Generating Well-Shaped *d*-dimensional Delaunay Meshes

Xiang-Yang Li

Dept. of Computer Science, Illinois Institute of Technology, 10 W. 31st Street, Chicago, IL 60616.

Abstract. A d-dimensional simplicial mesh is a Delaunay triangulation if the circumsphere of each of its simplices does not contain any vertices inside. A mesh is well-shaped if the maximum aspect ratio of all its simplices is bounded from above by a constant. It is a long-term open problem to generate well-shaped d-dimensional Delaunay meshes for a given polyhedral domain. In this paper, we present a refinement-based method that generates well-shaped d-dimensional Delaunay meshes for any PLC domain with no small input angles. Furthermore, we show that the generated well-shaped mesh has $O(n)$ d-simplices, where n is the smallest number of d-simplices of any almost-good meshes for the same domain. A mesh is almost-good if each of its simplices has a bounded circumradius to the shortest edge length ratio.

Keywords: Mesh generation, Delaunay triangulation, well-shaped, aspect ratio, radius-edge ratio, computational geometry, algorithms.

1 Introduction

Mesh generation is the process of breaking a geometry domain into a collection of primitive elements. In this paper we exclusively consider d-dimensional simplicial Delaunay meshes. The *aspect-ratio* of a mesh is the maximum aspect-ratio among all of its simplicial elements. A mesh is *well-shaped* if its aspect ratio is bounded from above by a small constant. The *aspect ratio* of a simplex is usually defined as the ratio of its circumradius to its inradius. Generating well-shaped d-dimensional Delaunay meshes is one of the long-term open problems in mesh generation when $d > 2$. An alternative but weaker quality measurement is to use the *radius-edge ratio* [8]. It is the circumradius divided by the shortest edge length of the simplex. The radius-edge ratio of a mesh is the maximum radius-edge ratio among all of its elements. A mesh is *almost-good* if it has a small radius-edge ratio. Numerous methods [3, 8–10] guarantee to generate almost-good 3-dimensional Delaunay meshes.

Bern *et al.* [1] showed that any set of n d-dimensional points had a Steiner Delaunay triangulation with $O(n^{\lceil d/2 \rceil})$ simplices, none of which has an obtuse dihedral angle. No bound depending only on n is possible if we require a bound on the minimum dihedral angle [1]. We assume a d-dimensional piecewise linear complex (PLC) domain Ω as the input. Shewchuk [11] showed that if each

k-dimensional constraining facet in Ω with $k \leq d-2$ is a union of strongly Delaunay k-simplices, then Ω has a d-dimensional constrained Delaunay triangulation. The Delaunay refinement method [3, 9, 10] can be extended to d-dimensions to generate almost-good Delaunay meshes if Ω satisfying the projection lemma [10] and having small input angles. However, it encounters significant difficulties in producing well-shaped Delaunay meshes in 3D. Chew [4] first proposed to add a vertex inside the circumsphere of any given badly-shaped tetrahedron (sliver) to remove it. The difficulty is in proving the existence of, and finding, a Steiner vertex that does not itself result in slivers. To make it possible, Chew defined a tetrahedron to be a sliver if it has a large aspect ratio, a small radius-edge ratio and its circumradius is no more than one unit. Without the restriction on the circumradius length, almost all results in [4] donot hold any more. In addition, with this restriction, the termination of his algorithm is straightforward. Recently, Li [6] extended this algorithm and showed how to generate well-shaped non-uniform Delaunay meshes in 3D.

In this paper, we show how to generate d-dimensional well-shaped meshes for a PLC domain with no small angles. In Section 2, we review the basic concepts and define what is a d-dimensional sliver simplex. [1] Then we give an algorithm in Section 3 to generate d-dimensional well-shaped Delaunay meshes. It basically adds a point around the circumcenter of each of d-simplices containing any k-dimensional sliver (hereafter, k-sliver). We prove its correctness in Section 4 and its termination guarantee in Section 5 by showing that the distance between the closest mesh vertices is just decreased by a constant factor compared with that of the input mesh. In Section 6, we show that the size of the generated mesh is within a constant factor of the size of any almost-good mesh generated for the same domain. Section 7 concludes the paper with discussions.

2 Preliminaries

After inserting a new vertex p, every new d-dimensional simplex created in the Delaunay triangulation of the new vertex set has p as one of its vertices. The new triangulation can be updated by efficient operations local to the vertex p. A sphere centered at a point c with a radius r is denoted as (c, r) hereafter. It is called *empty* if it does not contain any mesh vertices inside. The nearest neighbor graph defined by a d-dimensional vertex set is contained in the Delaunay triangulation of the vertex set. Thus the shortest edge length of the Delaunay triangulation is the closest distance among mesh vertices. This fact is used in proving the termination guarantee of our algorithm.

Delaunay refinement methods have been shown to be effective in generating almost-good meshes in 2 and 3 dimensions [9, 10]. There is also no much difficulty to extend them to d-dimensions, if the input domain satisfies a projection lemma [10], to generate meshes with radius-edge ratio no more than $\sqrt{2}^{d-1}$.

[1] Surprisingly, a simple extension of 3-dimensional sliver definition to d-dimensions does not work here.

For a k-simplex μ, its *min-circumsphere* is the smallest d-dimensional sphere containing the vertices of μ on its surface. A point *encroaches* the domain boundary if it is contained inside the min-circumsphere of a boundary k-simplex μ. Here a k-simplex μ is a boundary simplex if it belongs to the Delaunay triangulation of a k-dimensional input boundary polyhedral face. A k_1-simplex τ_1 *directly* encroaches another k_2-simplex τ_2 if the circumcenter of τ_1 encroaches τ_2. A k_1-simplex τ_1 *indirectly* encroaches another k_2-simplex τ_2 if the circumcenter of τ_1 encroaches a simplex μ and μ directly or indirectly encroaches τ_2. Assume the circumcenter of a k_1-simplex τ encroaches a boundary k_2-simplex μ. Call τ the *encroaching simplex* and μ the *encroached simplex.* Assume that μ contains the projection of the circumcenter of τ inside. Then Shewchuk [10] showed that the circumradii of τ and μ satisfy that $R_\mu \geq \frac{1}{\sqrt{2}} R_\tau$.

A d-simplex is *bad* if it has a large aspect ratio. Let's consider a k-simplex τ, where $1 \leq k \leq d$. For later convenience, we use R_τ, L_τ and $\rho(\tau) = R_\tau / L_\tau$ to denote the circumradius, the shortest edge length and the radius-edge ratio of τ. Let V be its volume. We define $\sigma = \sigma(\tau) = V/L_\tau^k$ as a measure of its quality. Let ϱ_0 and σ_k, $1 \leq k \leq d$ be positive constants that we specify later.

Definition 1. *A k-simplex τ is* well-shaped *if $\rho(\tau) \leq \varrho_0$ and $\sigma(\tau) \geq \sigma_k$.*

Definition 2. [SLIVER] *Call a k-simplex τ a k-*sliver *if $\rho(\tau) \leq \varrho_0$, $\sigma(\tau) < \sigma_k$, and each of its facets is well-shaped if $k > 3$.*

We call a d-simplex τ a *sliver-simplex* if it contains a k-sliver. If a d-simplex τ has a small radius-edge ratio and a small σ value, then it must have a face χ that is a sliver or τ itself is a d-sliver. It is not difficulty to prove that the volume of a d-simplex τ is at most $\varphi_d R_\tau^d$, where $\varphi_d = 2\prod_{i=2}^{d} \frac{(i-1)^{(i-1)/2}(i+1)^{(i+1)/2}}{i^{i+1}}$. The aspect ratio of τ is then at most $\frac{(1+d)\varphi_{d-1}\rho(\tau)^d}{d\sigma(\tau)}$. This verifies our definition of σ.

3 Refinement Algorithm

For a k-simplex τ, we call the intersection of its min-circumsphere with the affine space defined by its vertices as its k-sphere. When refine a k-simplex τ, we add a point p inside the shrinked k-sphere $(c_\tau, \delta R_\tau)$, where $\delta < 1$ is a constant. The point p is *good* if its insertion will not introduce any small slivers in the new Delaunay triangulation. Here a created sliver μ is *small* if $R_\mu \leq bR_\tau$ for a constant b specified later. We call the solid k-dimensional ball $(c_\tau, \delta R_\tau)$ the *picking region* of τ, denoted by $\mathcal{P}(\tau)$. Its volume is $\phi_k \delta^k R_\tau^k$, where $\phi_k = \frac{2\pi^{k/2}}{k\Gamma(\frac{k}{2})}$.

Algorithm: REFINEMENT (ϱ_0, σ_0, δ, b)

Enforce Empty Encroachment: Add the circumcenter c_τ of any encroached boundary simplex τ and update the Delaunay triangulation. If c_τ encroaches any lower dimensional boundary simplex μ, add c_μ instead of c_τ.

Clean Large Radius-Edge Ratio: Add the circumcenter c_τ of any d-simplex τ with a large $\rho(\tau)$ and update the Delaunay triangulation. If c_τ encroaches any boundary k-simplex, we apply the last rule instead of adding c_τ.

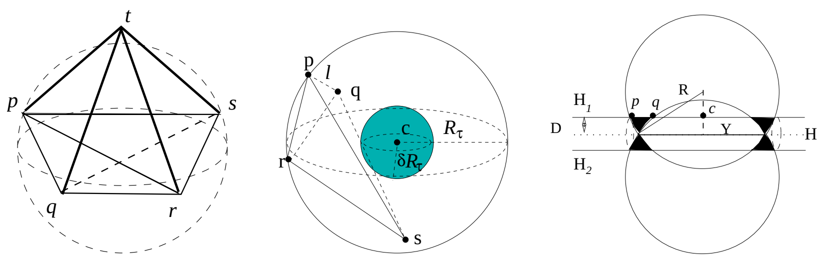

Fig. 1. Left: a 4-dimensional sliver example; Middle: the picking region of a simplex; Right: the forbidden region of μ with circumradius Y. Here $R = 2\varrho_0 Y$ and $D = 2k\sigma_0 Y$.

Clean the Slivers: For a sliver-simplex τ, add a good point $p \in \mathcal{P}(\tau)$ and update the Delaunay triangulation. If the circumcenter c_τ encroaches the domain boundary, we apply the following rules instead of finding p. If the insertion of p introduces some new d-simplices with large radius-edge ratio, we apply the previous rule to eliminate them immediately.

Encroach Boundary Simplices: If a boundary k-simplex μ is encroached directly or indirectly by a d-simplex with a large radius-edge ratio, add the circumcenter c_μ and update the Delaunay triangulation. However, if c_μ encroaches any other lower dimensional boundary simplex μ_1, we insert the circumcenter of μ_1 instead of adding c_μ.
If a boundary k-simplex μ is encroached directly or indirectly by a sliver-simplex, add a good point $p \in \mathcal{P}(\mu)$ and update the Delaunay triangulation. However, if c_μ encroaches any other lower dimensional boundary simplex μ_1, we add a good point from $\mathcal{P}(\mu_1)$ instead of adding p.

The key part of the algorithm is to find a good point p to refine a sliver-simplex τ. We select some k random points from $\mathcal{P}(\tau)$. Then choose the point that optimizes the quality σ of all created small simplices with a radius-edge ratio less than ϱ_0. Another approach is based on a randomized selection [4, 6, 7]. We randomly select a point $p \in \mathcal{P}(\tau)$ until we find a good point p. By defining sliver and small slivers properly, we can show that we are expected to find a good point p in constant rounds. Then the rest of the paper is devoted to prove the termination guarantee and the good grading guarantee of the algorithm.

4 Proofs of Correctness

4.1 Sliver Regions

Recall that a k-simplex τ is a sliver if $\rho(\tau) \leq \varrho_0$, $\sigma(\tau) \leq \sigma_k$, and all of its facets are well-shaped. The quality measure $\sigma(\tau)$ is related with a distance-radius ratio defined in [2, 5]. Consider any vertex p of τ. Let μ be the facet formed by other vertices of τ. Let D be the Euclidean distance of point p from the hyperplane passing through μ. Recall R_μ is the circumradius of μ.

Lemma 1. *For any k-simplex τ, $\frac{D}{R_\mu} \leq 2k\frac{\sigma(\tau)}{\sigma(\mu)}$.*

PROOF. It follows from the volume of τ: $V = \sigma(\tau)L_\tau^k = \frac{D}{k}\sigma(\mu)L_\mu^{k-1}$. ◻

A k-simplex τ is a sliver if $\rho(\tau) \leq \varrho_0$ and $\sigma(\tau) = \frac{V}{L_\tau^k} \leq \sigma_0^k$ for a small constant σ_0, i.e., $\sigma_k = \sigma_0^k$. Consequently, we have $\frac{D}{R_\mu} \leq 2k\sigma_0$ if τ is a k-sliver.

Given a well-shaped k-simplex τ, a point p that forms a $(k+1)$-sliver together with τ can not be anywhere. It must be around the circumsphere of τ and is not far-away from the hyperplane passing τ. We call the locus of such point p as the *sliver region* or *forbidden region* F_τ of the simplex τ. It is easy to show that F_τ is contained in the solid sun-hour glass shaped region as illustrated in the right figure of Figure 1. Let $|F_\tau|$ denote the volume of F_τ.

Lemma 2. *For any well-shaped k-simplex τ, $|F_\tau| \leq c_d\sigma_0^{d-k+1}R_\tau^d$, where c_d is a constant depending only on ϱ_0, k, and d.*

PROOF. We know that F_τ is inside the slab region defined by the two hyperplanes H_1 and H_2 illustrated by the right figure of Figure 1. Let r_2, r_1 be the circumradius of the two spheres intersected by H_1 and the circumsphere of τ. Here $r_2 = ||cp||$ and $r_1 = ||cq||$. It is easy to show that $r_2^2 \leq (1-(2k\sigma_0)^2 + 8k\sigma_0\varrho_0)\cdot Y^2$ and $r_1^2 \geq (1-(2k\sigma_0)^2 - 8k\sigma_0\varrho_0)\cdot Y^2$. The volume of a k-sphere with the radius r is $\phi_k r^k$. Then

$$|F_\tau| \leq \phi_k(r_2^k - r_1^k)\cdot\phi_{d-k}D^{d-k}.$$

Notice $\forall i \geq 1$, $r_2^{i+2} - r_1^{i+2} \leq (r_2^2 + r_1^2)(r_2^i - r_1^i)$. In addition, $r_2^2 + r_1^2 = 2(1-(2k\sigma_0)^2)Y^2 \leq 2Y^2$, and $r_2^2 - r_1^2 \leq 16k\varrho_0\sigma_0 Y^2$. The fact that $r_2 \geq Y$ implies that $r_2 - r_1 \leq 16k\varrho_0\sigma_0 Y$. By induction, we have $r_2^k - r_1^k \leq 2^{\lceil k/2\rceil+3}k\varrho_0\sigma_0 Y^k$. Consequently, we have $|F_\tau| \leq \phi_k\phi_{d-k}2^{\lceil k/2\rceil+3}k\varrho_0\sigma_0(2k\sigma_0)^{d-k}Y^d$. The lemma follows by setting $c_d = \phi_k\phi_{d-k}2^d k^{d-k+1}\varrho_0 > \phi_k\phi_{d-k}2^{\lceil k/2\rceil+3}k\varrho_0(2k)^{d-k}$. ◻

A good point p used to refine a k-simplex τ cannot be inside $F_\mu \cap \mathcal{P}(\tau)$ for any well-shaped m-simplex μ, We use $F_{\mu,k}$ to denote such intersection region.

Lemma 3. *For any k-simplex τ and a well-shaped m-simplex μ, $|F_{\mu,k}| \leq c_k\sigma_0 R_\mu^k$, where c_k is a constant that depends only on ϱ_0, d and k.*

4.2 Existence

We then show that given a k-simplex τ in an almost-good d-dimensional simplicial Delaunay mesh, there is a good point p in $\mathcal{P}(\tau)$. Let $S(\tau)$ be the set of simplices in any dimensions that each can form a small Delaunay sliver with a point from $\mathcal{P}(\tau)$. A created sliver μ is Delaunay if it belongs to the Delaunay triangulation by inserting a point from $\mathcal{P}(\tau)$; it is small if $R_\mu \leq bR_\tau$.

We first recall some results from Talmor *et al.* and extend it to d-dimensions. For a mesh vertex v, the *edge length variation*, denoted by $\nu(v)$, is defined as the length of the longest edge incident on it divided by the length of the shortest edge incident on it. Talmor[12] proved that, given an almost-good mesh in 3D, the edge length variation of each mesh vertex is at most a constant depending on the radius-edge ratio of the mesh. This result is extended to d-dimensions.

Lemma 4. *Given any vertex v of a d-dimensional almost good mesh, $\nu(v) \leq \nu_0$, where ν_0 is a constant depending on the mesh's radius-edge ratio ϱ_0.*

Lemma 5. *Given a k-simplex τ of an almost-good mesh, the number of the vertices of $S(\tau)$ is at most a constant depending on ν_0 and b.*

Proof. Assume that a well-shaped m-simplex μ forms a Delaunay sliver together with a point $p \in \mathcal{P}(\tau)$. The edges of μ have lengths at most $2R_\mu \leq 2bR_\tau$. The edges incident on any vertex q of μ before point p is introduced have length at least $2bR_\tau/\nu_0$. It implies that the closest distance among all vertices from $S(\tau)$ is at least $2bR_\tau/\nu_0$. It is simple to show that all such vertices are inside the sphere centered at c_τ with a radius $(\delta+2b)R_\tau$. Then by a volume argument, the number of vertices of $S(\tau)$ is a constant. ◻

Thus the number of simplices in $S(\tau)$ is also a constant, let's say, W.

Theorem 1. *Good point exists in any k-simplex τ of an almost-good mesh.*

Proof. Each simplex μ from $S(\tau)$ claims a forbidden region F_μ with volume at most $c_k \cdot \sigma_0 R_\mu^k$. The circumradius R_μ of μ is at most bR_τ. The volume of $\mathcal{P}(\tau)$ is $\phi_k \delta^k R_\tau^k$. Therefore, if we select σ_0 such that $W \cdot c_k \sigma_0 b^k R_\tau^k < \phi_k \delta^k R_\tau^k$, then the pigeonhole principal will guarantee the existence of a good point $p \in \mathcal{P}(\tau)$. ◻

5 Termination Guarantee

A k-sliver is called *created* if it contains at least one Steiner vertex; otherwise it is called *original*. We classify the bad d-simplices into three categories. The simplices containing an *original k-sliver* are called *original sliver-simplices*. The simplices containing only *created k-slivers* are called *created sliver-simplices*. The third category has d-simplices with a large radius-edge ratio. Let $e(\tau)$ be the shortest edge length introduced by eliminating a d-simplex τ. Here $e(\tau)$ could be less than L_τ for a sliver-simplex τ. For simplicity, we assume that all original sliver-simplices are removed first.

Lemma 6. *For an original sliver-simplex μ, $e(\mu) \geq \frac{(1-\delta)}{\sqrt{2}^{d+1}} L_\tau$.*

Lemma 7. *For a simplex τ with $\frac{R_\tau}{L_\tau} > \varrho_0$, $e(\tau) \geq \frac{\varrho_0}{\sqrt{2}^{d-1}} L_\tau$.*

The proofs are similar to the 3-dimensional counterpart, which are omitted. It remains to show that refining any created sliver-simplex will not introduce shorter edges to the mesh. Let's consider a sliver-simplex μ. Assume it contains a k-sliver τ created by inserting a point from the picking region of an element, say $f(\tau)$. Element $f(\tau)$ is called the *parent* of the k-sliver τ. There are three cases: $f(\tau)$ is a sliver-simplex; $f(\tau)$ is a d-simplex with a large radius-edge ratio; $f(\tau)$ is an encroached boundary simplex.

Lemma 8. *Assume a sliver-simplex μ contains a k-sliver τ created by splitting another sliver-simplex $f(\tau)$. Then* $e(\mu) \geq \frac{(1-\delta)b}{\sqrt{2}^{d+1}} \cdot L_{f(\tau)}$.

Lemma 9. *Assume a sliver-simplex μ contains a k-sliver τ, and parent $f(\tau)$ has a large radius-edge ratio. Then* $e(\mu) \geq \frac{(1-\delta)\varrho_0}{\sqrt{2}^{d+1}} \cdot L_{f(\tau)}$.

Lemma 10. *Assume a sliver-simplex μ contains a sliver τ created by splitting a boundary m-simplex χ that is encroached directly or indirectly by $p(\chi)$. Then*

1. $e(\mu) \geq \frac{(1-\delta)\varrho_0 L_{p(\chi)}}{2^d}$, *if d-simplex $p(\chi)$ has a large radius-edge ratio.*
2. $e(\mu) \geq \frac{(1-\delta)bL_\chi}{\sqrt{2}^{d+1}}$, *if d-simplex $p(\chi)$ is a sliver-simplex.*

Combining all above analysis, we know that the shortest distance between all mesh vertices is at least $\frac{1-\delta}{\sqrt{2}^{d+1}}$ factor of that of the original mesh if we choose b, δ and ϱ_0 such that $(1-\delta)b \geq \sqrt{2}^{d+1}$ and $(1-\delta)\varrho_0 \geq 2^d$. We then have the following theorem by a volume argument.

Theorem 2. *The algorithm terminates in generating well-shaped Delaunay meshes.*

6 Good Grading Guarantee

As [6, 7, 10], we study the relation between the nearest neighbor function defined by the final mesh and the local feature size function defined by the input domain. The local feature size $lfs(x)$ of a point x is the radius of the smallest ball intersecting two non-incident segments or vertices. The nearest neighbor $N(x)$ of a point x is its distance to the second nearest mesh vertex.

With each vertex v, we associate an *insertion edge length* e_v equal to the length of the shortest edge connected to v immediately after v is introduced into the Delaunay mesh. Notice that v may not have to be inserted into the mesh actually. If v is an input vertex, then e_v is the distance between v and its nearest input neighbor. So $e_v \geq lfs(v)$ from the definition of $lfs(v)$. If v is from the picking region of a sliver-simplex τ, then $(1+\delta)R_\tau \geq e_v \geq (1-\delta)R_\tau$. The inequality $(1+\delta)R_\tau \geq e_v$ comes from the fact that there is a mesh vertex on the circumsphere of τ.

For the sake of convenience of analyzing, we also define a parent vertex $p(v)$ for each vertex v, unless v is an input vertex. Intuitively, for any noninput vertex v, $p(v)$ is the vertex "responsible" for the insertion of v. We discuss in detail what means by responsible here. Assume v is inserted inside the picking region of a simplex μ. If simplex μ has $\rho(\mu) \geq \varrho_0$, then $p(v)$ is the most recently inserted end point of the shortest edge of μ. If μ is a sliver-simplex containing an original

k-sliver τ, then $p(v)$ is an end point of the shortest edge of τ. If μ is a created sliver-simplex containing a created k-sliver τ, then $p(v)$ is the most recently inserted vertex of τ. If μ is an encroached boundary k-simplex, then $p(v)$ is the encroaching vertex. For the sake of simplicity, we always use the almost-good mesh generated by Delaunay refinement method as input mesh. Therefore the boundary simplices can not be encroached by input vertices. Ruppert [9] and Shewchuk [10] showed that $N()$ defined on the mesh generated by Delaunay refinement method is within a constant factor of $lfs(v)$, i.e., $N(v) \sim lfs(v)$.

Lemma 11. *Let v be a vertex in the final mesh and let $p = p(v)$. Then we have $e_v \geq lfs(v)$ for an input vertex v; and $e_v \geq C \cdot e_p$ for a Steiner point v, where C is a constant specified in the proof.*

PROOF. If v is an original input vertex, then $e_v \geq lfs(v)$ from the definition of $lfs(v)$. Thus the theorem holds. Then consider a non-input vertex v. We first consider that v is selected from the picking region of a d-simplex; say μ.

Case 1.1: μ is a d-simplex with large radius-edge ratio $\rho(\mu) \geq \varrho_0$. The parent p is one of the end points of the shortest edge of μ. Let L_μ be the length of the shortest edge pq of μ. Then q is an original vertex or is inserted before p. Therefore, $e_p \leq \|p - q\| = L_\mu \leq \frac{R_\mu}{\varrho_0}$. Thus $e_v = R_\mu \geq \varrho_0 \cdot e_p$.

Case 1.2: μ is a sliver-simplex containing an original k-sliver τ. The parent $p = p(v)$ is one of the end points of the shortest edge of τ. Let L_τ be the length of the shortest edge pq of τ. Similar to the previous case, we have $e_p \leq \|p - q\| = L_\tau$. Notice that $R_\mu \geq R_\tau \geq L_\tau/2$. Thus $e_v \geq (1 - \delta)R_\mu \geq \frac{1-\delta}{2} \cdot e_p$.

Case 1.3: μ is a sliver-simplex containing a created k-sliver τ. There are three cases about the parent element $f(\tau)$ of τ: $f(\tau)$ is a sliver; $f(\tau)$ has a large radius-edge ratio; $f(\tau)$ is an encroached k-dimensional boundary simplex. Recall that the parent vertex $p = p(v)$ is the most recently inserted vertex of τ.

Subcase 1.3.1: $f(\tau)$ is a sliver. Recall that the insertion of p from the picking region of sliver $f(\tau)$ will always avoid creating small slivers. Thus, $R_\tau \geq bR_{f(\tau)}$, where $R_{f(\tau)}$ is the circumradius of $f(\tau)$. Notice that $e_p \leq (1 + \delta)R_{f(\tau)}$. Thus, $e_v \geq (1 - \delta)bR_{f(\tau)} \geq \frac{1-\delta}{1+\delta} b \cdot e_p$.

Subcase 1.3.2: $f(\tau)$ has $\rho(f(\tau)) \geq \varrho_0$. Let pq be an edge of simplex τ, where p is inserted from the picking region of $f(\tau)$. Notice that $e_p \leq \|p - q\|$. We also have $R_\tau \geq \|p - q\|/2$. Thus, $e_v \geq (1 - \delta)R_\mu \geq (1 - \delta)R_\tau \geq \frac{1-\delta}{2} \cdot e_p$.

Final subcase 1.3.3: $f(\tau)$ is an encroached boundary k-simplex. We first consider the scenario that $f(\tau)$ is encroached by a sliver-simplex directly or indirectly. We then know that the insertion of p from $f(\tau)$ will always avoid creating small slivers because the mesh is almost-good. Thus $R_\tau \geq bR_{f(\tau)}$. Notice that $e_p \leq (1 + \delta)R_{f(\tau)}$. Thus, $e_v \geq (1 - \delta)bR_{f(\tau)} \geq \frac{1-\delta}{1+\delta} b \cdot e_p$. We then consider the scenario that $f(\tau)$ is encroached by a d-simplex with a large radius-edge ratio directly or indirectly. Here parent p is selected from the picking region of $f(\tau)$. Let pq be an edge of τ. Notice that $e_p \leq \|p - q\|$. We also have $R_\tau \geq \|p - q\|/2$. Thus, $e_v \geq (1 - \delta)R_\mu \geq (1 - \delta)R_\tau \geq \frac{1-\delta}{2} \cdot e_p$.

Then we consider that v is selected from the picking region of a boundary k-simplex χ, which is encroached by a m-simplex μ.

Case 2.1: μ is encroached by a sliver-simplex directly or indirectly. Here parent p is always the circumcenter of μ. Notice that $R_\mu \leq \sqrt{2}R_\chi$ and $e_p \leq (1+\delta)R_\mu$. Thus, $e_v \geq (1-\delta)R_\chi \geq \frac{1-\delta}{\sqrt{2}(1+\delta)} \cdot e_p$.

Case 2.2: μ is encroached directly or indirectly by a d-simplex with a large radius-edge ratio. Parent p is the circumcenter of μ. We always have $R_\mu \leq \sqrt{2}R_\chi$ and $e_p = R_\mu$. Thus, $e_v = R_\chi \geq R_\mu\sqrt{2} \geq \frac{\sqrt{2}}{2} \cdot e_p$. ◻

For a vertex v, as [10], we define $D_v = \frac{lfs(v)}{e_v}$. We call D_v the *density ratio* at point v. Clearly, initially D_v is at most one for an input vertex v, and after inserting new vertices, D_v tends to become larger.

Lemma 12. *Let v be a vertex with a parent $p = p(v)$ if there is any. Assume that $e_v \geq C \cdot e_p$. If v is inserted due to eliminating sliver-simplex, then $D_v \leq \frac{1+\delta}{1-\delta} + \frac{D_p}{C}$. If v is inserted beause of eliminating a d-simplex with a large radius-edge ratio, then $D_v \leq 1 + \frac{D_p}{C}$.*

The proof is omitted, which is almost the same as the 3-dimensional counterpart.

Theorem 3. *There are fixed constants $D_k \geq 1$, $1 \leq k \leq d$ such at for any vertex v inserted or rejected at the picking region of a bad k-simplex, $D_v \leq D_k$. Specifically, the values of D_i should satisfy the following conditions:*

$$D_d \geq max\{\frac{\varrho_0 A + \sqrt{2}^{d+1} B_{d-2}}{\varrho_0(1-\delta) - 2^d \alpha^{d-1}}, \frac{\alpha b + \alpha^2 + \alpha B_{d-2}}{b - \sqrt{2}^{d-1}\alpha^d}, \alpha + \frac{2}{1-\delta}\},$$

$$D_{d-k} = B_{k-1} + \sqrt{2}^k \alpha^k D_d,$$

where $1 \leq k \leq d-1$, $A = \alpha(1-\delta) + \frac{\sqrt{2}^{d+1}-2}{\sqrt{2}-1}$, and $B_i = \sum_{j=1}^{i}(\sqrt{2}^j \alpha^{j+1})$. Hence, there is a constant $D = \max_{k=1}^{d}\{D_k\}$ such that $D_v \leq D$ for all mesh vertex v.

The proof, which is based on induction, of the theorem is omitted due to space limit. The following theorem concludes that the generated mesh has good grading, i.e., for any mesh vertex v, $N(v)$ is at least some constant factor of $lfs(v)$.

Theorem 4. *For any mesh vertex v generated by refinement algorithm, the distance connected to its nearest neighbor vertex u is at least $\frac{lfs(v)}{D+1}$.*

The proof is omitted here, which is the same as the three-dimensional counterpart. Thus, if $\varrho_0 > 2^d\alpha^d$ and $b > \sqrt{2}^{d-1}\alpha^d$, our algorithm generates well-shaped Delaunay meshes with good grading. Ruppert showed that the nearest neighbor value $N(v)$ of a mesh vertex v of any almost-good mesh is at most a constant factor of $lfs(v)$, where the constant depends on the radius-edge ratio of the mesh. The above Lemma 4 shows that the nearest neighbor $N(v)$ for the well-shaped Delaunay mesh is at least some constant factor of $lfs(v)$. Notice that the number of vertices of an almost-good mesh is $O(\int_{x \in \Omega} \frac{1}{N(x)^d} dx)$. Then we have the following theorem.

Theorem 5. *Given a d-dimensional almost-good mesh with n vertices, the generated well-shaped mesh has $O(n)$ vertices, where the constant depends on d and the radius-edge ratios of the meshes.*

7 Discussions

In this paper, we present a refinement-based method that guarantees to remove all slivers in a d-dimensional almost-good simplicial mesh. Notice that the σ_0 derived from all the proofs may be too small for any practical use. We would like to conduct some experiments to see what σ_0 can guarantee that there is no small slivers created. In addition, we could have different definitions about slivers depending upon the location of the simplex: inside or near the domain boundary. This could also improve the bound on σ_0. Let ρ_d be the minimum radius-edge ratio of a d-simplex. It is easy to show that $\rho_2 = \frac{\sqrt{3}}{3}$ and $\rho_d = \frac{1}{2\sqrt{1-\rho_{d-1}^2}}$, which is much less than the radius-edge ratio bound ϱ_0 achieved by our algorithm (and also the Delaunay refinement). We would like to know if we can get better radius-edge ratio bound on the d-simplices of the generated mesh.

References

1. Bern, M., Chew, P., Eppstein, D., and Ruppert, J. Dihedral bounds for mesh generation in high dimensions. In *Proc. 6th ACM-SIAM Symposium on Discrete Algorithms* (1995), ACM, pp. 189–196.
2. Cheng, S. W., Dey, T. K., Edelsbrunner, H., Facello, M. A., and Teng, S. H. Sliver exudation. In *Proc. 15th ACM Symposium on Computational Geometry* (1999), pp. 1–13.
3. Chew, L. P. Constrained Delaunay triangulations. *Algorithmica 4* (1989), 97–108.
4. Chew, L. P. Guaranteed-quality delaunay meshing in 3d (short version). In *13th ACM Sym. on Comp. Geometry* (1997), pp. 391–393.
5. Edelsbrunner, H., Li, X. Y., Miller, G., Stathopoulos, A., Talmor, D., Teng, S. H., ÜNGÖR", A., and Walkington, N. Smoothing and cleaning up slivers. In *ACM Symposium on Theory of Computing (STOC00)* (2000).
6. Li, X. Y. *Sliver-free Three Dimensional Delaunay Mesh Generation.* PhD thesis, University of Illinois at Urbana-Champaign, 2000.
7. Li, X. Y., and Teng, S. H. Sliver-free three dimensional delaunay mesh generation. In *Symposium on Discrete Algorithms (SODA)* (2001).
8. Miller, G. L., Talmor, D., Teng, S. H., and Walkington, N. A delaunay based numerical method for three dimensions: generation, formulation, and partition. In *Proc.* 27^{th} *Annu. ACM Sympos. Theory Comput.* (1995), pp. 683–692.
9. Ruppert, J. A new and simple algorithm for quality 2-dimensional mesh generation. In *Third Annual ACM-SIAM Symposium on Discrete Algorithms* (1992), pp. 83–92.
10. Shewchuk, J. R. *Delaunay Refinement Mesh Generation.* PhD thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, 1997.
11. Shewchuk, J. R. A condition guaranteeing the existence of higher-dimensional constrained delaunay triangulations. In *Proc. 14th Annual Symposium on Computational Geometry* (1998), ACM, pp. 76–85.
12. Talmor, D. *Well-Spaced Points for Numerical Methods.* PhD thesis, Carnegie Mellon University, 1997.