two assigned weights, one to check if $\widehat{u}$ encroaches upon a boundary element and possibly another if u participates in an actual pumping in the final stage. Let $\omega_0 < 1$ is a constant chosen in advance. The value of ω_0 will be determined in Section 5. We use the weight $\omega_0^2 f(u)^2$ for encroachment check and the interval $[0, \omega_0^2 f(u)^2]$ to assign weights during pumping. It is important that there is enough space around u when it is pumped. Recall that the vertex $\widehat{u}$ with weight U^2 is equivalent to a sphere centered at u

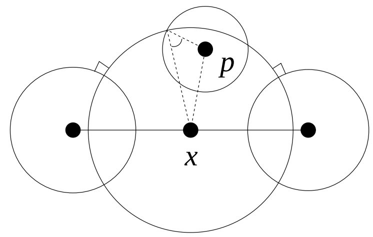

Figure 2: A subsegment is encroached by $\widehat{p}$. Note that the angle shown between $\widehat{p}$ and the smallest orthosphere of the subsegment is less than $\pi/2$. The subsegment is split by its orthocenter x.

with radius U. The assigned weights also impose a requirement that no other vertex is allowed within an Euclidean distance of $2\omega_0 f(u)$ from u. Consequently, after pumping, $\widehat{u}$ can reach up to half of the Euclidean distance to its nearest neighbor. So no two weighted vertices intersect after pumping. We provide an exact statement of this property below.

> VERTEX GAP PROPERTY: For each vertex u in $\mathcal{V}$, the weight of u used for encroachment checking or pumping is at most $\omega_0^2 f(u)^2$ and the Euclidean nearest neighbor distance of u in $\mathcal{V}$ is at least $2\omega_0 f(u)$.

Note that we need to maintain the vertex gap property throughout the algorithm as new vertices are inserted (added to $\mathcal{V}$). As the weights assigned to the vertices are not large compared with inter-point distances, the resulting weighted Delaunay triangulation satisfies many properties of the unweighted one, and many results of the Delaunay refinement carry over to the weighted Delaunay refinement (Lemmas 3.1–3.5). We will prove the vertex gap property in Section 5.

3.4 Locations of centers

Whenever a subsegment is encroached, we split it by inserting its midpoint. Whenever a subfacet is encroached, we split it by inserting its circumcenter. This requires the circumcenter to lie on the facet of $\mathcal{P}$ containing that subfacet. Whenever there is a skinny tetrahedron, we insert its circumcenter. This requires the circumcenter to lie inside the input domain to prevent perpetual growth of the mesh. Although at this point we need these results for Delaunay triangulations, we prove them for weighted Delaunay triangulations which we will need in section 7 for pumping.

Lemma 3.1 *Suppose that the vertex gap property holds. If no weighted-subsegment or weighted-subfacet is encroached, no weighted vertex $\widehat{p}$ intersects a segment or a facet that does not contain p.*

Proof. Assume to the contrary that the lemma does not hold. First of all, $\widehat{p}$ cannot enclose any vertex other than p because of the vertex gap property. Let F be a segment intersected by $\widehat{p}$ if there is any. Otherwise, let F be a facet intersected by $\widehat{p}$. Let q be the projection of p on F. Observe that any point on F lies inside the smallest orthosphere of some weighted-subsegment or weighted-subfacet, or inside some weighted vertex on F (viewed as a sphere). Suppose that q lies inside the smallest orthosphere $\widehat{x}$ of a weighted-subsegment or weighted-subfacet τ. We have $\|p-x\|^2 = \|q-x\|^2 + \|p-q\|^2 < X^2 + P^2$ as $\|q-x\| < X$ and $\|p-q\| < P$. This implies that $\widehat{p}$ encroaches upon τ, a contradiction. Suppose that q lies inside a weighted vertex $\widehat{v}$. Similarly, we get $\|p - v\|^2 = \|q - v\|^2 + \|p - q\|^2 < V^2 + P^2$. This implies that $\widehat{p}$ and $\widehat{v}$ intersect, which contradicts the vertex gap property. □

Lemma 3.2 *Suppose that the vertex gap property holds.*

(i) A weighted-subsegment contains its orthocenter.

(ii) If no weighted-subsegment is encroached, a facet contains the orthocenter of any weighted-subfacet on it.

(iii) If no weighted-subsegment or weighted-subfacet is encroached, the input domain contains the orthocenter of any weighted Delaunay tetrahedron inside the input domain.

Proof. Because of the vertex gap property, (i) is obvious. We present a proof that works for both (ii) and (iii). Let Ω be a facet or the input domain. Let τ be a weighted-subfacet on Ω or a weighted Delaunay tetrahedron inside Ω correspondingly. Let $\widehat{s}$ be the smallest orthosphere of τ. Assume to the contrary that s lies outside Ω.

Let p be a vertex of τ such that ps crosses $\partial\Omega$. By Lemma 3.1, $\widehat{p}$ does not cross $\partial\Omega$, so $\widehat{s}$ must cross $\partial\Omega$ in order to be orthogonal to $\widehat{p}$. Let E be the element in $\partial\Omega$ closest to s. As $\widehat{s}$ is empty of vertices, E is either a segment or facet. Let L_E be the affine hull of E. Let $\widehat{y}$ be the diametral/equatorial sphere of $\widehat{s} \cap L_E$. We make four observations. Refer to Figure 3.

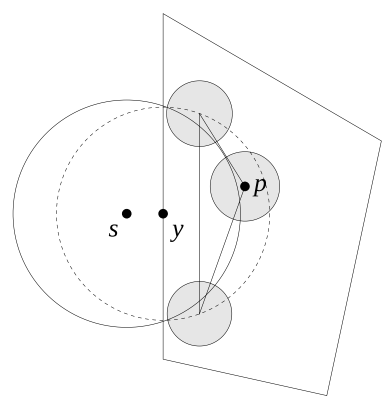

Figure 3: Illustration for Lemma 3.2 when Ω is a facet. The shaded disks are the vertices of the weighted-subfacet whose orthocenter is s. The solid circle is $\widehat{s}$ and the dashed one is $\widehat{y}$.

First, y lies in the interior of E.

Second, the bisector plane of $\widehat{s}$ and $\widehat{y}$ contains E. Note that p and s lie on opposite sides of this bisector plane. So $\pi(\widehat{p}, \widehat{y}) < \pi(\widehat{p}, \widehat{s}) = 0$.

Third, for any vertex v on E, we have $\pi(\widehat{v}, \widehat{y}) \geq 0$ because $\pi(\widehat{v}, \widehat{s}) = \pi(\widehat{v}, \widehat{y})$ and $\widehat{s}$ is not closer than orthogonal to $\widehat{v}$.

Fourth, we claim that the projection q of p onto L_E lies in the interior of E. Assume to the contrary that q does not lie in the interior of E. By the first observation, y lies in the interior of E. So qy intersects an endpoint of E if L_E is a line or qy intersects a weighted-subsegment in ∂E if L_E is a plane. Let γ denote the endpoint/weighted-subsegment that qy intersects. Let $\widehat{x}$ denote $\widehat{\gamma}$ if γ is a vertex or the smallest orthosphere of γ if γ is a weighted-subsegment. Let H be the bisector plane of $\widehat{x}$ and $\widehat{y}$. Consider $\pi(\widehat{v}, \widehat{x})$ for each vertex v of γ. If γ is a vertex, then $\pi(\widehat{v},\widehat{x}) = \pi(\widehat{x},\widehat{x}) = -2X^2 \leq 0$; otherwise, γ is a weighted-subsegment and $\pi(\widehat{v},\widehat{x}) = 0$. On the other hand, by the third observation, $\pi(\widehat{v},\widehat{y}) \geq 0$ for each vertex v of γ. We conclude that for each vertex v of γ, $\pi(v,\widehat{y}) \geq \pi(v,\widehat{x})$. So γ lies in the halfspace H^+ bounded by H where $\pi(a,\widehat{x}) \leq \pi(a,\widehat{y})$ for each point $a \in H^+$. The ray emitting from x through y must shoot outside H^+ because a point sufficiently far in this direction is closer to $\widehat{y}$ than $\widehat{x}$. Coupled with the fact that qy intersects γ, we conclude that $q \in H^+$. Note that pq is parallel to H. So $p \in H^+$ which implies that $\pi(\widehat{p},\widehat{x}) \leq \pi(\widehat{p},\widehat{y})$. By the second observation, $\pi(\widehat{p},\widehat{y}) < 0$ and so $\pi(\widehat{p},\widehat{x}) < 0$. If γ is a vertex, this contradicts the vertex gap property; otherwise, it contradicts the assumption that no weighted-subsegment is encroached. This proves the claim.

Let τ be the weighted-subsegment or weighted-subfacet in the interior of E that contains q. Let $\widehat{z}$ be the smallest orthosphere of τ. Let H_1 be the bisector plane of $\widehat{y}$ and $\widehat{z}$. By the third observation, for each vertex v of τ, $\pi(\widehat{v},\widehat{y}) \geq 0 = \pi(\widehat{v},\widehat{z})$. So τ lies inside the halfspace H_1^+ bounded by H_1 where $\pi(a,\widehat{z}) \leq \pi(a,\widehat{y})$ for each point $a \in H_1^+$. As $q \in \tau$ and pq is parallel to H_1, we have $p \in H^+$. So $\pi(\widehat{p},\widehat{z}) \leq \pi(\widehat{p},\widehat{y})$ which is negative by the second observation. But then $\widehat{p}$ encroaches upon τ, a contradiction. □

3.5 Adjacent elements and encroachment

A key property in ordinary Delaunay refinement (without weight assignment) is that vertices on one element of $\mathcal{P}$ cannot encroach upon subsegments and subfacets on an adjacent element of $\mathcal{P}$, 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 the weighted case.

Lemma 3.3 *If the vertex gap property holds, then for any weighted-subsegment ab on an edge e of $\mathcal{P}$, ab cannot be encroached by any vertex that lies on an edge adjacent to e or a facet adjacent but non-incident to e.*

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

Lemma 3.4 *Let abc be a weighted-subfacet on a facet F of $\mathcal{P}$. If there is no encroached weighted-subsegment, then abc cannot be encroached by any vertex that lies on a facet adjacent to F or an edge adjacent but non-incident to F.*

Proof. Let H be the plane containing F. Let T denote the two-dimensional weighted Delaunay triangulation of the vertices on F. Let V_T denote the subdivision of H induced by T. Let Σ be the set of the smallest orthospheres of the triangles in T, the smallest orthospheres of weighted-subsegments in ∂F, and a sphere centered at infinity in H with infinite radius. By duality, V_T is the intersection of H and the weighted Voronoi diagram of Σ. Let u be a vertex that lies on a facet adjacent to F or an edge adjacent but non-incident to F. Because no input angle is less than $\pi/2$, the orthogonal projection of u onto H falls outside F or on ∂F. Consider any weighted-subfacet abc on F. The directed segment from the projection of u to a intersects a sequence of cells of V_T, including some weighted-subsegment vw in ∂F. This yields a corresponding sequence of spheres in Σ owning the cells intersected. The weighted distances from $\widehat{u}$ to the spheres in the sequence increase along the sequence. It follows that if $\widehat{u}$ encroaches upon abc, $\widehat{u}$ also encroaches upon vw, a contradiction. □

3.6 Projection

Instead of finding any encroached subfacet, we will focus on one that contains the projection of its encroaching vertex. Such a result has been proved by Shewchuk [22] for encroachments by unweighted points in the unweighted Delaunay triangulations. We will need a weighted version of this result.

Lemma 3.5 *If no weighted-subsegment is encroached and $\widehat{p}$ encroaches upon some weighted-subfacet on a facet F, then there exists a weighted-subfacet h on F which is encroached upon by $\widehat{p}$ and h contains the orthogonal projection of p on F.*

Proof. Let H be the plane containing F. Let T denote the two-dimensional weighted Delaunay triangulation of the vertices on F. Let V_T denote the subdivision of H induced by T. Let Σ be the set of the smallest orthospheres of the triangles of T, the smallest orthospheres of the subsegments in ∂F, and a sphere centered at infinity in H with infinite radius. V_T is the intersection of H and the weighted Voronoi diagram of Σ. Let t_1 be a weighted-subfacet of F that is encroached upon by $\widehat{p}$. Suppose that p projects to a cell t_2 in V_T. When we walk along the directed segment from the projection of p (inside t_2) to an interior point of t_1, we encountered a sequence of cells in V_T. This yields a corresponding sequence of spheres in Σ owning the cells encountered. The weighted distances from $\widehat{p}$ to the spheres in this sequence increase along the sequence. 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 weighted-subsegment in ∂F. Because $\widehat{p}$ 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 orthosphere of t_1. It follows that $\widehat{p}$ encroaches upon $t \cap t'$ which is a contradiction. □

3.7 Algorithm

The input to our algorithm QualMesh is a PLC. The PLC bounds a convex domain and the PLC may contain vertices, segments, and facets within the domain. We also assume that no input angle is less than $\pi/2$. This includes all angles between two segments sharing a vertex, a segment and a facet sharing one vertex, or two facets sharing a vertex or a segment. The assumption of a convex bounded domain is not a serious restriction because any PLC can be enclosed within a large enough box whose elements are included in the extended PLC. After the meshing is finished, one can choose to retain the desired tetrahedra. This standard technique has been used before [12, 21].

In the algorithm below we have a refinement step which is done in the unweighted Delaunay triangulation though some of the refinements may be triggered by a weighted point. Subsequent to this refinement, the vertices are pumped to eliminate slivers. Obviously, this is carried out in the weighted Delaunay triangulation. The results proved so far for the weighted Delaunay triangulation also remain valid for the unweighted case (where all weights are assumed to be zero) as well as when only some vertices are weighted.

QUALMESH($\mathcal{P}$)

1. Compute the Delaunay triangulation of the input vertices of $\mathcal{P}$;
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. The parameters ρ_0, σ_0 and ω_0 will be determined later.

   RULE 1(SUBSEGMENT REFINEMENT). If there is an encroached subsegment, insert its midpoint.

   RULE 2(SUBFACET REFINEMENT). If there is an encroached subfacet, there exists an encroached subfacet h that contains the projection of its encroaching vertex on the facet containing h. Insert the circumcenter of h provided that it does not encroach upon any subsegment. Otherwise, reject the circumcenter and apply rule 1 to split an encroached subsegment.

   RULE 3(TETRAHEDRON REFINEMENT). Assume that there is a tetrahedron with radius-edge ratio exceeding ρ_0 and circumcenter z. If z does not encroach upon any subsegment or subfacet, insert z. Otherwise, reject z and perform one of the following actions:

   - If z encroaches upon some subsegment(s), use rule 1 to split one.
   - Otherwise, z encroaches upon some subfacet(s) and use rule 2 to split one that contains the projection of z.

   RULE 4(WEIGHTED ENCROACHMENT). Let $\mathrm{Del}\,\mathcal{V}$ be the Delaunay triangulation of the current vertex set $\mathcal{V}$. Take a vertex v that is incident on a sliver τ (i.e., $\sigma(\tau) \leq \sigma_0$). Let $\widehat{v}$ be the weighted vertex v with weight $\omega_0^2 f(v)^2$.

   - If $\widehat{v}$ encroaches upon some subsegment in $\mathrm{Del}\,\mathcal{V}$ that does not lie on the same segment as v, use rule 1 to split one.
   - If $\widehat{v}$ encroaches upon some subfacet in $\mathrm{Del}\,\mathcal{V}$ that does not lie on the same facet as v, use rule 2 to split one that contains the projection of v.

   Notice that we always maintain an unweighted Delaunay triangulation in step 2 and $\widehat{v}$ is used only for checking encroachments.

3. For each vertex v incident on a sliver τ (i.e., $\sigma(\tau) \leq \sigma_0$), pump v with weight in $[0, \omega_0^2 f(v)^2]$ until no sliver is incident to v. Maintain the weighted Delaunay triangulation during the pumping. We claim that no pumped vertex encroaches upon any weighted-subsegment and weighted-subfacet.

3.8 Time analysis

We analyze the time complexity of the algorithm in terms of n, the number of input vertices, and N, the number of output vertices. We will prove in Section 7.4 that N is no more than a constant factor of the minimum number of vertices possible.

Consider the lifting map $\mu : \mathbb{R}^3 \to \mathbb{R}^4$ which maps a point $x = (x_1, x_2, x_3) \in \mathbb{R}^3$ to a point $\mu(x) = (x_1, x_2, x_3, x_1^2 + x_2^2 + x_3^2) \in \mathbb{R}^4$. For a point set $\mathcal{V}$ in three dimensions, let $\mu(\mathcal{V}) = \{\mu(v), v \in \mathcal{V}\}$. The Delaunay triangulation of $\mathcal{V}$ is the projection of the convex hull of $\mu(\mathcal{V})$ [11]. So the first step can be done in $O(n^2)$ time using Chazelle's convex hull algorithm [4]. The second step is a loop and a new vertex is added in each iteration. So there are less than

N iterations. After each vertex insertion in Rule 1, 2, 3 or 4, we have to update the Delaunay triangulation. An inserted point p is the center of a circumsphere, circumcircle or a segment. In each case we get a tetrahedron which is destroyed after inserting p. We explore in the Delaunay data structure in a depth first manner to collect all tetrahedra that are destroyed with the insertion of p. Once these tetrahedra are identified, p is connected to the boundary of the union of them to update the Delaunay triangulation. If D_p is the number of deleted tetrahedra, the complexity of this update is $O(D_p)$. So, the total time of all updates over the entire algorithm is upper bounded by the number of deleted tetrahedra. We argue that this number is $O(N^2)$.

In the lifted diagram in four dimensions, the insertion of p can be viewed as follows. The point $\mu(p)$ is below the convex hull of $\mu(\mathcal{V})$ and let T be the set of tetrahedra on this convex hull visible to $\mu(p)$. Insertion of $\mu(p)$ destroys all tetrahedra in T and creates new tetrahedra on the updated convex hull by connecting $\mu(p)$ to the boundary of the union of tetrahedra in T. Let us call the union of new tetrahedra incident to $\mu(p)$ as its *cap*. The space between the cap of $\mu(p)$ and T can be triangulated by connecting $\mu(p)$ to each tetrahedron in T. Thus, assuming that the convex hull of the initial point set is triangulated, one can maintain a triangulation in the lifted diagram after each insertion, which contains the lifted deleted tetrahedra. Therefore, all tetrahedra deleted by QUALMESH can be mapped to tetrahedra in the triangulation of N points in four dimensions. Since the size of any triangulation of N points in four dimensions is only $O(N^2)$ (Theorem 1.2 [10]), the same bound applies to the number of deleted tetrahedra.

Each insertion is preceded by a search of an encroached subsegment, subfacet or a skinny tetrahedron. We argue that this search can also be done in $O(N^2)$ total time. We maintain a stack of all skinny tetrahedra. Whenever an update is performed in the triangulation, we update the stack and mark any tetrahedron through pointers in the stack which is deleted. Thus, a skinny tetrahedron can be obtained by popping the stack until the popped tetrahedron is not marked. The time to create the initial stack is $O(n^2)$, the complexity of the initial Delaunay triangulation. The time to update the stack can be absorbed in the triangulation update time which is $O(N^2)$ in total. Next, we need to account for searching the encroached subsegments and subfacets. This encroachement may occur by an inserted or rejected point. Since each rejected point leads to an insertion, the total number of inserted and rejected points is $O(N)$. For each such point we can scan all subfacets and subsegments to determine the encroachments. At any time of the algorithm, the total number of subsegments and subfacets is $O(N)$. This is because the subfacets and subsegments on a planar facet create a plane graph whose complexity is linear in terms of the number of vertices, and each input edge can be incident to only constant number of facets due to the input angle constraint. Therefore, counting over all points, all encroachements can be determined in $O(N^2)$ time.

Other than vertex insertion, we also need to compute $f(v)$ for some vertex v in each application of Rule 4. This can be done in $O(n)$ time by checking all the input elements. So the total time spent in computing local feature sizes is $O(Nn)$. Hence, the total time taken by the second step is $O(N^2 + Nn + n^2) = O(N^2)$. The third step pumps at most N vertices. Each pumping requires updating the mesh connectivity. The complexity of these updates again can be counted through the lifted diagram to be $O(N^2)$. Thus, the third step takes $O(N^2)$ time. In all, we have the following theorem.

Theorem 3.2 *The time complexity of* QUALMESH *is $O(N^2)$ where N is the minimum number of points required to mesh the input domain with tetrahedra that have bounded aspect ratio.*

4 Insertion radii

We define a notion of *insertion radius* for each vertex inserted or rejected by QUALMESH. The main result in this section is a relation among insertion radii and local feature sizes which allows us to prove a lower bound on inter-vertex distances in Section 5. For any vertex x in the input PLC, the insertion radius r_x is the Euclidean distance from the nearest input vertex. For any vertex x that QUALMESH inserts or rejects, the insertion radius r_x is the distance of x from the nearest vertex in the current $\mathcal{V}$. In Section 5, we will see that a lower bound on the insertion radii of vertices implies a lower bound on the inter-vertex distances. So a packing argument shows that QUALMESH must terminate as the domain has bounded volume.

The analysis uses a parent-child relation among all vertices that are input vertices or inserted/rejected by QUALMESH. This is similar to the parent-child relation defined by Shewchuk [22] for the three-dimensional Delaunay refinement, but we need some modifications because of refinement triggered by weighted points. Parents of input vertices are undefined. If QUALMESH 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: Among the two endpoints of the shortest edge of the tetrahedron split by x, the vertex p that appeared in $\mathcal{V}$ the latest is the parent of x.

2. x has type 1 or 2: Then x is the circumcenter of a subsegment or subfacet τ. There are two cases.

   (a) If τ is encroached by a weighted vertex $\widehat{p}$ in rule 4, the parent of x is p.

   (b) If τ is encroached by an unweighted vertex, we choose the parent of x to be the encroaching vertex p nearest to x. If p was not rejected, then $r_x = \|p-x\|$; otherwise, $r_x = X \geq \|p-x\|$.

Our goal is to lower bound r_x in terms of $f(x)$ and r_p, where p is the parent of x. (This recurrence will be useful in an inductive proof to lower bound r_x in terms of $f(x)$ only.) To this end, we prove two technical lemmas. Lemma 4.1 lower bounds $\|p-x\|$ in terms of $f(x)$, $f(p)$, and r_p. Lemma 4.2 lower bounds r_x in terms of $\|p-x\|$. Both apply under some special conditions. Then we employ these two lemmas and analyze the remaining cases to obtain the lower bound on r_x in terms of $f(x)$ and r_p.

Lemma 4.1 *Suppose that the vertex gap property holds. Let x be a vertex of type 1 or 2. Let p be the parent of x. Let $\widehat{x}$ be the smallest circumsphere of the subsegment or subfacet centered at x.*

(i) If p is an input vertex, p has type 1, or p has type 2 and x has type 2, then $\|p-x\| \geq \max\{f(x), f(p)\}$.

(ii) Suppose that p has type 2 and x has type 1, or p has type 3. If p does not lie inside $\widehat{x}$, then $\|p-x\| \geq r_p/\sqrt{2}$.

Proof. Let F be the segment containing x if x has type 1, otherwise let F be the facet containing x. We first claim that p does not lie on F. If the insertion/rejection of x is induced in rule 4, the claim is enforced by the algorithm. Otherwise, because the unweighted Delaunay triangulation of vertices on a segment or facet gives the subsegments and subfacets, an unweighted vertex on a segment (facet) cannot encroach upon subsegments (subfacets) on the same segment (facet).

Case 1: p is an input vertex. The two balls centered at p and x with radius $\|p-x\|$ intersect two non-adjacent input elements (F and p). So $\|p-x\| \geq f(p)$ and $\|p-x\| \geq f(x)$.

Case 2: p has type 1, or p has type 2 and x has type 2. Let F' be the segment/facet containing p. We already know that $F' \neq F$. We argue that F and F' are non-adjacent. If x has type 1 and p has type 1, then Lemma 3.3 implies that F' is non-adjacent from F. Suppose that x has type 2. The insertion/rejection of x could be induced in an application of rule 4 or it could be induced in a direct application of rule 2. In any case, there is no encroached subsegment, otherwise rule 1 would have been invoked instead. So Lemma 3.4 implies that F' and F are non-adjacent. As F and F' are non-adjacent, $\|p-x\| \geq f(x)$ and $\|p-x\| \geq f(p)$.

Case 3: p has type 2 and x has type 1, or p has type 3. (When p has type 2 and x has type 1, p and x may lie on the same facet.) By assumption, p does not lie inside $\widehat{x}$. This implies that the insertion/rejection of x is induced in rule 4. Also, QUALMESH enforces two conditions. First, p was inserted by QUALMESH. Second, if x has type 1 (resp. type 2) and lies on a segment (resp. facet), p does not lie on the same segment (resp. facet). It suffices to lower bound the distance from p to F. Go back to the time when p was inserted by QUALMESH.

Case 3.1: x has type 1. So F is a segment. Let ab be the subsegment on F that contains x at that time. Observe that p lies outside the smallest circumsphere of ab; otherwise, p would have been rejected for encroaching upon a subsegment because p has type 2 or 3. If the nearest point of ab to p is a or b, then $\|p-x\| \geq \min\{\|p-a\|, \|p-b\|\} \geq r_p$ because a and b exist when p is inserted. Otherwise, the nearest point lies in the interior of ab. Assume that the orthogonal projection of p onto ab lies on ay, where y is the midpoint of ab. Then $\sin \angle pay \geq 1/\sqrt{2}$. So $\|p-x\| \geq \|a-p\| \cdot \sin \angle pay \geq r_p/\sqrt{2}$.

Case 3.2: x has type 2. So p has type 3 and F is a facet. Observe that p lies outside the smallest circumsphere of any subsegment in ∂F or any subfacet on F. Otherwise, p would have been rejected for encroaching upon a subsegment or subfacet because p has type 3.

If the nearest point of F to p lies on a segment F' in ∂F, by applying the analysis in Case 3.1 to F', we conclude that $\|p-x\| \geq r_p/\sqrt{2}$. Otherwise, the nearest point q on F to p lies inside some subfacet abc on F. Note that q is the orthogonal projection of p onto abc. The Voronoi diagram of a, b and c divides abc into three regions, each containing a vertex of abc. Assume that q lies inside the region containing a. Let $\widehat{z}$ be the smallest circumsphere of abc. Let H be the plane that is perpendicular to F and passes through a and p. $H \cap \widehat{z}$ is a circle and p lies outside it. Let y be the center of $H \cap \widehat{z}$. Observe that q lies on ay and so $\sin \angle pay \geq 1/\sqrt{2}$. The Euclidean distance from p to F is $\|a-p\| \cdot \sin \angle pay \geq r_p/\sqrt{2}$.

□

Lemma 4.2 *Suppose that the vertex gap property holds. Let x be a vertex of type 1 or 2. Let p be the parent of x. Then $r_x \geq \|p-x\|/\sqrt{2}$.*

Proof. Let $\widehat{x}$ be the circumsphere of which x is the center. If p lies inside $\widehat{x}$, p is unweighted and it follows from the definition of parent that $r_x \geq \|p-x\|$. Assume that p does not lie inside $\widehat{x}$. For

p to be the parent of x, the insertion/rejection of x must be induced in rule 4 by the weighted vertex $\widehat{p}$ with weight $\omega_0^2 f(p)^2$. This also implies that $p \in \mathcal{V}$ at that time. By Lemma 4.1, $\|p-x\| \geq f(p)$ or $\|p-x\| \geq r_p/\sqrt{2}$. If $\|p-x\| \geq f(p)$, then $\|p-x\|/\sqrt{2} \geq \omega_0 f(p)$ as $\omega_0 \leq 1/2$. If $\|p-x\| \geq r_p/\sqrt{2}$, as $r_p \geq 2\omega_0 f(p)$ by the vertex gap property, we also get $\|p-x\|/\sqrt{2} \geq \omega_0 f(p)$. Because $\widehat{p}$ was closer than orthogonal from $\widehat{x}$, we have $X^2 + \omega_0^2 f(p)^2 > \|p-x\|^2$. Substituting $\omega_0^2 f(p)^2$ by $\|p-x\|^2/2$ and rearranging terms, we get $r_x = X > \|p-x\|/\sqrt{2}$. □

We are ready to lower bound r_x in terms of $f(x)$ and the insertion radius of the parent of x. This is the main result of this subsection.

Lemma 4.3 *Suppose that the vertex gap property holds. Let x be an input vertex or a vertex inserted or rejected. Let p be the parent of x, if it exists.*

(i) If x is an input vertex or p is an input vertex, then $r_x \geq f(x)/\sqrt{2}$.

(ii) Otherwise, $r_x \geq f(x)/\sqrt{2}$ or $r_x \geq c \cdot r_p$, where $c = 1/2$ if x has type 1 or 2 and $c = \rho_0$ if x has type 3.

Proof. If x is an input vertex, then $r_x \geq f(x)$ by the definition of local feature size. Suppose that QUALMESH inserts x or rejects x by rule 1, 2, or 3. Recall that x is the center of the smallest circumsphere of a subsegment, subfacet, or tetrahedron. Let $\widehat{x}$ denote this circumsphere.

Consider the case where p is an input vertex. If x has type 1 or 2, then $\|p-x\| \geq f(x)$ by Lemma 4.1. Note that $r_x \geq \|p-x\|/\sqrt{2}$ by Lemma 4.2. So we get $r_x \geq f(x)/\sqrt{2}$. Suppose that x has type 3. Let τ be the skinny tetrahedron split by x. One endpoint of the shortest edge of τ is p. Let q denote the other endpoint. Because p is an input vertex, q is also an input vertex. This is because p did not appear in $\mathcal{V}$ earlier than q by the definition of parent-child relation which can only happen if q is also an input vertex. This means the empty ball centered at x with radius r_x has two input vertices, namely p and q, on its boundary. Therefore, $r_x = \|p-x\| \geq f(x)$.

Consider the case where p is not an input vertex. 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. If p has type 1, or p has type 2 and x has type 2, then $\|p-x\| \geq f(x)$ by Lemma 4.1 and $r_x \geq \|p-x\|/\sqrt{2}$ by Lemma 4.2. So $r_x \geq f(x)/\sqrt{2}$. The remaining cases are that p has type 2 and x has type 1, or p has type 3.

Case 1: p lies inside $\widehat{x}$. So p was rejected by QUALMESH and $r_x = X$. Let τ be the subsegment or subfacet of which $\widehat{x}$ is the smallest circumsphere.

Case 1.1: τ is a subsegment ab. Let a be the vertex of τ nearest to p. So $\angle pxa \leq \pi/2$. It follows that $\|p-a\| \leq \sqrt{2}X = \sqrt{2}r_x$. The vertex a was in $\mathcal{V}$ when p was rejected. So $r_p \leq \|p-a\|$. Hence, $r_x \geq r_p/\sqrt{2}$.

Case 1.2: τ is a subfacet abc. The Voronoi diagram of a, b and c divides abc into three regions, each owned by a vertex of abc. Rule 2 enforces that τ contains the orthogonal projection of p. So we can assume that the projection of p lies in the region, say owned by a. Let H be the plane that is perpendicular to abc and passes through a and p. Let y be the center of the circle $H \cap \widehat{x}$. As the projection of p lies inside the region owned by a, $\angle pya \leq \pi/2$. This implies that $\|p-a\|^2 \leq \|y-p\|^2 + \|y-a\|^2 \leq \|y-p\|^2 + X^2$. As p lies inside $\widehat{x} \cap H$, $\|y-p\| < \text{radius}(H \cap \widehat{x}) \leq X$. It follows that $\|p-a\| \leq \sqrt{2}X = \sqrt{2}r_x$. The vertex a was in $\mathcal{V}$ when p was rejected. So $r_p \leq \|p-a\|$. Hence, $r_x \geq r_p/\sqrt{2}$.

Case 2: p does not lie inside $\widehat{x}$. We have $\|p-x\| \geq r_p/\sqrt{2}$ by Lemma 4.1(ii) and $r_x \geq \|p-x\|/\sqrt{2}$ by Lemma 4.2. So $r_x \geq r_p/2$.

□

5 Vertex-to-vertex distances

In this section, we apply Lemma 4.3 to prove lower bounds on the insertion radii and inter-vertex distances. In the process, we also prove that the vertex gap property holds throughout the algorithm. We will use these results in Section 7 to prove several guarantees provided by weighted Delaunay refinement. We need the following relation involving local feature sizes and insertion radii.

Lemma 5.1 *Let x be a vertex with parent p. If $r_x \geq c \cdot r_p$, then $f(x)/r_x \leq f(p)/(c \cdot r_p) + \sqrt{2}$.*

Proof. Recall that when x is inserted, x is the center of the smallest circumsphere of a subsegment, subfacet, or tetrahedron. Let $\widehat{x}$ denote this circumsphere. If x has type 3, then $r_x = \|p - x\|$. If x has type 1 or type 2, then $r_x \geq \|p - x\|/\sqrt{2}$ by Lemma 4.2. Starting with the Lipschitz condition, we get

$$\begin{aligned} f(x) &\leq f(p) + \|p - x\| \\ &\leq \frac{f(p)}{c \cdot r_p} \cdot r_x + \sqrt{2} r_x \end{aligned}$$

which implies that $f(x)/r_x \leq f(p)/(c \cdot r_p) + \sqrt{2}$. □

The following are the constants of proportionality in Lemma 5.2, the main result in this subsection.

$$C_1 = \frac{7\sqrt{2}\rho_0}{\rho_0 - 4} \qquad C_2 = \frac{3\sqrt{2}\rho_0 + 2\sqrt{2}}{\rho_0 - 4} \qquad C_3 = \frac{\sqrt{2}\rho_0 + 3\sqrt{2}}{\rho_0 - 4} \qquad \omega_0 = \frac{1}{2(1 + C_1)}$$

Note that whenever $\rho_0 > 4$, we have $C_1 > C_2 > C_3 > \sqrt{2}$.

Lemma 5.2 *Let x be a vertex of $\mathcal{P}$ or a vertex inserted or rejected by* QUALMESH. *We have the following invariants for $\rho_0 > 4$.*

(i) If x is a vertex of $\mathcal{P}$ or the parent of x is a vertex of $\mathcal{P}$, then $r_x \geq f(x)/\sqrt{2} > f(x)/C_3$. Otherwise, if x has type i, for $1 \leq i \leq 3$, then $r_x \geq f(x)/C_i$.

(ii) For any other vertex y that appears in $\mathcal{V}$ currently, $\|x-y\| \geq \max\{f(x)/C_1, f(y)/(1+C_1)\}$.

(iii) If x is inserted by QUALMESH, *the vertex gap property holds afterwards.*

Proof. We prove by induction. Invariant (i) holds before QUALMESH starts (the basis case). Clearly, invariant (i) is not affected by pumping a vertex. So it suffices to prove invariant (i) where x is inserted or rejected by QUALMESH. Let p be the parent of x. If p is an input vertex, Lemma 4.3 implies that $r_x \geq f(x)/\sqrt{2}$. Suppose that p is not an input vertex. We assume inductively that invariant (i) holds for p and we conduct a case analysis. If x has type 3, then

$r_x \geq \rho_0 \cdot r_p$ by Lemma 4.3. By induction, we get $f(p) \leq C_1 r_p$ regardless of the type of p. By Lemma 5.1, we get

$$\frac{f(x)}{r_x} \leq \frac{C_1}{\rho_0} + \sqrt{2} = C_3.$$

If x has type 2, the proof of Lemma 4.3 reveals that $r_x \geq f(x)/\sqrt{2} > f(x)/C_2$ when p has type 1 or 2. When p has type 3, case 1 and case 2 in the proof of Lemma 4.3 reveal that $r_x \geq r_p/2$. By induction assumption, $f(p) \leq C_3 r_p$. Then by Lemma 5.1, we get

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

If x has type 1, then the proof of Lemma 4.3 reveals that $r_x \geq f(x)/\sqrt{2} > f(x)/C_1$ when p has type 1. When p has type 2 or 3, case 1 and case 2 in the proof of Lemma 4.3 reveal that $r_x \geq r_p/2$. By induction assumption, $f(p) \leq C_2 r_p$. Then by Lemma 5.1, we get

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

This proves that invariant (i) holds in general.

Consider invariant (ii). For any vertex y that appears in $\mathcal{V}$ currently, $\|x-y\| \geq r_x \geq f(x)/C_1$. Because of $f(x) \geq f(y) - \|x-y\|$, we also get $\|x-y\| \geq f(x)/C_1 \geq f(y)/C_1 - \|x-y\|/C_1$. It follows that $\|x-y\| \geq f(y)/(1+C_1)$.

Consider invariant (iii). It follows from invariant (ii) that for any vertices u and v in $\mathcal{V}$, $\|u-v\| \geq f(u)/(1+C_1)$. By our choice of values of C_1 and ω_0, $f(u)/(1+C_1) = 2\omega_0 f(u)$. This proves that the vertex gap property still holds. □

6 Effect of pumping

Let $\mathcal{V}$ denote a set of unweighted points. Let $\operatorname{Conv}\mathcal{V}$ denote the convex hull of $\mathcal{V}$. Let $\widehat{\mathcal{V}}$ be the weighted points obtained after some weight assignment to points in $\mathcal{V}$. We have already defined $N(x)$ to be the distance to the nearest neighbor in $\mathcal{V}$ for any $x \in \mathcal{V}$. We extend the definition for any $x \in \mathbb{R}^3$ by letting $N(x)$ denote the Euclidean distance from x to its *second* nearest neighbor in $\mathcal{V}$ for any point $x \in \mathbb{R}^3$. If $x \in \mathcal{V}$, then $N(x)$ still denotes the nearest neighbor distance of x. We say $\widehat{\mathcal{V}}$ has *weight property* $[\omega]$ for some $\omega \in (0, 1/2)$ if $U \leq \omega N(u)$ for each $\widehat{u} \in \widehat{\mathcal{V}}$. Let $\operatorname{Del}\widehat{\mathcal{V}}$ denote the weighted Delaunay triangulation of $\widehat{\mathcal{V}}$. $\operatorname{Del}\widehat{\mathcal{V}}$ has *ratio property* $[\rho]$ if the orthoradius-edge ratio of every tetrahedron in $\operatorname{Del}\widehat{\mathcal{V}}$ is at most ρ.

The work of Cheng et al. [8] suggests that $\operatorname{Del}\mathcal{V}$ and $\operatorname{Del}\widehat{\mathcal{V}}$ behave similarly given the ratio and weight properties:

Lemma 6.1 (Claim 7 in [8]) *Let $\mathcal{V}$ be a periodic point set. If* $\operatorname{Del}\mathcal{V}$ *has ratio property* $[\rho]$ *and* $\widehat{\mathcal{V}}$ *has weight property* $[\omega]$*, then* $\operatorname{Del}\widehat{\mathcal{V}}$ *has ratio property* $[\rho']$ *for some* ρ' *depending on* ρ *and* ω.

In this section, we prove a version of the Lemma 6.1 to deal with a finite point set (see Lemma 6.6). There are two differences in the result. First, $\mathcal{V}$ is a finite point set instead of a periodic point set. Second, we need an extra condition that the orthocenter of each tetrahedron in $\operatorname{Del}\widehat{\mathcal{V}}$ lies inside $\operatorname{Conv}\mathcal{V}$. Then the rest of Lemma 6.1 carries over.

We need three results from Talmor's thesis [18]. We state them below and include the proof of Lemma 6.3 as we need an inequality in the proof later. For a point $p \in \mathcal{V}$, let $V_p(\mathcal{V})$ denote the Voronoi cell owned by p in the Voronoi diagram of $\mathcal{V}$. Let b_p and B_p be balls centered at p such that $\text{radius}(b_p) = N(p)/2$ and $\text{radius}(B_p) = \rho L \cdot N(p)$ where L is a constant used in the next lemma.

Lemma 6.2 (Lemma 3.4.3 in [18]) *If* $\text{Del}\,\mathcal{V}$ *has ratio property* $[\rho]$*, the lengths of two adjacent edges of* $\text{Del}\,\mathcal{V}$ *differ by at most some constant factor* L *depending on* ρ.

Lemma 6.3 (Lemma 3.5.1 in [18]) *Assume that* $\text{Del}\,\mathcal{V}$ *has ratio property* $[\rho]$*. For each* $p \in \mathcal{V}$*,* $V_p(\mathcal{V})$ *contains* b_p *and* B_p *contains all vertices of* $V_p(\mathcal{V})$.

Proof. It is obvious that $b_p \subseteq V_p(\mathcal{V})$ as $\text{radius}(b_p) = N(p)/2$, Let $v \in \mathcal{V}$ such that $\|p - v\| = N(p)$. So pv is an edge of $\text{Del}\,\mathcal{V}$. Let τ be some tetrahedron in $\text{Del}\,\mathcal{V}$ incident to p. Let pq be an edge of τ. Using Lemma 6.2, we get

$$\|p - q\| \leq L \cdot \|p - v\| = L \cdot N(p). \tag{6.1}$$

Let z be the circumcenter of τ, i.e., z is a vertex of $V_p(\mathcal{V})$. The ratio property implies that

$$\|p - z\| \leq \rho \cdot \|p - q\| \leq \rho L \cdot N(p) = \text{radius}(B_p).$$

Thus, B_p contains all vertices of $V_p(\mathcal{V})$. □

Lemma 6.4 (Theorem 3.6.2 in [18]) *Assume that* $\text{Del}\,\mathcal{V}$ *has ratio property* $[\rho]$*. Let* xz *be a line segment lying inside* $\bigcup_{p \in \mathcal{V}} V_p(\mathcal{V}) \cap B_p$*. Let* $\widehat{z}$ *be the sphere centered at* z *with radius* $\|x - z\|$*. Then there is a constant* $C > 0$ *such that if* $\widehat{z}$ *is empty, then* $N(z) \leq C \cdot N(x)$.

Next, we apply Lemma 6.3 to show that B_p is so large that $V_p(\mathcal{V}) \cap B_p$ contains $V_p(\mathcal{V}) \cap \text{Conv}\,\mathcal{V}$. The implication is that $\bigcup_{p \in \mathcal{V}} V_p(\mathcal{V}) \cap B_p$ contains the Voronoi diagram of $\mathcal{V}$ clipped within $\text{Conv}\,\mathcal{V}$.

Lemma 6.5 *Assume that* $\text{Del}\,\mathcal{V}$ *has ratio property* $[\rho]$*. For each* $p \in \mathcal{V}$*,* $V_p(\mathcal{V}) \cap \text{Conv}\,\mathcal{V} \subseteq V_p(\mathcal{V}) \cap B_p$.

Proof. We first prove that $V_p(\mathcal{V}) \cap \text{Conv}\,\mathcal{V} \subseteq B_p$. If $V_p(\mathcal{V})$ is bounded, then $V_p(\mathcal{V}) \subseteq B_p$ by Lemma 6.3. It follows that $V_p(\mathcal{V}) \cap \text{Conv}\,\mathcal{V} \subseteq B_p$. Suppose that $V_p(\mathcal{V})$ is unbounded. Then p is an extreme vertex of $\text{Conv}\,\mathcal{V}$. Let T be the set of boundary triangles of $\text{Conv}\,\mathcal{V}$ incident to p. For each triangle $t \in T$, let H_t denote the supporting plane of t and let H_t^+ denote the halfspace bounded by H_t that contains $\text{Conv}\,\mathcal{V}$. By convexity, $V_p(\mathcal{V}) \cap \text{Conv}\,\mathcal{V} \subseteq V_p(\mathcal{V}) \cap \bigcap_{t \in T} H_t^+$. We show that B_p contains the vertices of $V_p(\mathcal{V}) \cap \bigcap_{t \in T} H_t^+$. A vertex z of $V_p(\mathcal{V}) \cap \bigcap_{t \in T} H_t^+$ has one of three types:

1. z is a vertex of $V_p(\mathcal{V})$. By Lemma 6.3, $z \in B_p$.

2. z is the intersection of some edge pq in $\text{Del}\,\mathcal{V}$ with a facet of $V_p(\mathcal{V})$. By (6.1), $\|p - q\| \leq L \cdot N(p) \leq \text{radius}(B_p)$, so $z \in B_p$.

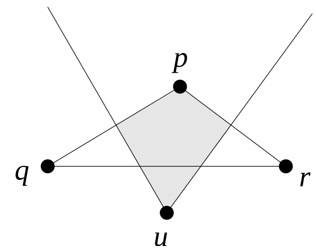

Figure 4: The shaded convex quadrilateral is $\mathcal{Q}$.

3. z is the intersection of H_t for some $t \in T$ and some edge of $V_p(\mathcal{V})$. Let q and r be the other two vertices of t. Let $\mathcal{Q}$ be the convex quadrilateral on H_t bounded by the bisector plane of p and q, the bisector plane of p and r, pq, and pr. See Figure 4. Let u be the vertex of $\mathcal{Q}$ diagonally opposite p. Observe that z lies inside $\mathcal{Q}$. So we are done if we can show that $\mathcal{Q} \subseteq B_p$. By (6.1), $\|p-q\| \leq \text{radius}(B_p)$ and $\|p-r\| \leq \text{radius}(B_p)$. It follows that p, q, and r lie inside B_p. Let $\theta = \angle pqr$ and $\beta = \angle prq$. The angle of $\mathcal{Q}$ at u is $\theta + \beta$. After splitting this angle into two with the diagonal pu, let γ denote the one on the same side as pq. Without loss of generality, assume that $\gamma \geq (\theta+\beta)/2$ which is at least $\min\{\theta, \beta\}$. By the ratio property, $\rho \geq \min\{1/(2\sin\theta), 1/(2\sin\beta)\}$. It follows that $\rho \geq 1/(2\sin\gamma)$. We have

$$\begin{aligned} \|p-u\| &= \frac{\|p-q\|}{2\cdot\sin\gamma} \\ &\leq \rho\cdot\|p-q\| \\ &\overset{(6.1)}{\leq} \rho L\cdot N(p) \\ &= \text{radius}(B_p) \end{aligned}$$

Therefore, $\mathcal{Q} \subseteq B_p$ and hence z lies inside B_p.

Note that $V_p(\mathcal{V}) \cap \bigcap_{t\in T} H_t^+$ is bounded. So we conclude that $V_p(\mathcal{V}) \cap \operatorname{Conv}\mathcal{V} \subseteq V_p(\mathcal{V}) \cap \bigcap_{t\in T} H_t^+ \subseteq B_p$. Finally, $V_p(\mathcal{V}) \cap \operatorname{Conv}\mathcal{V} = V_p(\mathcal{V}) \cap V_p(\mathcal{V}) \cap \operatorname{Conv}\mathcal{V} \subseteq V_p(\mathcal{V}) \cap B_p$. □

Finally, we apply Lemmas 6.4 and 6.5 to prove the main result of this subsection.

Lemma 6.6 *Let $\mathcal{V}$ be a finite point set. Assume that* $\operatorname{Del}\mathcal{V}$ *has ratio property* $[\rho]$*,* $\widehat{\mathcal{V}}$ *has weight property* $[\omega]$*, and the orthocenter of each tetrahedron in* $\operatorname{Del}\widehat{\mathcal{V}}$ *lies inside* $\operatorname{Conv}\mathcal{V}$*. Then* $\operatorname{Del}\widehat{\mathcal{V}}$ *has ratio property* $[\rho']$ *for some constant* ρ' *depending on* ρ *and* ω*.*

Proof. Let $\widehat{z}$ be the orthosphere of a tetrahedron τ in $\operatorname{Del}\widehat{\mathcal{V}}$. That is, z is the orthocenter of τ. Let qr be the shortest edge of τ. Let x be the intersection point $qz \cap \widehat{z}$.

By assumption, z lies inside $\operatorname{Conv}\mathcal{V}$. So we have $xz \subseteq \operatorname{Conv}\mathcal{V}$ by convexity. Using the fact that $\bigcup_{p\in\mathcal{V}} V_p(\mathcal{V}) = \mathbb{R}^3$, we get $xz \subseteq \operatorname{Conv}\mathcal{V} \cap \bigcup_{p\in\mathcal{V}} V_p(\mathcal{V}) = \bigcup_{p\in\mathcal{V}} V_p(\mathcal{V}) \cap \operatorname{Conv}\mathcal{V}$. Lemma 6.5 further implies that $xz \subseteq \bigcup_{p\in\mathcal{V}} V_p(\mathcal{V}) \cap B_p$. As $\widehat{z}$ is empty, we can apply Lemma 6.4 to xz and $\mathcal{V}$. We get

$$Z \leq N(z) \leq C\cdot N(x). \tag{6.2}$$

By the Lipschitz property, $N(x) \leq N(q) + \|q-x\|$. Because $\widehat{q}$ and $\widehat{z}$ intersect and $x = qz \cap \widehat{z}$, x lies inside $\widehat{q}$. By the weight property, the radius of $\widehat{q}$ is at most $\omega N(q)$ and so $\|q-x\| \leq \omega N(q)$.

Thus, we have $N(x) \le (1+\omega)N(q)$. As q and r are vertices in $\mathcal{V}$, $N(q) \le \|q-r\|$. It follows that

$$N(x) \le (1+\omega)\cdot\|q-r\|. \tag{6.3}$$

Substituting (6.3) into (6.2), we get $Z \le C\cdot(1+\omega)\cdot\|q-r\|$. Recall that qr is the shortest edge of τ. Hence, ρ' can be set to be $C\cdot(1+\omega)$. □

7 Guarantees

We establish four guarantees for QUALMESH. The first two guarantees, termination and gradedness, follow from Lemma 5.2. The absence of sliver and the size optimality are guaranteed using Lemma 5.2, Lemma 6.6, and some other results to be proved.

7.1 Termination and gradedness

Theorem 7.1 QUALMESH *terminates with a graded mesh.*

Proof. It follows from Lemma 5.2(ii) that any two vertices u and v at any stage of the algorithm must satisfy $\|u-v\| \ge f(u)/(1+C_1) \ge f_{min}/(1+C_1)$ where f_{min} is the minimum local feature size in the domain. If we center disjoint balls of radii $f_{min}/(2+2C_1)$ at the mesh vertices, then we can pack at most $24(1+C_1)^3\text{volume}(D)/(4\pi f_{min}^3)$ such balls inside a bounded domain D. So the algorithm must terminate. Gradedness follows from the vertex gap property. □

7.2 Conformity

After pumping in step 3, is $\text{Del}\,\widehat{\mathcal{V}}$ conforming? The vertices in $\widehat{\mathcal{V}}$ partition the input segments into weighted-subsegments. For any input facet F, the two-dimensional weighted Delaunay triangulation of vertices on F partition F into weighted-subfacets. We show that $\text{Del}\,\widehat{\mathcal{V}}$ is conforming by showing that no weighted-subsegment or weighted-subfacet is encroached.

Theorem 7.2 *No weighted-subsegment or weighted-subfacet is encroached upon the completion of* QUALMESH.

Proof. Assume to the contrary that a vertex $\widehat{v}$ in $\text{Del}\,\widehat{\mathcal{V}}$ encroaches upon a weighted-subsegment or weighted-subfacet τ.

Consider the case where τ is a weighted-subsegment. Observe that the smallest circumsphere C of τ encloses the smallest orthosphere O of τ. Thus, the bisector H of C and O avoids C and so H avoids τ too. τ lies in the halfspace H^+ bounded by H such that $\pi(p,C) \le \pi(p,O)$ for all point $p \in H^+$. Observe that in order that $\widehat{v}$ encroaches up on τ, τ contains the orthogonal projection of v onto the segment containing τ. Thus, $v \in H^+$ and so $\pi(\widehat{v},C) \le \pi(\widehat{v},O) < 0$. The weight of $\widehat{v}$ is at most the weight used for v in rule 4. However, this contradicts the fact that, v with that weight, did not encroach upon τ.

Consider the case where τ is a weighted-subfacet on a facet F. By Lemma 3.5, we can assume that τ contains the projection of v. Let O be the smallest orthosphere of τ. The two-dimensional unweighted Delaunay triangulation of vertices in $\mathcal{V}$ on F partitions F into subfacets. Let τ' be the subfacet that contains the projection of v on τ. Let C be the smallest circumsphere of τ'. We claim that the bisector plane H of C and O avoids τ'. Otherwise,

H intersects C which implies that C and O intersect (H passes through their intersection). Because τ' has vertices on both sides of H, a vertex of τ' must lie inside O, a contradiction. Observe that τ' lies in the halfspace H^+ bounded by H such that $\pi(p, C) \leq \pi(p, O)$ for all point $p \in H^+$. Thus, $\pi(\widehat{v}, C) \leq \pi(\widehat{v}, O) < 0$. The weight of $\widehat{v}$ is at most the weight used for v in rule 4. However, this contradicts the fact that, v with that weight, did not encroach upon τ'. $\square$

7.3 No sliver

Slivers incident to an unweighted vertex p are eliminated by pumping p. In sliver exudation, p is pumped within the weight interval $[0, \omega_0^2 N(p)^2]$. During the pumping, tetrahedra incident to p change at discrete instances. For pumping to work, two conditions (Lemmas 7.1 and 7.2) must hold over the entire interval of pumping. First, the lengths of edges incident to p are within a constant factor. Second, only a constant number of tetrahedra can be incident to p. These two conditions are proved by Cheng et al. [8] for periodic point set using Lemma 6.1. Lemma 6.6 is the analog of Lemma 6.1 for finite point set. By Theorem 7.2 and Lemma 3.2(iii), the conditions of Lemma 6.6 are satisfied. So we can prove the two conditions for finite point set using exactly the same proofs in [8]. For completeness, we sketch the proofs below. Let $K(\mathcal{V})$ be the graph consisting of edges in $\mathrm{Del}\,\widehat{\mathcal{V}}$ for all $\widehat{\mathcal{V}}$ with weight property $[\omega_0]$.

Lemma 7.1 (Claim 10 in [8]) *Assume that* $\mathrm{Del}\,\mathcal{V}$ *has ratio property* $[\rho_0]$. *The lengths of any two adjacent edges in* $K(\mathcal{V})$ *is within a constant factor* $\nu_0 \geq 1$ *depending on* ρ_0 *and* ω_0.

Proof. (Sketch) Let p be a vertex in $\mathcal{V}$. First, consider a triangle $pqu \in \mathrm{Del}\,\widehat{\mathcal{V}}$ for some $\widehat{\mathcal{V}}$ with weight property $[\omega_0]$. Let Z be the radius of the smallest orthosphere of pqu. By Lemma 6.6, $Z \leq \rho' \cdot \|p - q\|$ and $Z \leq \rho' \cdot \|p - u\|$. On other hand, by the weight property, a constant fraction of pq lies outside $\widehat{p}$ and $\widehat{q}$. This fraction of pq lies inside the smallest orthosphere of pqu and so $Z = \Omega(\|p - q\|)$. Similarly, $Z = \Omega(\|p - u\|)$. It follows that $\|p - q\|$ and $\|p - u\|$ differ by at most some constant factor k_1.

Second, consider the case where pq and pu are two edges in $K(\mathcal{V})$ such that $\angle qpu$ is less than some constant angle bound η. Assume that $\widehat{\mathcal{V}_1}$ and $\widehat{\mathcal{V}_2}$ denote the two weighted versions of $\mathcal{V}$ such that $pq \in \mathrm{Del}\,\widehat{\mathcal{V}_1}$ and $pu \in \mathrm{Del}\,\widehat{\mathcal{V}_2}$. Let H be the plane passing through pqu. We pick the orthosphere $\widehat{z}$ of a tetrahedron τ in $\mathrm{Del}\,\widehat{\mathcal{V}_1}$ that is incident on pq. We intersect $\widehat{z}$ with H to obtain a circle $\widehat{y}$ centered at y with radius Y. Note that $\widehat{y}$ is orthogonal to the circles $\widehat{p} \cap H$ and $\widehat{q} \cap H$. By Lemma 6.6, the radius of $\widehat{z}$ is at most $\rho' \cdot \|p - q\|$. Clearly, Y is at most the radius of $\widehat{z}$. So $Y \leq \rho' \cdot \|p - q\|$. This implies that pq cuts deeply into $\widehat{y}$. As u lies outside $\widehat{y}$, for sufficiently small η (dependent on ρ' and ω_0), pu cannot be much shorter than pq. By symmetry, pq cannot be much shorter than pu. Thus, $\|p - q\|$ and $\|p - u\|$ differ by at most some constant factor k_2.

Next, we deal with all incident edges of p. Let S be a unit sphere centered at p. We take a maximal packing of disjoint spherical caps with angular radii $\eta/4$ on S. The number of such caps is a constant m dependent on η. Then we expand the angular radius of each cap to $\eta/2$. The expanded caps cover S. Each incident edge pq projects radially to a vertex q' on S. Each triangle pqr in $\mathrm{Del}\,\widehat{\mathcal{V}}$ for some $\widehat{\mathcal{V}}$ with weight property $[\omega_0]$ projects radially to an arc $q'r'$ on S. This yields a connected graph embedded on S. Suppose that we walk from q' to an arbitrary vertex u' within the graph. If the walk stays within a cap, by our second observation, the edge length increases by at most a factor of k_2. If the walk enters a new cap, by our first observation, the edge length increases by at most a factor of k_1. If the walk returns to a cap visited before, the whole detour increases the edge length by at most a factor of k_2. As there are m caps, we conclude that $\|p-q\|$ and $\|p-u\|$ differ by at most a factor of $k_2^m k_1^{m-1}$. □

Lemma 7.2 (Claim 11 in [8]) *Assume that* $\mathrm{Del}\,\mathcal{V}$ *has ratio property* $[\rho_0]$. *The degree of every vertex in* $K(\mathcal{V})$ *is bounded by some constant* δ_0 *depending on* ρ_0 *and* ω_0.

Proof. (Sketch) Let p be a vertex in $\mathcal{V}$. Let L be the length of the longest incident edge of p in $K(\mathcal{V})$. Let r be a neighbor of p in $K(\mathcal{V})$. By Lemma 7.1, $\|p-r\| \geq L/\nu_0$. The nearest neighbor of r is a Delaunay neighbor of r. So by Lemma 7.1 again, the nearest neighbor distance of r is at least L/ν_0^2. It follows that at the neighbors of p, we can center disjoint balls of radius $L/(2\nu_0^2)$. Observe that all these balls lie inside the ball centered at p with radius $L + L/(2\nu_0^2)$. Thus, a packing argument shows that p has $O(1)$ neighbors. □

Cheng et al.[8] proved that a sliver incident to p can remain weighted Delaunay only within a subinterval of width $O(\sigma_0 N(p)^2)$ during pumping. As the number of tetrahedra in $K(\mathcal{V})$ incident to p is bounded by a constant, if we choose σ_0 properly, the intervals over which slivers remain incident to a vertex p can be made small enough so that there is a subinterval over which no sliver is incident to p. This is the key idea in [8] in the proof that pumping can eliminate slivers for periodic point sets. The freedom in choosing σ_0 reveals that it is unnecessary to use exactly the weight interval $[0, \omega_0^2 N(p)^2]$. It is equally good that the weight interval contains $[0, \overline{\omega}_0^2 N(p)^2]$ for some constant $\overline{\omega}_0 \leq \omega_0$.

We would like to employ Lemmas 7.1 and 7.2 with $\mathcal{V}$ being a finite point set. To this end, we first need to guarantee that the pumping in step 3 of QUALMESH uses a weight interval $[0, \overline{\omega}_0^2 N(p)^2]$ and that the resulting weighted point set $\widehat{\mathcal{V}}$ has weight property $[\omega_0]$ for some constants $\overline{\omega}_0 \leq \omega_0$. The weight property $[\omega_0]$ follows from the vertex gap property. Lemma 7.3 tells us how to set $\overline{\omega}_0$.

Lemma 7.3 *Let* $\mathcal{M}$ *be the mesh obtained at the end of step 2 of* QUALMESH. *For any vertex* v *in* $\mathcal{M}$, *its nearest neighbor distance is at most* $2\sqrt{2}f(v)$.

Proof. Let B_v denote the ball centered at v of radius $N(v)/2$. Assume to the contrary that $2\sqrt{2}f(v) < N(v)$. Then B_v intersects two disjoint elements of $\mathcal{P}$. B_v cannot contain any vertex in $\mathcal{V}$. So B_v intersects the interior of subsegments or subfacets. Let E be the nearest subsegment or subfacet to v that B_v intersects. Let $\tilde{v}$ be the orthogonal projection of v onto the affine hull of E. Note that $\tilde{v}$ lies on E. The Voronoi diagram of the vertices of E partition E into regions, each owned by one vertex of E. Assume that $\tilde{v}$ lies in the region owned by the vertex w of E. Because $\|v-w\| \geq \mathrm{radius}(B_v) > \sqrt{2}f(v)$ and $\|v-\tilde{v}\| \leq f(v)$, we have

$$\|v-\tilde{v}\| < \|v-w\|/\sqrt{2}. \tag{7.4}$$

Let $\widehat{x}$ be the smallest circumsphere of E. We have $\|v-x\|^2 = \|\tilde{v}-x\|^2 + \|v-\tilde{v}\|^2$. Because $\tilde{v}$ lies in the region owned by w, $\angle x\tilde{v}w \geq \pi/2$ which implies that $\|\tilde{v}-x\|^2 \leq \|w-x\|^2 - \|\tilde{v}-w\|^2$. Therefore,

$$\begin{aligned}\|v-x\|^2 &\leq \|w-x\|^2 - \|\tilde{v}-w\|^2 + \|v-\tilde{v}\|^2 \\ &= \|w-x\|^2 - \|v-w\|^2 + 2\|v-\tilde{v}\|^2 \\ &\overset{(7.4)}{<} \|w-x\|^2\end{aligned}$$

But then v lies inside $\widehat{x}$ and so v encroaches upon E, a contradiction. $\square$

In step 3 of QUALMESH, we pump p using the weight interval $[0, \omega_0^2 f(p)^2]$. This interval contains $[0, \omega_0^2 N(p)^2/8]$ by Lemma 7.3. Using Lemmas 7.1 and 7.2, the same proof in [8] shows that pumping eliminates slivers for the finite point set $\mathcal{V}$. For completeness, we sketch the proof below. Algorithmically, we can use flips to generate new tetrahedra as pumping progresses and stop when no sliver is incident to p.

Theorem 7.3 *There is a constant $\sigma_0 > 0$ such that $\sigma(\tau) > \sigma_0$ for every tetrahedron τ in the output mesh of* QUALMESH.

Proof. (Sketch) Let $pqrs$ be a sliver in some $\widehat{\mathcal{V}}$ with weight property $[\omega_0]$. We are to analyze what happens to $pqrs$ when p is pumped with weight from the interval $[0, \omega_0^2 f(p)^2]$. Let W_{qrs} be the subinterval such that $pqrs$ may remain weighted Delaunay when $P^2 \in W_{qrs}$.

We claim that $|W_{qrs}| = O(\sigma_0 N(p)^2)$. Let L be the shortest edge length of $pqrs$. Let $\widehat{z}$ be the orthosphere of $pqrs$. Let $H(P)$ be the signed distance of z from the plane passing through qrs when p has weight P^2. $H(P)$ is positive if p and z lie on the same side and $H(P)$ is negative otherwise. Observe that $Z^2 = H(P)^2 + Y^2$, where Y is the radius of smallest orthosphere of qrs. By Lemma 6.6, $Z \leq \rho' L$. The circumradius of qrs is at least $L/2$. Also, by Claim 4 in [8], the circumradius of qrs is at most $Y/\sqrt{1 - 4\omega_0^2}$. It follows that $H(P)^2 = Z^2 - Y^2 = O(L^2)$ or $H(P) \in [-kL, kL]$ for some constant k. By Claim 13 in [8], $H(P) = H(0) - P^2/(2D)$, where D is the distance of p from the plane passing through qrs. Substituting into $H(P) \in [-kL, kL]$ and rearranging terms, we get $2D \cdot H(0) - 2kDL \leq P^2 \leq 2D \cdot H(0) + 2kDL$. Thus, $|W_{qrs}| \leq 4kDL$. Note that volume$(pqrs) = \Theta(L^2 D)$ and volume$(pqrs)/L^3 \leq \sigma_0$ as $pqrs$ is a sliver. This implies that $D = O(\sigma_0 L)$ and so $|W_{qrs}| = O(\sigma_0 L^2)$. The nearest neighbor of p is Delaunay neighbor of p. So by applying Lemma 7.1 to p and then to q, r and s, we conclude that $L = \Theta(N(p))$. Hence, $|W_{qrs}| = O(\sigma_0 N(p)^2)$.

Finally, by Lemma 7.2, there are at most δ_0^3 slivers incident to p throughout the entire pumping. So there are at most δ_0^3 forbidden subintervals. Their total length is at most $k'\sigma_0\delta_0^3 N(p)^2$ for some constant k'. By Lemma 7.3, the weight interval $[0, \omega_0^2 f(p)^2]$ contains $[0, \omega_0^2 N(p)^2/8]$. It follows that if $\sigma_0 < \omega_0^2/(8k'\delta_0^3)$, p can be assigned a weight within $[0, \omega_0^2 f(p)^2]$ such that p is not incident to any sliver. $\square$

7.4 Size optimality

We prove that the size of our Delaunay mesh is within a constant factor of the size of any mesh that has bounded aspect ratio. Our proof is a combination of ideas in Ruppert's proof for the two-dimensional case [21] and ideas in Mitchell and Vavasis's proof for their octtree algorithm in higher dimensions [20].

Let T be a triangulation of the input domain that conforms to $\mathcal{P}$ and has 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. For a vector x, we use $\|x\|$ to denote its L_2-norm. For a square matrix A, we use $\|A\|$ to denote its spectral norm, i.e., the square root of the maximum eigenvalue of $A^t A$.

Lemma 7.4 (Lemma 2 in [20]) *For any tetrahedron τ in T, $\|(M_\tau^{-1})^t\| = \|M_\tau^{-1}\| \le k_1/h(\tau)$ for some constant $k_1 \ge 1$.*

Proof. It is proved in [20] that $\|M_\tau^{-1}\| \le k_1/h(\tau)$ for some constant k_1. $\|A\| = \|A^t\|$ for any square matrix A. □

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. The result is analogous to Theorem 2 in [20].

Lemma 7.5 *Let p and q be two points in the interior of T. Let τ_p and τ_q be the tetrahedra of T containing p and q respectively. If τ_p and τ_q do not share any vertex, then there is a tetrahedron τ in T intersecting pq such that*

(i) $h(\tau) \le k_2\|p-q\|$ for some constant $k_2 \ge 1$.

(ii) τ shares a vertex with τ_q. (τ can also be forced to share a vertex with τ_p instead, but τ cannot be guaranteed to share vertices with both τ_p and τ_q.)

Proof. Define a simplicial map $\psi : T \to \mathbb{R}$ by setting $\psi(v) = 1$ for each vertex v of τ_q and $\psi(w) = 0$ for all other vertices w. It follows that $\psi(x) = 1$ for all point $x \in \tau_q$ and $\psi(x) = 0$ for all point $x \in \tau_p$. As 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,

$$\| \bigtriangledown \psi\| \ge \frac{1}{\|p-q\|} \tag{7.5}$$

on τ. This implies that τ shares a vertex with τ_q, otherwise ψ and $\bigtriangledown\psi$ would be identically zero on τ which is a contradiction. This proves (ii). 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 $0 \le |r_i| \le 1$.

Claim 7.1 *For any point z in τ, $\psi(z) - \psi(v_0) = (r_1, r_2, r_3)M_\tau^{-1}(z - v_0)$.*

Proof. 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 $M_\tau = (v_1 - v_0, v_2 - v_0, v_3 - v_0)$. As $M_\tau^{-1} M_\tau = I$, $\alpha_k \cdot (v_k - v_0) = 1$ and $\alpha_k \cdot (v_j - v_0) = 0$ for all $j \ne k$. It follows that

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

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

The claim implies that $\nabla\psi = (M_\tau^{-1})^t(r_1, r_2, r_3)^t$ on τ and so

$$\begin{aligned} \| \nabla \psi\| &\leq \|(M_\tau^{-1})^t\| \cdot \|(r_1, r_2, r_3)^t\| \\ &\leq \sqrt{3} \cdot \|(M_\tau^{-1})^t\|, \quad \text{as } 0 \leq |r_i| \leq 1 \\ &\overset{\text{Lemma 7.4}}{\leq} \sqrt{3}k_1/h(\tau). \end{aligned}$$

Combining the above with (7.5), we obtain $h(\tau) \leq \sqrt{3}k_1\|p - q\|$. This proves (i). □

As T has bounded aspect ratio, it enjoys properties similar to that stated in Lemma 6.2. In particular, two tetrahedra sharing a vertex have similar minimum heights.

Lemma 7.6 ([19, 20]) *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 $k_3 \geq 1$.*

Next, we prove that the minimum heights of tetrahedra in T change fairly smoothly. The result is analogous to Lemma 11 in [20].

Lemma 7.7 *Let p and q be two points and let τ_p and τ_q be the two tetrahedra in T that contain them respectively. Then $h(\tau_q) \leq k_4 \max\{h(\tau_p), \|p - q\|\}$ for some constant $k_4 \geq 1$.*

Proof. If τ_q shares a vertex with τ_p, then $h(\tau_q) \leq k_3 h(\tau_p)$ by Lemma 7.6. Consider the case where τ_q does not share a vertex with τ_p. By Lemma 7.5, there is a tetrahedron τ in T intersecting pq such that $h(\tau) \leq k_2\|p - q\|$ and τ shares a vertex with τ_q. Starting with Lemma 7.6, we get $h(\tau_q) \leq k_3 h(\tau) \leq k_2 k_3\|p - q\|$. □

The following lemma shows that the minimum heights of tetrahedra in T are also related to the local feature sizes. The result is analogous to Lemma 5 in [21].

Lemma 7.8 *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 $k_5 \geq 1$.*

Proof. Let B be the ball centered at x of radius $f(x)$. B contains two points p and q on two disjoint elements of $\mathcal{P}$. By Lemma 7.5, 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 7.7 to u and x, we have

$$\begin{aligned} h(\tau) &\leq k_4 \max\{h(\tau'), \|u - x\|\} \\ &\leq k_4 \max\{k_2\|p - q\|, \|u - x\|\} \\ &\leq k_4 \max\{k_2 \cdot 2f(x), f(x)\} \\ &\leq 2k_2k_4 f(x) \end{aligned}$$

□

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

Theorem 7.4 *The output size of* QualMesh *is within a constant factor of the size of any mesh of bounded aspect ratio for the same domain.*

Proof. Let $\mathcal{D}$ 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 $\int_{\mathcal{D}} \frac{dx}{f(x)^3}$. As $N(v) \geq f(v)/(1+C_1)$, 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. Observe that $f(x) \leq f(v) + d_v$ for all $x \in B_v$. Therefore,

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

Therefore, we have

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

By Lemma 7.2, 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

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

for some constant k_6.

Next, consider a mesh T of $\mathcal{D}$ of bounded aspect ratio. For each point x in $\mathcal{D}$, define $\ell(x)$ to be the minimum height of the tetrahedra in T containing x. By Lemma 7.8, $\ell(x) \leq k_5 f(x)$ and so

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

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

8 Conclusions

A series of developments starting with Chew [5] and Ruppert [21], and continuing with Shewchuk [22] and Cheng et al. [8] brought the difficult problem of quality three-dimensional Delaunay meshing of bounded domains close to the solution. Li and Teng [16] recently developed a randomized point-placement strategy to generate a provably good three-dimensional Delaunay mesh of bounded domains. This paper introduces a new paradigm, weighted Delaunay refinement, which gives the first deterministic algorithm for the problem. We believe that we will add fewer points in practice because weighted Delaunay refinement uses pumping to eliminate slivers instead of point-placement.

Of course, as with previous algorithms the constants derived for the theory are miserably unsatisfactory for all practical purposes. For example, the constant $\rho_0 > 4$ 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 [14]. Will the same remain true when we combine the two into our weighted Delaunay refinement algorithm?

In QUALMESH we need to compute the local feature size $f(v)$ while assigning weight to a vertex v. Although the computation of $f(v)$ is feasible, it is better if we can avoid computing it in practice. Towards this end, one can gradually increase the weight to v and check the quality of tetrahedra incident to v as they change at discrete moments. Experiments should be performed to see what bound on angles do we get in practice with this strategy. We plan future experiments to answer these questions.

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 P. Plassman. Mesh generation. *Handbook of Computational Geometry* (eds. J. Sack and J. Urrutia), Elsevier, 2000.

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

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

[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 mesh generation for curved surfaces. *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, (1993), 274–280.

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

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

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

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

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

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

[13] 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.

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

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

[16] 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.

[17] 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.

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

[19] 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.

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

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

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

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

[24] S.-H. Teng and C.-W. Wong. Unstructured mesh generation: Theory, practice and perspectives. *Intl. J. Comput. Geom. Appl.*, 10 (2000), 227–266.