On the Sizes of Delaunay Meshes*

Siu-Wing Cheng
Department of Computer Science
The Hong Kong University of Science and Technology

July 25, 2005

Abstract

Let $\mathcal{P}$ be a polyhedral domain occupying a convex volume. We prove that the size of a graded mesh of $\mathcal{P}$ with bounded vertex degree is within a factor $O(H_{\mathcal{P}}^3)$ of the size of any Delaunay mesh of $\mathcal{P}$ with bounded radius-edge ratio. The term $H_{\mathcal{P}}$ depends on the geometry of $\mathcal{P}$ and it is likely a small constant when the boundaries of $\mathcal{P}$ are fine triangular meshes. There are several consequences. First, among all Delaunay meshes with bounded radius-edge ratio, those returned by Delaunay refinement algorithms have asymptotically optimal sizes. This is another advantage of meshing with Delaunay refinement algorithms. Second, if no input angle is acute, the minimum Delaunay mesh with bounded radius-edge ratio is not much smaller than any minimum mesh with aspect ratio bounded by a particular constant.

Key words: mesh generation, Delaunay triangulation, radius-edge ratio, aspect ratio, gradedness

1 Introduction

Generating meshes of polyhedral domains are important steps in numerical simulations in scientific and engineering applications. A mesh is a partition of the domain into elements of simple shape, and three-dimensional triangulations are popular meshes.

The mesh quality is often measured by the shape of the tetrahedra, the edge lengths, and the mesh size. A popular shape measure of a triangle or tetrahedron τ is the *aspect ratio* $\frac{L^{\dim(\tau)}}{\text{volume}(\tau)}$, where L is the longest edge length of τ. A tetrahedron τ is well-shaped if its aspect ratio is upper bounded by a constant. A weaker measure is the *radius-edge ratio* which is the ratio of the circumradius of τ to its shortest edge length. A mesh has bounded aspect ratio (resp., radius-edge ratio) if the aspect ratio (resp., radius-edge ratio) of all tetrahedra are bounded. Bounded radius-edge ratio eliminates almost all tetrahedra with large aspect ratio except for a class known as *slivers* [2]. Still, bounded radius-edge ratio is good enough for some applications [13]. A mesh is *graded* if the distance from a vertex v to the nearest vertex is at least a constant factor of the local feature size at v. Gradedness is instrumental in proving the optimality of mesh sizes when there is no sliver [1, 9, 15].

Delaunay meshes are popular in theory and practice. A basic technique to generate a Delaunay mesh is called *Delaunay refinement*, which was studied in $\mathbb{R}^2$ by Chew [5] and Ruppert [15], and

*Research supported by the Research Grant Council, Hong Kong, China (project no. HKUST6190/02E).

then extended to $\mathbb{R}^3$ by Shewchuk [16]. There are two major challenges in $\mathbb{R}^3$ for the ordinary Delaunay refinement technique. First, only bounded radius-edge ratio can be guaranteed and so slivers may remain in the mesh. Second, the algorithm may not terminate when there are acute input angles, which include dihedral angles between faces, angles between edges, and angles between faces and non-incident edges.

The sliver problem has been addressed in a series of works. Cheng et al. [2] proposed the *sliver exudation* technique to eliminate slivers for unbounded domains. Chew [6] outlined a randomized point insertion strategy to handle slivers. Later, Li and Teng [9], and Cheng and Dey [1] developed algorithms to deal with boundaries. All the above works assume that no input angle is acute. Cheng and Poon [4] proposed an algorithm to deal with acute input angles. A graded Delaunay mesh with bounded radius-edge ratio is guaranteed. Note that the radius-edge ratio is dependent on the input small angles in the vicinity of them. More recently, Cheng, Dey, Ramos, and Ray [3] proposed a new algorithm for polyhedra that has similar guarantees as that of [4], but is much simpler to implement. The mesh is graded and most tetrahedra have bounded radius-edge ratio except possibly those that are close to small input angles. Pav and Walkington [11] presented another algorithm for domains with non-manifold boundaries.

Regarding the size analysis of meshes, the focus has been to show that the meshes computed have asymptotically optimal sizes when compared with any minimum mesh with bounded aspect ratio [1, 9, 14]. Although it makes sense to compare against a mesh with bounded aspect ratio, there are other interesting questions about mesh sizes. Can we compare against a mesh with a weaker shape measure such as the bounded radius-edge ratio? This question is meaningful for Delaunay refinement algorithms that generate Delaunay meshes with bounded radius-edge ratio. In other words, is the mesh size obtained the best possible among all Delaunay meshes with bounded radius-edge ratio?

Our main result is that given a polyhedral domain $\mathcal{P}$ occupying a convex volume, the size of a graded mesh with bounded vertex degree is within a factor $O(H_{\mathcal{P}}^3)$ of the size of any Delaunay mesh with bounded radius-edge ratio. The term $H_{\mathcal{P}}$ depends on the geometry of $\mathcal{P}$ and it is likely a small constant when the boundaries of $\mathcal{P}$ are fine triangular meshes. For a non-convex domain, we can place it inside a large box and mesh the inside of the box as in previous approaches [1, 7, 15]. Thus the requirement of convexity is not a severe restriction.

There are several consequences. First, among all Delaunay meshes with bounded radius-edge ratio, those returned by Delaunay refinement algorithms have asymptotically optimal sizes. This is another advantage of meshing with Delaunay refinement algorithms, in addition to its simplicity and shape guarantees. Second, if no input angle is acute, the minimum Delaunay mesh with bounded radius-edge ratio is not much smaller than any minimum mesh with aspect ratio bounded by a particular constant $A_{\mathcal{P}}$ (to be defined in Section 4). This further implies that the Delaunay meshes returned by Delaunay refinement algorithms have similar sizes as any minimum mesh with aspect ratio bounded by $A_{\mathcal{P}}$. Hence, although radius-edge ratio is a less stringent shape measure than aspect ratio, it does not lead to a significantly more compact mesh.

The paper is organized as follows. Section 2 classifies the shape of triangles and tetrahedra. Section 3 proves our main result. In Section 4, we give two applications of our main result, and we argue that a small $H_{\mathcal{P}}$ does not necessarily lead to bounded aspect ratio. We conclude in Section 5.

2 Shape measure

The poorly shaped elements of a mesh refer to the triangles and tetrahedra with large aspect ratio. We introduce a classification due to Cheng et al. [2]. A poorly shaped triangle has either one or two tiny angles. See Figure 1. Observe that poorly shaped triangles have large radius-edge ratio. Tetrahedra with large aspect ratio can be classified into two categories. The first category consists of those with vertices lying close to a line. See Figure 2. Each of the tetrahedra shown has one or more poorly shaped triangles as faces. If the vertices of a tetrahedron with large aspect ratio do not lie close to a line, they must lie close to a plane. See Figure 3. The left two tetrahedra in the top row contain one or more poorly shaped triangles as faces. Since bounded radius-edge ratio eliminates all poorly shaped triangles, it eliminates all the poorly shaped tetrahedra mentioned so far. The other two tetrahedra in the top row in Figure 3 are exceptions. The left one has the top vertex very close to the base triangle, so its circumradius is much larger than its edge lengths. Thus bounded radius-edge ratio still eliminates this type of tetrahedra. The rightmost figure in the top row in Figure 3 is called a sliver. The bottom figure in Figure 3 shows that a sliver actually has bounded radius-edge ratio. Therefore, bounded radius-edge ratio is a weaker shape measure than bounded aspect ratio. Still, bounded radius-edge ratio is good enough for some applications [13].

Figure 1: Poorly shaped triangles.

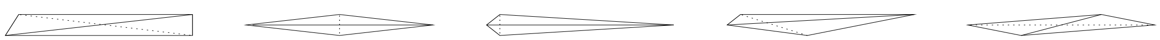

Figure 2: Tetrahedra with vertices close to a line.

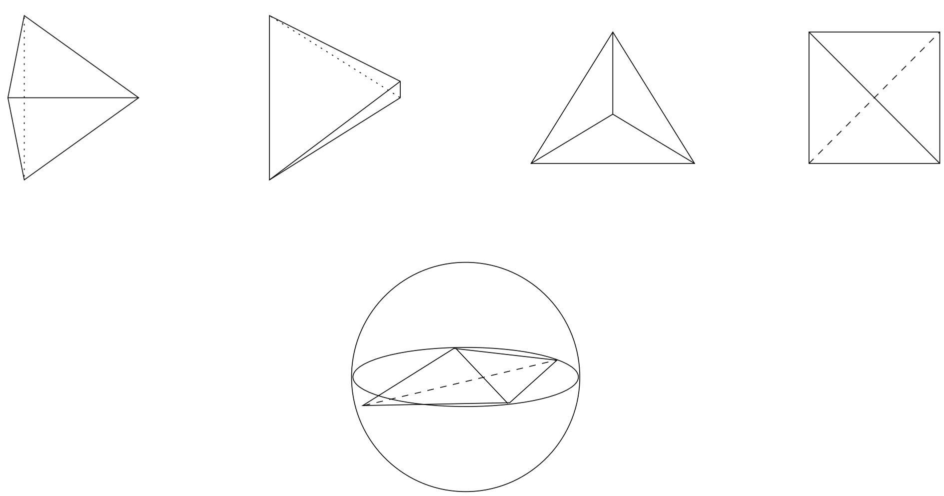

Figure 3: Tetrahedra with vertices close to a plane.

3 Comparison of mesh sizes

We use $\mathcal{P}$ to denote the input polyhedral domain. The domain is specified by a *piecewise linear complex* which is a collection of vertices, edges and faces, called the *elements* of $\mathcal{P}$. We use $\partial\mathcal{P}$ to denote this collection of elements. The intersection of any two elements is either empty or a collection of lower-dimensional elements. Since we assume that $\mathcal{P}$ occupies a convex volume, we call $\mathcal{P}$ a convex polyhedral domain for simplicity. The domain $\mathcal{P}$ has to be bounded by facets in $\partial\mathcal{P}$, i.e., the convex outer boundary of $\mathcal{P}$ comprises of some facets in $\partial\mathcal{P}$. It should be kept in mind that some elements in $\partial\mathcal{P}$ may not lie on the outer convex boundary.

The *local feature size* 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 disjoint input elements. The function f is 1-Lipschitz, i.e., $|f(x) - f(y)| \leq \|x - y\|$ for any points $x, y \in \mathbb{R}^3$. A mesh is *graded* if the distance from a vertex v to the nearest vertex is $\Omega(f(v))$.

Let $\mathcal{M}$ be the Delaunay mesh of $\mathcal{P}$ with bounded radius-edge ratio that we would like to compare against. A common strategy is to show that the size of $\mathcal{M}$ is $\Omega(\int_{\mathcal{P}} \frac{dx}{f(x)^3})$, which has been done for the case that $\mathcal{M}$ has bounded aspect ratio [1]. Since bounded radius-edge ratio is a weaker measure, we need to add a parameter $H_{\mathcal{P}}$ defined as follows. Define two functions $r, h : \mathbb{R}^3 \to \mathbb{R}$ where $r(x)$ is the distance from x to the nearest vertex of $\mathcal{P}$ (including x itself), and $h(x) = \max\{1, r(x)/f(x)\}$. Then $H_{\mathcal{P}}$ is defined as $\max_{x \in \partial\mathcal{P}} h(x)$.

Our goal is show that the size of $\mathcal{M}$ is $\Omega((1/H_{\mathcal{P}}^3) \cdot \int_{\mathcal{P}} \frac{dx}{f(x)^3})$. Although $H_{\mathcal{P}}$ may be large for some input, it can be a small constant in some situations. For example, it is becoming popular to build a digital model of an object from dense point samples obtained by laser scanning. In this case, $\partial\mathcal{P}$ consists of fine triangular meshes returned by some surface reconstruction algorithm. Then the distance between a point $x \in \partial\mathcal{P}$ and the nearest vertex would not be large when compared with $f(x)$.

We mentioned that a non-convex domain can be handled by placing it inside a large box and then meshing the inside of the box. That is, we treat the box facets as input elements as well. For our results to be meaningful, we need to ensure that $H_{\mathcal{P}}$ will not be affected much. Such a box can be chosen as follows. Let C be the smallest bounding cube of the domain $\mathcal{P}$. Fix the center of C and expand its side length $2\sqrt{3}+1$ times to obtain the desired bounding box B of $\mathcal{P}$. For any point $z \in \partial\mathcal{P}$, since z lies in C, the distance from z to the boundary of B is at least $\sqrt{3} \cdot \text{sidelength}(C)$, which is the length of the diagonal of C. Thus, if a ball centered at z intersects the boundary of B, this ball must contain C and hence $\mathcal{P}$. It follows that for any point $x \in \partial\mathcal{P}$, $r(x)$, $f(x)$ and $h(x)$ remain the same despite the introduction of the box B. For any point x on the boundary of B, $r(x)$ is at most half the diagonal length of a facet of B, which is $(\sqrt{6} + 1/\sqrt{2}) \cdot \text{sidelength}(C)$. The distance from x to C is at least $\sqrt{3} \cdot \text{sidelength}(C)$, and the distance between two disjoint input elements on B is at least $(2\sqrt{3} + 1) \cdot \text{sidelength}(C)$. So $f(x) \geq \sqrt{3} \cdot \text{sidelength}(C)$. Thus $h(x) = \max\{1, r(x)/f(x)\} < 2$ for any point x on the boundary of B. Hence $H_{\mathcal{P} \cup B} \leq \max\{2, H_{\mathcal{P}}\}$.

In the following, we use $B(p, R)$ to denote a ball centered at p with radius R.

3.1 Covers of tetrahedra

We study the cover of a tetrahedron τ in $\mathcal{M}$ by balls. Such a cover was first introduced by Edelsbrunner for triangular meshes in $\mathbb{R}^2$ [7]. For any vertex v in $\mathcal{M}$, let ℓ_v denote the length of the shortest edge incident to v. Let $c_0 < 1$ be an arbitrary constant. We are to put four balls centered at the vertices of τ. At the vertex v of τ, the ball radius is set to be $c_0\ell_v$. It turns out that we can shrink the circumball of τ by a constant factor and yet τ is still covered by the shrunk ball and the four vertex balls. We need the following result by Miller, Talmor, Teng, and Walkington [10].

Lemma 1 *Let $\mathcal{M}$ be a Delaunay mesh with bounded radius-edge ratio. There is a constant $\nu \geq 1$ such that for any two edges pq and pu, $\|p - q\| \leq \nu \cdot \|p - u\|$.*

We can then prove that the vertex balls and the shrunk circumball cover the tetrahedron.

Lemma 2 *Let $\mathcal{M}$ be a Delaunay mesh with bounded radius-edge ratio. Let z and R be the circumcenter and circumradius, respectively, of a tetrahedron τ in $\mathcal{M}$. For any $0 < c_0 < 1$, there is a constant $0 < c_2 < 1$ such that for any point $p \in \tau$, $p \in B(z, c_2R)$ or $p \in B(v, c_0\ell_v)$ for some vertex v of τ.*

Proof. Let $p \in \tau$ be a point that lies outside the balls at the vertices of τ. The Voronoi diagram of the four vertices partitions τ into four regions. Assume that p lies inside the region owned by the vertex v. This implies that $\angle zpv \geq \pi/2$. See Figure 4. So $\|p - z\|^2 \leq R^2 - \|p - v\|^2$. Since $p \notin B(v, c_0\ell_v)$, $\|p - v\| > c_0\ell_v$. Assume that the tetrahedra in $\mathcal{M}$ have radius-edge ratio bounded by some constant $\rho \geq 1$. Then Lemma 1 implies that $\|p - v\| > c_0R/(\nu\rho)$. It follows that

$$\begin{aligned} \|p - z\| &\leq \sqrt{R^2 - \|p - v\|^2} \\ &\leq \frac{\sqrt{\nu^2\rho^2 - c_0^2}}{\nu\rho} \cdot R. \end{aligned}$$

So we can set $c_2 = \sqrt{\nu^2\rho^2 - c_0^2}/(\nu\rho)$. ∎

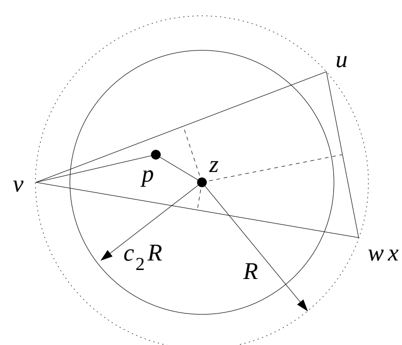

Figure 4: We look at the tetrahedron $uvwx$ along the direction of wx. The partition of $uvwx$ by the Voronoi diagram is shown.

The next result shows that the local feature sizes of points inside the cover cannot be arbitrarily small.

Lemma 3 *Let $\mathcal{M}$ be a Delaunay mesh with radius-edge ratio bounded by some constant $\rho \geq 1$. Let z and R be the circumcenter and circumradius, respectively, of a tetrahedron τ in $\mathcal{M}$. Assume that $c_0 \leq 1/(4\rho)$.*

(i) For any point $x \in B(z, c_2R)$, $f(x) \geq (1 - c_2)R/(4H_{\mathcal{P}})$.

(ii) Let v be a vertex of τ. For any point $x \in B(v, c_0\ell_v)$, $f(x) \geq c_0\ell_v$.

Proof. Consider (i). Let σ and σ' be the input elements that define $f(x)$. Let B be the ball $B(x, f(x))$. If B does not lie inside $B(z, (1 + c_2)R/2)$, then $f(x) \geq (1 - c_2)R/2$. See Figure 5. Otherwise, let p be the point on $\sigma \cap B$ that achieves the minimum distance D between $\sigma \cap B$

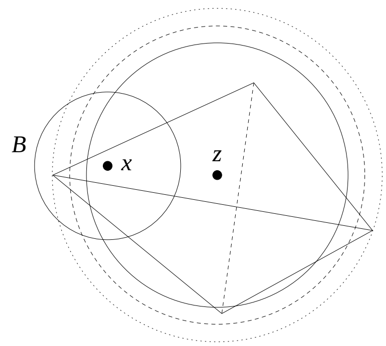

Figure 5: The dotted ball is the circumball of τ. The dashed ball is $B(z, (1 + c_2)R/2)$. The solid ball is $B(z, c_2R)$.

and $\sigma' \cap B$. See Figure 6. As σ and σ' are disjoint, $D \geq f(p)$. Since $\mathcal{M}$ is Delaunay, the

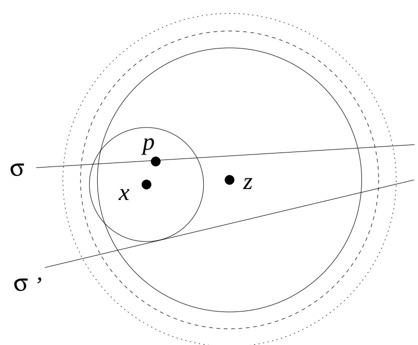

Figure 6: B lies inside $B(z, (1 + c_2)R/2)$.

nearest vertex to p lies outside the circumball of τ. Thus $r(p) \geq (1 - c_2)R/2$. It follows that $D \geq r(p)/h(p) \geq (1 - c_2)R/(2H_{\mathcal{P}})$. Hence $f(x) \geq D/2 \geq (1 - c_2)R/(4H_{\mathcal{P}})$.

Consider (ii). We claim that $f(v) \geq 2c_0\ell_v$. This claim implies that for any point $x \in B(v, c_0\ell_v)$, $f(x) \geq f(v) - c_0\ell_v \geq c_0\ell_v$ as desired. The ball $B(v, f(v))$ touches some input element. Since $\mathcal{P}$ is convex, the radius of $B(v, f(v))$ connecting v to this input element intersects the face $u_1u_2u_3$ of some tetrahedron $vu_1u_2u_3$ incident to v. We show that the distance from v to the triangle $u_1u_2u_3$ is at least $\ell_v/(2\rho)$ and thus $f(v) \geq \ell_v/(2\rho) \geq 2c_0\ell_v$.

Let s be the shortest segment connecting v to the triangle $u_1u_2u_3$. If s connects to u_i, then clearly length$(s) \geq \ell_v$. If s connects to the interior of an edge u_iu_j, then s is perpendicular to u_iu_j and length$(s) = |vu_i| \cdot \sin \angle vu_iu_j \geq \ell_v \cdot \sin \angle vu_iu_j$. Since the radius-edge ratio of the triangle vu_iu_j is also bounded by ρ, $\sin \angle vu_iu_j \geq 1/(2\rho)$. Thus length$(s) \geq \ell_v/(2\rho)$. The remaining possibility is

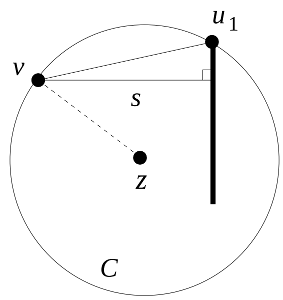

Figure 7: The bold line segment is the cross-section of the triangle $u_1u_2u_3$.

that s connects to the interior of $u_1u_2u_3$. Then s is perpendicular to $u_1u_2u_3$. Let L be the support line of s. Let H_i the halfplane that contains u_i and is bounded by L. Let $B(y,r)$ be the circumball of the tetrahedron $vu_1u_2u_3$. Let H be the halfplane that contains y and is bounded by L (if $y \in L$, we take H to be an arbitrary halfplane bounded by L). Notice that H_1, H_2 and H_3 divide $\mathbb{R}^3$ into three wedges, each with angle less than π. Then the dihedral angle between H and some H_i, say H_1, lies in the range $[\pi/2, \pi]$; otherwise, the angle of one of the three wedges would be greater than π. So the orthogonal projection z of y onto the plane containing H_1 does not lie strictly inside H_1. Take the cross-section C of $B(y,r)$ on the plane containing H_1. Note that u_1 and z do not lie on the same side of s in the cross-section. See Figure 7. We have $\text{length}(s) = |vu_1| \cdot \cos\phi$, where ϕ is the angle between vu_1 and s. Observe that $\phi \leq \angle u_1vz$ and $\cos\angle u_1vz = |vu_1|/(2|vz|) \geq |vu_1|/(2r) \geq 1/(2\rho)$. Therefore, $\text{length}(s) = |vu_1| \cdot \cos\phi \geq |vu_1| \cdot \cos\angle u_1vz \geq \ell_v/(2\rho)$. ◻

3.2 Analysis of mesh sizes

We use Lemma 3 to derive a lower bound on the size of $\mathcal{M}$. Then we upper bound the size of any mesh $\mathcal{T}$ of $\mathcal{P}$ that is graded and has bounded vertex degree. These two bounds imply our main result.

Lemma 4 *Let $\mathcal{M}$ be a Delaunay mesh of $\mathcal{P}$ with bounded radius-edge ratio. The number of tetrahedra in $\mathcal{M}$ is $\Omega((1/H_{\mathcal{P}}^3) \cdot \int_{\mathcal{P}} \frac{dx}{f(x)^3})$.*

Proof. Take the covers of tetrahedra in $\mathcal{M}$. For each ball B_i in the covers, let r_i denote its radius and let f_i denote the minimum local feature size of points inside B_i. By Lemma 3, we have $f_i \geq kr_i/H_{\mathcal{P}}$ for some constant k. This implies that

$$\begin{aligned}\int_{\mathcal{P}} \frac{dx}{f(x)^3} &\leq \sum_i \int_{B_i} \frac{dx}{f(x)^3} \\ &\leq \sum_i \frac{H_{\mathcal{P}}^3}{k^3 r_i^3} \cdot \text{vol}(B_i) \\ &= \sum_i \frac{4\pi H_{\mathcal{P}}^3}{3k^3}.\end{aligned}$$

The number of tetrahedra in $\mathcal{M}$ is at least $\sum_i \frac{1}{5}$ which is at least $\frac{3k^3}{20\pi H_{\mathcal{P}}^3} \int_{\mathcal{P}} \frac{dx}{f(x)^3}$. ◻

We need a technical result that there is a constant $\epsilon_0 > 0$ such that for any point x inside $\mathcal{P}$ and for any $\epsilon \leq \epsilon_0$, a constant fraction of $B(x, \epsilon f(x))$ lies inside $\mathcal{P}$. We first define some terminologies. At each vertex v on the convex outer boundary of $\mathcal{P}$, the incident facets of v delimit an unbounded convex polyhedral cone Q_v with apex v. We use C_v to denote the unbounded circular cone inside Q_v with apex v that has the largest angular aperture. Let θ_1 be the minimum angular aperture of C_v among all vertices v on the outer convex boundary of $\mathcal{P}$. Let θ_2 be the minimum angle between two adjacent edges, an edge and a facet sharing only one vertex, two facets sharing only one vertex, and two facets sharing edge(s) in $\partial\mathcal{P}$. We define $\theta = \min\{\theta_1, \theta_2\}$.

Lemma 5 *There is a constant $\epsilon_0 > 0$ such that for any point x inside $\mathcal{P}$ and for any $\epsilon \leq \epsilon_0$, a fraction λ of $B(x, \epsilon f(x))$ lies inside $\mathcal{P}$, where λ depends on θ.*

Proof. Let S be the outer convex boundary surface of $\mathcal{P}$. We first deal with points on S and in its vicinity. Then we deal with points in the interior of $\mathcal{P}$. We will need the following two claims.

Claim 1 *For any $\alpha > 0$, if $q \notin B(p, \alpha f(p))$, then $\|p - q\| \geq \frac{\alpha}{1+\alpha} \cdot f(q)$.*

Proof. The Lipschitz property implies that $f(q) \leq f(p) + \|p - q\| \leq \frac{1+\alpha}{\alpha} \cdot \|p - q\|$. □

Claim 2 *For any $0 < \alpha \leq 1$, if $q \in B(p, \frac{\alpha f(p)}{2})$, then $B(q, \frac{\alpha f(q)}{2+\alpha}) \subseteq B(p, \alpha f(p))$.*

Proof. By the Lipschitz property, $f(q) \leq f(p) + \|p - q\| \leq (2+\alpha) f(p)/2$. So the distance from q to the boundary of $B(p, \alpha f(p))$ is at least $\alpha f(p)/2 \geq \alpha f(q)/(2 + \alpha)$. □

Case 1. Let v be a vertex on S. Clearly, $B(v, f(v))$ does not contain any other vertex, and all edges and faces stabbing $B(v, f(v))$ are incident to v. Notice that $B(v, f(v)) \cap Q_v \subset \mathcal{P}$, so $B(v, f(v)) \cap C_v \subset \mathcal{P}$. By elementary integration, the area of the spherical cap of $B(v, f(v)) \cap C_v$ is at least a fraction $(1 - \cos(\theta/2))/2$ of the area of $\partial B(v, f(v))$, which implies that $B(v, f(v)) \cap C_v$ is at least a fraction $(1 - \cos(\theta/2))/2$ of $B(v, f(v))$. Introduce the ball $B_v = B(v, f(v)/2)$. Notice that $B_v \cap C_v$ lies inside $\mathcal{P}$ and it is at least a fraction $(1 - \cos(\theta/2))/2$ of B_v. For any point $w \in B_v$ inside $\mathcal{P}$, by Claim 2, the ball $B_w = B(w, f(w)/3) \subseteq B(v, f(v))$. We make a copy Q' of Q_v and translate Q' so that its apex is w. Since Q' is a copy of Q_v, Q' also contains a circular cone C' with apex w and aperture at least θ. Since $w \in Q_v$ and Q_v is convex, we have $Q' \subset Q_v$. Then $B_w \cap Q' \subset B(v, f(v)) \cap Q_v \subset \mathcal{P}$, which implies that $B_w \cap C'$ lies inside $\mathcal{P}$. As before, $B_w \cap C'$ is at least a fraction $(1 - \cos(\theta/2))/2$ of B_w.

Case 2. Let e be an edge on S. The edge e bounds two halfplanes that contain the two facets on S incident to e. These two halfplanes delimit an unbounded convex wedge W_e. The angle of W_e is at least θ. Let x be a point on e that lies outside B_v for all vertices v on S. Grow a ball B centered at x until B intersects an input element σ different from e and the input facets incident to e. If σ is disjoint from e, then radius$(B) = f(x)$. If σ is an endpoint of e, then by Claim 1, radius$(B) \geq f(x)/3$. Otherwise, σ is an input edge or facet incident to an endpoint v of e. Since the angle between e and σ is at least θ, Claim 1 implies that radius$(B) \geq \|v - x\| \sin\theta \geq f(x) \sin\theta/3$. In all three cases, the wedge W_e cuts out a slice of B inside $\mathcal{P}$ which is at least a fraction $\theta/(2\pi)$ of B. Let $\epsilon_1 = \sin\theta/6$. Introduce the ball $B_x = B(x, \epsilon_1 f(x))$. Since $B_x \subset B$, $B_x \cap W_e$ lies inside

$\mathcal{P}$ and it is at least a fraction $\theta/(2\pi)$ of B_x. For any point $y \in B_x$ inside $\mathcal{P}$, by Claim 2, the ball $B_y = B(y, \frac{\epsilon_1 f(y)}{1+\epsilon_1}) \subseteq B$. We make a copy W' of W_e and translate W' so that its sharp edge passes through y. Since $y \in W_e$ and W_e is convex, we have $W' \subset W_e$. Thus $B_y \cap W' \subset B \cap W_e \subset \mathcal{P}$. Since W' has angle at least θ, $B_y \cap W'$ is at least a fraction $\theta/(2\pi)$ of B_y.

Case 3. We deal with the facets on S similarly. Let F be a facet on S. Let p be a point on F that lies outside B_x for all points x in ∂F that we have introduced. Grow a ball B centered at p until B intersects an input element σ different from F. If σ is disjoint from F, then radius$(B) = f(p)$. If σ is a vertex or edge of F, let x be the point at which B touches σ. Then $\|p - x\| \geq \epsilon_1 f(x)$ which implies that radius$(B) = \|p - x\| \geq \frac{\epsilon_1}{1+\epsilon_1} f(p)$. Otherwise, σ is an input edge or face adjacent to F and let x be the point on $F \cap \sigma$ closest to p. Since the angle between F and σ is at least θ, radius$(B) \geq \|p - x\| \sin\theta \geq \frac{\epsilon_1 \sin\theta}{1+\epsilon_1} f(p)$. In all three cases, half of B lies inside S. Let $\epsilon_2 = \frac{\epsilon_1 \sin\theta}{2+2\epsilon_1}$. Introduce the ball $B_p = B(p, \epsilon_2 f(p))$. Since $B_p \subseteq B$, half of B_p lies inside S. For any point $q \in B_p$ inside $\mathcal{P}$, by Claim 2, the ball $B_q = B(q, \frac{\epsilon_2 f(q)}{1+\epsilon_2}) \subseteq B$. So at least half of it lies inside S too.

Case 4. Take a point z inside S that lies outside the balls B_p for all points p on S. Let p be the point on S closest to z. So $\|p - z\| \geq \epsilon_2 f(p)$ which implies that $\|p - z\| \geq \frac{\epsilon_2}{1+\epsilon_2} f(z)$ by the Lipschitz property. Thus the ball $B_z = B(z, \frac{\epsilon_2 f(z)}{1+\epsilon_2})$ lies completely inside S.

Finally, set $\epsilon_0 = \epsilon_2/(1+\epsilon_2)$. So for any point x inside $\mathcal{P}$ and for any $\epsilon \leq \epsilon_0$, $B(x, \epsilon f(x))$ is a subset of the ball B_x that we constructed in the above. Let $\lambda_1 = (1 - \cos(\theta/2))/2$, $\lambda_2 = \theta/(2\pi)$, $\lambda_3 = 1/2$, and $\lambda_4 = 1$. It is easy to see that for $1 \leq i \leq 4$, if B_x is constructed in case i above, $B(x, \epsilon f(x)) \cap \mathcal{P}$ is at least a fraction λ_i of $B(x, \epsilon f(x))$. ◻

We are ready to upper bound the size of any graded mesh with bounded vertex degree.

Lemma 6 *Let $\mathcal{T}$ be a graded mesh of $\mathcal{P}$ with bounded vertex degree. The number of tetrahedra in $\mathcal{T}$ is $O(\int_{\mathcal{P}} \frac{dx}{f(x)^3})$.*

Proof. Since the vertex degree is bounded, it suffices to bound the number of vertices in $\mathcal{T}$. Since $\mathcal{T}$ is graded, for any vertex v of $\mathcal{T}$, its distance from the nearest vertex is at least $cf(v)$ for some constant $c > 0$. By Lemma 5, a fraction λ of the ball $B(v, \epsilon_0 f(v))$ lies inside $\mathcal{P}$. Let $k = \min\{\epsilon_0, c\}$. Thus if we place at each vertex v a ball B_v of radius $kf(v)/2$, the resulting balls are disjoint. Let $D_v = B_v \cap \mathcal{P}$. By Lemma 5, $\mathrm{vol}(D_v) \geq \lambda \cdot \mathrm{vol}(B_v)$. Therefore,

$$\begin{aligned}
\int_{\mathcal{P}} \frac{dx}{f(x)^3} &\geq \sum_v \int_{D_v} \frac{dx}{f(x)^3} \\
&\geq \sum_v \frac{\mathrm{vol}(D_v)}{f(v)^3(1+k/2)^3} \\
&\geq \lambda \cdot \sum_v \frac{\mathrm{vol}(B_v)}{f(v)^3(1+k/2)^3} \\
&= \lambda \cdot \sum_v \frac{k^3\pi}{6(1+k/2)^3},
\end{aligned}$$

which is a constant times the number of vertices in $\mathcal{T}$. ◻

Our main result follows immediately from Lemmas 4 and 6.

Theorem 1 *Let $\mathcal{P}$ be a convex polyhedral domain. Let $\mathcal{M}$ be a Delaunay mesh of $\mathcal{P}$ with bounded radius-edge ratio. Let $\mathcal{T}$ be a graded mesh of $\mathcal{P}$ with bounded vertex degree. Then the size of $\mathcal{T}$ is within a factor $O(H_{\mathcal{P}}^3)$ of the size of $\mathcal{M}$.*

4 Case studies

Delaunay refinement. There are two Delaunay refinement algorithms due to Cheng and Poon [4] and Shewchuk [16] that deal with inputs with and without acute angles, respectively. Both guarantee that the output mesh is graded and has bounded vertex degree and bounded radius-edge ratio. When $H_{\mathcal{P}}$ is a small constant, it immediately follows from Theorem 1 that the meshes computed by these two algorithms have asymptotically optimal sizes among all Delaunay meshes with bounded radius-edge ratio.

Well-shaped mesh. Assume that $\mathcal{P}$ does not have an acute input angle. Cheng and Dey [1] developed a weighted Delaunay refinement algorithm which produces a graded mesh with bounded vertex degree and bounded aspect ratio. Moreover, its size is asymptotically optimal when compared with any mesh with bounded aspect ratio.

Let $A_{\mathcal{P}}$ be the maximum aspect ratio in their mesh of the domain $\mathcal{P}$. Clearly, the size of their mesh is at least the size of the minimum mesh with aspect ratio at most $A_{\mathcal{P}}$. Then Theorem 1 implies that the size of any minimum mesh with aspect ratio at most $A_{\mathcal{P}}$ is within a factor $O(H_{\mathcal{P}}^3)$ of the size of any Delaunay mesh with bounded radius-edge ratio. This implies that when $H_{\mathcal{P}}$ is a small constant, a minimum Delaunay mesh with bounded-radius edge cannot be much smaller than any minimum mesh with aspect ratio at most $A_{\mathcal{P}}$.

On the other hand, it is known that the size of the Delaunay mesh returned by Delaunay refinement is within a constant factor of any mesh with bounded aspect ratio [12]. We conclude that the Delaunay mesh returned by Delaunay refinement has similar size as any minimum mesh with aspect ratio at most $A_{\mathcal{P}}$. Thus although radius-edge ratio is a less stringent shape measure than aspect ratio, it does not lead to a significantly more compact mesh when $H_{\mathcal{P}}$ is a small constant. It is unclear what can be said about meshes with aspect ratio larger than $A_{\mathcal{P}}$ though.

$H_{\mathcal{P}}$ and aspect ratio. Is it possible that a small $H_{\mathcal{P}}$ will force the Delaunay mesh $\mathcal{M}$ with bounded radius-edge ratio to have bounded aspect ratio as well? If this were the case, Theorem 1 has basically been discovered before [1, 9].

Edelsbrunner and Guoy [8] performed experiments with sliver exudation. They used models with very fine and nicely shaped boundary triangulations, which probably have small $H_{\mathcal{P}}$. It was found that after all tetrahedra with large radius-edge ratio were eliminated by Delaunay refinement, there were quite a number of tetrahedra with dihedral angles less than 5° remaining. This apparently supports our claim that small $H_{\mathcal{P}}$ does not forbid the production of slivers. But what if we restrict our attention to minimum Delaunay meshes with bounded radius-edge ratio? In the following, we describe a domain such that the minimum Delaunay mesh has bounded radius-edge ratio as well as a sliver with arbitrarily small dihedral angle.

Consider the convex polyhedron $\mathcal{P}$ in Figure 8. It has six vertices, four of which form a tetrahedron *abcd* so flat that it is close to being a square. Arrange *abcd* so that *bcd* lies at the equator of the circumsphere, and *a* lies slightly above the equator. We put *p* and *q* at the north and south poles of the circumsphere of *abcd*. We can easily enforce that p and q project vertically onto ab and cd, respectively. So *abcd* is Delaunay, and the triangles *abp* and *cdq* are vertical. Note that the angle between the triangles *abp* and *abc* is obtuse. Thus the circumsphere of *abpc* does not contain d and *abpc* is Delaunay. Similarly, *abpd*, *cdqa*, and *cdqb* are also Delaunay. Therefore, they form a Delaunay mesh $\mathcal{M}$ of $\mathcal{P}$. Clearly, $H_{\mathcal{P}}$ is a small constant, and $\mathcal{M}$ has bounded radius-edge ratio. But the aspect ratio of $\mathcal{M}$ is arbitrarily large as we can make *abcd* arbitrarily flat.

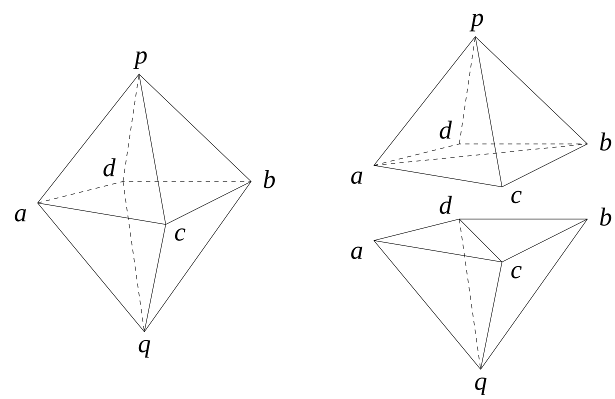

Figure 8: $\mathcal{P}$ and its minimum Delaunay mesh $\mathcal{M}$. We split *abcd* to show a better view of $\mathcal{M}$.

We argue that $\mathcal{M}$ has the minimum number of tetrahedra among all Delaunay meshes of $\mathcal{P}$. Let $\mathcal{T}$ be an arbitrary Delaunay mesh of $\mathcal{P}$. If $\mathcal{T}$ does not have any extra vertex, then $\mathcal{T} = \mathcal{M}$. If all the extra vertices lie on $\partial\mathcal{P}$, they split $\partial\mathcal{P}$ into nine or more triangles. Any tetrahedron in $\mathcal{T}$ has at most two of these triangles as its faces. Thus, there are at least $\lceil 9/2 \rceil = 5$ tetrahedra in $\mathcal{T}$. Suppose that a vertex v lies inside $\mathcal{P}$. Let the number of tetrahedra incident to v be Δ which must be at least three. Assume that δ of these Δ tetrahedra have faces on $\partial\mathcal{P}$. Then the number of tetrahedra in $\mathcal{T}$ is at least $\Delta + \lceil (8-\delta)/2 \rceil$. As $\Delta \geq 3$ and $0 \leq \delta \leq \Delta$, it can be verified that this quantity is at least six.

5 Conclusion

There have been significant advance in generating Delaunay meshes with shape guarantees. We attempt to augment this progress with a better understanding of the sizes of the Delaunay meshes returned by the Delaunay refinement algorithms. We show that given a convex polyhedral domain $\mathcal{P}$, the size of a graded mesh with bounded vertex degree is within a factor $O(H_{\mathcal{P}}^3)$ of the size of any Delaunay mesh with bounded radius-edge ratio. Although $H_{\mathcal{P}}$ depends on the geometry of $\mathcal{P}$, $H_{\mathcal{P}}$ is likely a small constant when the boundaries of $\mathcal{P}$ are fine triangular meshes. The consequences are that the Delaunay refinement algorithms produce meshes with asymptotically optimal size, and that although radius-edge ratio is a less stringent shape measure than aspect ratio, it does not lead to a significantly more compact mesh.

Acknowledgment

I thank the anonymous referee for comments and suggestions that improve the paper.

References

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

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

[3] S.-W. Cheng, T. K. Dey, E. Ramos, and T. Ray. Quality Meshing for Polyhedra with Small Angles. *Proceedings of the 20th Annu. ACM Sympos. Comput. Geom.*, 2004, 290–299.

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

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

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

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

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

[9] 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.

[10] 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.

[11] S. Pav and N. Walkington. 3D robust Delaunay refinement. *14th Intl. Meshing Roundtable* 2005.

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

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

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

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

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