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. Here a mesh is almost-good if each of its simplices has a bounded circumradius to the shortest edge length ratio.

Key words: Mesh generation, Delaunay triangulation, d-dimensional meshes, 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. A mesh is Delaunay if the circumsphere of each of its simplices does not contain any mesh vertices inside. 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 real 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* [1]. 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 [1–4] guarantee to generate almost-good 3-dimensional Delaunay meshes.

Bern *et al.* [5] 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 [5]. We assume a d-dimensional piecewise linear complex (PLC) domain Ω as the input. Shewchuk [**?**] 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 [2–4] can be extended to d-dimensions to generate almost-good Delaunay meshes if Ω satisfying the the above conditions, the projection lemma [4], and having small input angles. However, it encounters significant difficulties in producing well-shaped Delaunay meshes in 3D. Chew [6] first proposed to add a vertex inside the circumsphere of any given badly-shaped tetrahedron (sliver) to remove it. The tetrahedron is then no longer in the Delaunay triangulation of the augmented set of vertices. 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 [6] don't hold any more. In addition, with this restriction, the termination of his algorithm is straightforward. Recently, Li [7] 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.

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

2 Preliminaries

2.1 Delaunay Triangulation

Delaunay triangulations are widely used to generate d-dimensional simplicial meshes, especially the triangular and tetrahedral meshes. There are abundant well-studied algorithms to construct the Delaunay triangulations [8,9]. We will extensively use the following properties about Delaunay triangulations in our algorithm. 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. If the circumsphere of a d-simplex τ is empty, then τ is in the Delaunay triangulation and we call τ the *Delaunay simplex*. 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.

2.2 Delaunay Refinement

Delaunay refinement methods have been shown to be effective in generating almost-good meshes in 2 and 3 dimensions [3,4]. There is also no much difficulty to extend them to d-dimensions, if the input domain satisfies a projection lemma [4], to generate meshes with radius-edge ratio no more than $\sqrt{2}^{d-1}$.

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 [4] showed that the circumradii of τ and μ satisfy that $R_\mu \geq \frac{1}{\sqrt{2}} R_\tau$. It is of the Delaunay triangulation property that if the circumcenter c of a Delaunay d-simplex τ is not inside the domain, there must exist a boundary k-simplex whose circumsphere contains c inside. In a mesh generated by Delaunay refinement, the circumsphere of each boundary k-simplex is empty.

2.3 Slivers

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_τ be its volume. We define $\sigma(\tau) = V_\tau / L_\tau^k$ as a measure of its quality. 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. In other words, a simplex τ has a small aspect ratio if and only if it has a small $\sigma(\tau)$ value. Let ϱ_0 and σ_k, $1 \leq k \leq d$ be real positive constants that we will specify later.

Definition 2.1 [WELL-SHAPED] *A k-simplex τ is* well-shaped *if $\rho(\tau) \leq \varrho_0$ and $\sigma(\tau) \geq \sigma_k$.*

Definition 2.2 [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 $k > 3$.*

A necessary condition that τ is a k-sliver is that all its $(k-1)$-dimensional faces are well-shaped. Observe that, we do not have such requirement in three dimensions [10]. However, this condition is inherently satisfied so it passes away our attention. Recall that if a tetrahedron has a radius-edge ratio no more than ϱ_0, then so do all its triangle faces; a triangle μ with a radius-edge ratio at most ϱ_0 has a bounded $\sigma(\mu) = V/L_\mu^2$ value. It is easy to show that if a simplex has radius-edge ratio at most ϱ_0, then so do all of its faces.

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. The reason is as follows. If all of its $d-1$ dimensional facets are well-shaped, then then d-simplex τ itself is a d-sliver. Otherwise, there must exist a $d-1$ dimensional facet, say μ, that is not well-shaped. Thus, $\sigma(\mu) \leq \sigma_{d-1}$ because μ already has a small radius-edge ratio. Then recursively we know either μ itself is a $d-1$ dimensional sliver or it contains a lower dimensional sliver. Notice that, the above argument does not depend on the values of σ_i, for $1 \leq i \leq d$.

Consequently, we know that each d-simplex either is well-shaped, or it has a large radius-edge ratio, or it is a sliver-simplex.

3 Refinement Algorithm

In this section, we present our refinement-based algorithm to generate well-shaped Delaunay meshes.

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 shrunk k-sphere $(c_\tau, \delta R_\tau)$, where $\delta < 1$ is a real constant. The point p is *good* if its insertion will not introduce any d-dimensional sliver-simplices in the new Delaunay triangulation. Here a created sliver-simplex μ is *small* if $R_\mu \leq bR_\tau$ for a real 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 $\phi_k = \frac{2\pi^{k/2}}{k\Gamma(\frac{k}{2})}$. Here $\Gamma(n) = \prod_{i=1}^{n-1} i$, and $\Gamma(n+1/2) = \frac{\prod_{i=1}^{n}(2i-1)}{2^n}\sqrt{\pi}$.

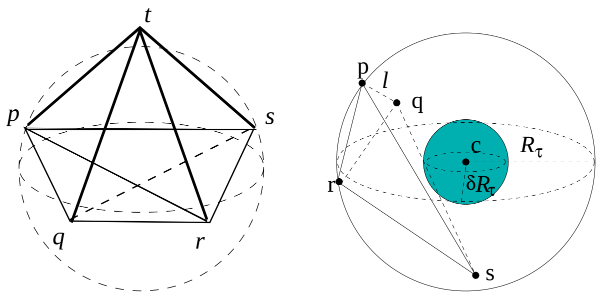

Fig. 1. Left: a 4-dimensional sliver example; Right: the shaded area is the picking region of a simplex.

The algorithm has the following basic steps. Firstly, we add the circumcenter of any d-simplex with a large radius-edge ratio until the mesh is almost-good. We then remove any d-simplex τ with a small σ value by adding a good point p from its picking region. Here the insertion of point p must avoid creating small d-dimensional simplex that contains some k-sliver in any dimension k. A created d-simplex μ is *small* if its circumradius is no more than b times that of the circumradius of τ. The existence of such point p is guaranteed by the Theorem 4.6. If the insertion of point p creates any d-simplex with large radius-edge ratio, we remove that new simplex first by adding its circumcenter. In other words, the d-simplices with large radius-edge ratio have priority over the sliver-simplices to be refined. If the circumcenter of the bad simplex encroaches the domain boundary, we split the encroached boundary simplices instead.

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

Enforce Empty Encroachment: Generate the Delaunay triangulation of all boundary facets. Add the circumcenter c_τ of any encroached boundary k-simplex τ and update the Delaunay triangulation. If c_τ encroaches any lower dimensional boundary simplex μ, add the circumcenter 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.

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 last rule 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.

After this step, the circumsphere of each boundary k-simplex is also empty.

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 d-simplices with a radius-edge ratio less than ϱ_0. Another approach is based on a randomized selection [6,7,11]. 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 $\rho(\tau) \leq \varrho_0$, $\sigma(\tau) < \sigma_k$, and all of its facets are well-shaped. The quality measure $\sigma(\tau)$ is related with a distance-radius ratio defined in [12,10]. 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. The following lemma shows that the value D/R_μ is no more than a constant factor of $\frac{\sigma(\tau)}{\sigma(\mu)}$.

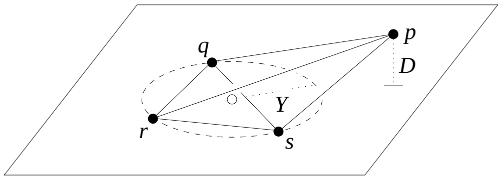

Fig. 2. D/R_μ depends on the ordering of the vertices of τ, but all such values are no more than a constant factor of $\frac{\sigma(\tau)}{\sigma(\mu)}$. Here $Y = R_\mu$.

Lemma 4.1 *For any k-simplex τ, $\frac{D}{R_\mu} \le 2k\frac{\sigma(\tau)}{\sigma(\mu)}$.*

PROOF. Notice that the volume of τ is $V = \sigma(\tau)L_\tau^k$, which also equals to $\frac{D}{k}\sigma(\mu)L_\mu^{k-1}$. Thus, from $R_u > L_\tau/2$, we have

$$\frac{D}{R_\mu} = k\frac{L_\tau}{R_u}\frac{\sigma(\tau)}{\sigma(\mu)} \le 2k\frac{\sigma(\tau)}{\sigma(\mu)}.$$

□

Notice that a simplex is a sliver if its volume is small compared with its size. Thus, we expect that the distance D should be bounded from above by a small constant factor of R_μ if we want to have a sliver τ. In other words, we need $\frac{\sigma(\tau)}{\sigma(\mu)}$ to be a bounded from above by a small constant. Thus, we have the following revisited definitions.

Definition 4.1 [SLIVER] *A k-simplex τ is a sliver if $\rho(\tau) \le \varrho_0$, $\sigma(\tau) = \frac{V_\tau}{L_\tau^k} \le \sigma_0^k$ for a small positive constant σ_0, i.e., $\sigma_k = \sigma_0^k$, and all of its facets are well-shaped. Here typically we require $2d\sigma_0 \le 1/2$.*

Definition 4.2 [WELL-SHAPED] *A k-simplex τ is* well-shaped *if $\rho(\tau) \le \varrho_0$ and $\sigma(\tau) \ge \sigma_0^k$.*

Consequently, we have $\frac{D}{R_\mu} \le 2k\sigma_0$ if τ is a k-sliver.

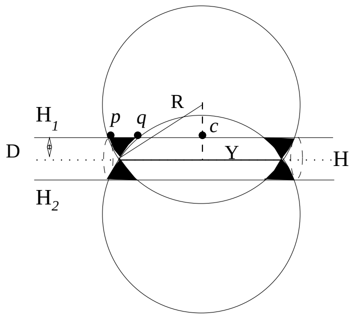

Fig. 3. The forbidden region of μ with circumradius Y. Here $R = 2\varrho_0 Y$ and $D = 2k\sigma_0 Y$.

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 by Figure 3. Notice that the forbidden region F_τ is a d-dimensional region. Let $|F_\tau|$ denote the volume of F_τ, which is bounded by the following lemma.

Lemma 4.2 [FORBIDDEN VOLUME] *For any well-shaped k-simplex τ, $|F_\tau| \le c_d \sigma_0^{d-k+1} R_\tau^d$, 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 Figure 3. Let R be the largest possible radius of d-simplex with a small radius-edge ratio ($\le \varrho_0$) such that it could have a facet τ with circumradius Y. Then $R < 2\varrho_0 Y$. 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 $r_1 = ||cq||$. Notice that

$$r_2^2 = R^2 - (\sqrt{R^2 - Y^2} - D)^2,$$

and

$$r_1^2 = R^2 - (\sqrt{R^2 - Y^2} + D)^2.$$

Using $R < 2\varrho_0 Y$ and $D \le 2k\sigma_0 Y$, it is easy to show that

$$r_2^2 \le (1 - (2k\sigma_0)^2 + 8k\sigma_0\varrho_0) \cdot Y^2$$

and

$$r_1^2 \ge (1 - (2k\sigma_0)^2 - 8k\sigma_0\varrho_0) \cdot Y^2.$$

The volume of a k-sphere with the radius r is $\phi_k r^k$. Then a simple gemetry observation implies that

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

Notice $\forall i \ge 1$, $r_2^{i+2} - r_1^{i+2} \le (r_2^2 + r_1^2)(r_2^i - r_1^i)$. In addition,

$$r_2^2 + r_1^2 = 2(1 - (2k\sigma_0)^2)Y^2 \le 2Y^2,$$

and

$$r_2^2 - r_1^2 \le 16k\varrho_0\sigma_0 Y^2.$$

The fact that $r_2 \ge Y$ implies that $r_2 - r_1 \le 16k\varrho_0\sigma_0 Y$. By induction, we have

$$r_2^k - r_1^k \le 2^{\lceil k/2 \rceil + 3} k\varrho_0\sigma_0 Y^k.$$

Consequently, we have

$$\begin{aligned} |F_\tau| &\le \phi_k\phi_{d-k} 2^{\lceil k/2 \rceil + 3} k\varrho_0\sigma_0 (2k\sigma_0)^{d-k} Y^d \\ &< \phi_k\phi_{d-k} 2^d k^{d-k+1} \varrho_0 \sigma_0^{d-k+1} Y^d \end{aligned}$$

The lemma follows by setting

$$\begin{aligned} c_d &= \phi_k \phi_{d-k} 2^d k^{d-k+1} \varrho_0 \\ &> \phi_k \phi_{d-k} 2^{\lceil k/2 \rceil + 3} k \varrho_0 (2k)^{d-k}. \end{aligned}$$

□

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

The proof is omitted here. Basically, if $k > m$, the proof is the same as the previous lemma. Otherwise, we find the k-dimensional hyperplane which has the maximum intersection volume with the forbidden region of m-simplex μ. This happens when the d-dimensional hyperplane H_1 or H_2 contains the k-simplex. The largest radius of the intersected region is at most $\sqrt{r_2^2 - r_1^2} \leq 4\sqrt{k \varrho_0 \sigma_0} Y$. Thus, the volume $|F_{\mu,k}|$ is at most $\phi_k (4\sqrt{k \varrho_0 \sigma_0} Y)^k$.

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 $R_\mu \leq b R_\tau$.

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[13] 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 4.4 [EDGE LENGTH VARIATION] *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.*

The proof is omitted here. It is almost the same as the three-dimensional case. We then show that $S(\tau)$ is at most a constant for any k-simplex τ of an almost-good mesh.

Lemma 4.5 [CONSTANT SMALL SLIVERS] *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 4.6 [EXISTENCE] *Good point exists in any k-simplex τ of an almost-good mesh, if we define sliver using σ_0 such that $W \cdot c_k \sigma_0^k b^k < \phi_k \delta^k$. In other words, we require*

$$\sigma_0 < (\frac{\phi_k}{W \cdot c_k})^{1/k} \cdot \frac{\delta}{b}.$$

PROOF. Each simplex μ from $S(\tau)$ claims a forbidden region F_μ with volume at most $c_k \cdot \sigma_0^k R_\mu^k$. 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

$$W \cdot c_k \sigma_0^k b^k R_\tau^k < \phi_k \delta^k R_\tau^k,$$

then the pigeonhole principal will guarantee the existence of a good point $p \in \mathcal{P}(\tau)$. ◻

5 Termination Guarantee

We then prove that the algorithm will terminate if we define what is bad simplex element properly.

For the sake of simplicity of analysis, we assume that the input mesh is already almost-good and the circumsphere of each boundary simplex is empty. For any bad d-simplex τ, it either has a large radius-edge ratio or it contains a k-dimensional face which is a k-sliver. 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. Notice that the newly created edges by inserting a point p from the picking region of τ to refine τ are all incident on p. Recall that L_τ is the shortest edge length of τ. Here $e(\tau)$ could be less than L_τ for a sliver-simplex τ. However, on the other hand, we will show that the shortest distance among mesh vertices is at least a constant factor of that of the original mesh. Thus, our algorithm terminates by a simple volume argument.

5.1 Original Sliver-Simplices

Let's first study the case of eliminating original sliver-simplices in the input mesh. Notice that in our algorithm, we do not distinguish the original sliver-simplices from the created sliver-simplices by point insertion. For simplicity, we assume that all original sliver-simplices are removed first.

Lemma 5.1 [ORIGINAL SLIVER-SIMPLEX] *Assume we refine an original sliver-simplex μ that contains an orginal k-sliver τ. Then $e(\mu) \geq \frac{(1-\delta)}{\sqrt{2}^{d+1}} L_\tau$.*

PROOF. Assume a point p is inserted to refine μ. If $p \in \mathcal{P}(\mu)$, then for any mesh vertex v

$$||v-p|| \geq (1-\delta)R_\mu \geq (1-\delta)R_\tau \geq (1-\delta)L_\tau/2.$$

If p is selected from the picking region of an encroached boundary simplex χ, then we have $R_\chi \geq R_\mu/\sqrt{2}^{d-1}$. Notice that $R_\mu \geq R_\tau \geq L_\tau/2$. Thus

$$||v-p|| \geq (1-\delta)R_\chi \geq (1-\delta)L_\tau/\sqrt{2}^{d+1}.$$

Then the lemma follows. ◻

5.2 Large Radius-Edge Ratio

Then let's study the scenario when we insert points to remove a d-simplex τ with a large radius-edge ratio. Simplex τ could be original or be created by refining other bad simplices.

Lemma 5.2 [LARGE RADIUS-EDGE RATIO] *Assume we refine a simplex τ with $\frac{R_\tau}{L_\tau} > \varrho_0$. Then $e(\tau) \geq \frac{\varrho_0}{\sqrt{2}^{d-1}} L_\tau$.*

PROOF. If the circumcenter c_τ of τ is inserted to remove τ, then

$$e(\tau) = R_\tau \geq \varrho_0 L_\tau.$$

Then assume that c_τ encroaches the circumsphere (c_χ, R_χ) of a boundary k-simplex χ directly or indirectly. Recall that $R_\chi \geq R_\tau / \sqrt{2}^{d-k}$. Thus

$$e(\tau) = R_\chi \geq \frac{R_\tau}{\sqrt{2}^{d-k}} \geq \frac{\varrho_0 L_\tau}{\sqrt{2}^{d-1}}.$$

Notice that it may need split the domain boundary several times till the simplex τ is eliminated. The above argument is true for each boundary splitting. Thus $e(\tau) \geq \frac{\varrho_0 L_\tau}{\sqrt{2}^{d-1}}$. □

Thus, if $\varrho_0 \geq \sqrt{2}^{d-1}$, the shortest edge length is not decreased by refining any d-simplex with a radius-edge ratio larger than ϱ_0.

5.3 Created Sliver-Simplices

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 $f(\tau)$. In other words, $f(\tau)$ is responsible for creating the new k-sliver τ. Element $f(\tau)$ is called the *parent* of the k-sliver τ. There are three cases: $f(\tau)$ is a sliver-simplex; and $f(\tau)$ is a d-simplex with a large radius-edge ratio; $f(\tau)$ is an encroached boundary simplex.

Let us first consider the case that $f(\tau)$ is a sliver-simplex. Recall that when we split $f(\tau)$, the mesh is almost-good. Then $R_\tau \geq bR_{f(\tau)}$ because the insertion of some good point in the picking region of $f(\tau)$ will always avoid creating small slivers.

Lemma 5.3 [LARGE SLIVER-SLIVER] *Assume a sliver-simplex μ contains a k-sliver τ created by splitting another sliver-simplex $f(\tau)$. Then*

$$e(\mu) \geq \frac{(1-\delta)b}{\sqrt{2}^{d+1}} \cdot L_{f(\tau)},$$

where $L_{f(\tau)}$ is the shortest edge length of $f(\tau)$.

PROOF. First consider the case that the circumcenter c_μ of μ does not encroach the boundary. Obviously, $e(\mu) \geq (1-\delta)R_\mu$, which is at least

$$(1-\delta)R_\tau \geq (1-\delta)bR_{f(\tau)} \geq (1-\delta)bL_{f(\tau)}/2.$$

Then consider the case that c_μ encroaches a boundary m-simplex χ directly or indirectly. Notice that $R_\chi \geq R_\mu/\sqrt{2}^{d-m}$. Thus $e(\mu)$ is at least

$$(1-\delta)R_\chi \geq \frac{1-\delta}{\sqrt{2}^{d-m}} R_\tau \geq \frac{(1-\delta)b}{\sqrt{2}^{d-m+2}} L_{f(\tau)}.$$

Then the lemma follows from $1 \leq m \leq d-1$. ◻

Then we study the second case that the parent $f(\tau)$ is a d-simplex with a large radius-edge ratio.

Lemma 5.4 [Large Sliver-Large R/L] *Assume a sliver-simplex μ contains a k-sliver τ, and parent $f(\tau)$ has a large radius-edge ratio. Then*

$$e(\mu) \geq \frac{(1-\delta)\varrho_0}{\sqrt{2}^{d+1}} \cdot L_{f(\tau)},$$

where $L_{f(\tau)}$ is the shortest edge length of $f(\tau)$.

Proof. We first consider the case that the circumcenter c_μ of μ does not encroach the boundary. Then $e(\mu) \geq (1-\delta)R_\mu \geq (1-\delta)R_\tau$. Let $p = c_{f(\tau)}$ and pq be an edge of τ. The shortest edge connected to p has length at least $R_{f(\tau)}$. From $R_\tau \geq \frac{\|p-q\|}{2}$, we have $R_\tau \geq \frac{R_{f(\tau)}}{2}$. Then

$$e(\mu) \geq \frac{(1-\delta)R_{f(\tau)}}{2}.$$

We then consider the case that c_μ encroaches a boundary m-simplex χ directly or indirectly. We insert a point from $(c_\chi, \delta R_\chi)$. Notice that $R_\chi \geq R_\mu/\sqrt{2}^{d-m}$. Then $e(\mu) = (1-\delta)R_\chi \geq \frac{1-\delta}{\sqrt{2}^{d-m}} R_\mu$. Similarly, from $R_\mu \geq R_{f(\tau)}/2$, we have

$$e(\mu) \geq \frac{(1-\delta)R_{f(\tau)}}{\sqrt{2}^{d+1}}.$$

Then $R_{f(\tau)} \geq \varrho_0 L_{f(\tau)}$ completes the proof. ◻

Then let's consider a sliver-simplex μ that contains a k-sliver τ created by splitting a boundary m-simplex χ.

Lemma 5.5 [Large Sliver-Boundary] *Assume a sliver-simplex μ contains a sliver τ created by splitting a boundary m-simplex χ that is encroached directly or indirectly by a d-simplex $p(\chi)$. Then*

(1) $e(\mu) \geq \frac{(1-\delta)\varrho_0 L_{p(\chi)}}{2^{d+1}}$*, 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.*

PROOF. Recall that we split χ because there is a bad d-simplex encroaches χ directly or indirectly. Let $p(\chi)$ be that bad d-simplex. There are two cases about $p(\chi)$: it is a sliver-simplex or it has a large radius-edge ratio.

Let us first study the case that $p(\chi)$ has a radius-edge ratio larger than ϱ_0. Then as proved by previous lemmas, the length $e(\mu)$ of the shortest edge introduced by splitting μ is at least $\frac{(1-\delta)R_\mu}{\sqrt{2}^{d-1}}$. Notice that here, we may need split some boundary simplices if the circumcenter of μ encroaches the domain boundary. It is always true that $R_\chi \geq R_{p(\chi)}/\sqrt{2}^{d-1}$. From $R_{p(\chi)} \geq \varrho_0 L_{p(\chi)}$, we have $R_\chi \geq \varrho_0 L_{p(\chi)}/\sqrt{2}^{d-1}$, where $R_{p(\chi)}$ and $L_{p(\chi)}$ are the circumradius and the shortest edge length of $p(\chi)$ respectively. Notice that $R_\mu \geq R_\tau \geq R_\chi/2$. Consequently, we have

$$e(\mu) \geq \frac{(1-\delta)R_\mu}{\sqrt{2}^{d-1}} \geq \frac{(1-\delta)\varrho_0 L_{p(\chi)}}{2^{d+1}}.$$

We then study the case that $p(\chi)$ is a sliver-simplex, i.e., one of its faces is a sliver. Recall that the mesh should be almost-good when we split sliver-simplex $p(\chi)$. Then the Existence Theorem 4.6 guarantees that τ is not a small sliver, i.e., $R_\tau \geq bR_\chi$. Notice that the length $e(\mu)$ of the shortest edge introduced by splitting μ is at least $\frac{(1-\delta)R_\tau}{\sqrt{2}^{d-1}}$, which is then at least $\frac{(1-\delta)bR_\chi}{\sqrt{2}^{d-1}}$. The fact that $R_\chi \geq L_\chi/2$ implies that

$$e(\mu) \geq \frac{(1-\delta)bR_\chi}{\sqrt{2}^{d-1}} \geq \frac{(1-\delta)bL_\chi}{\sqrt{2}^{d+1}}.$$

$\square$

The above lemma implies that, if $(1-\delta)b \geq \sqrt{2}^{d+1}$ and $(1-\delta)\varrho_0 \geq 2^{d+1}$, eliminating a sliver-simplex μ will not introduce shorter edges.

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

6 Good Grading Guarantee

This section is devoted to analyze the mesh size of the generated mesh. The algorithms by Ruppert and Shewchuk produce nicely graded meshes in the sense that *a small feature in one part of the mesh does not unreasonably reduce the edge lengths at places sufficiently far away* [14]. As showed by following analysis, the sliver-free d-dimensional Delaunay meshing algorithm presented in this paper also produces nicely graded meshes. In addition, the mesh size is within a small constant factor of the mesh generated by classic Delaunay refinement method for a same domain. The main idea behind the proof is to relate the spacing functions defined by the final mesh and by the input domain. Because the insertion of a point in the picking region of a sliver-simplex will likely reduce the shortest edge length locally, previous proofs by Ruppert and Shewchuk can not be applied directly here. Recall that we avoid creating small slivers when refining an element from an almost-good mesh.

As [7,11,4], 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. For example, if v is the circumcenter of a bad d-simplex, and v encroaches some boundary k-simplices then it is rejected. Consider what this definition means in each possible circumstance. 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 a vertex inserted in 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. When v is selected from the picking region of an encroached boundary simplex, it is defined similarly. Notice that if a vertex is considered for insertion but rejected because of an encroachment, its insertion edge length is defined the same way, even though it is never actually inserted. We have to emphasize that e_v is defined just immediately after v is introduced to the mesh; it is not defined based on the final mesh. In other words, if there is a new vertex inserted near v such that the shortest edge connected to v becomes shorter, e_v does not change.

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 non–input 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 [3] and Shewchuk [4] showed that $N()$ defined on the mesh generated by Delaunay refinement method is within a constant factor of $lfs(v)$, i.e., $N(v) \sim lfs(v)$.

We then show that the insertion edge length e_v for any introduced mesh vertex v is related to that of its parent vertex $p(v)$. Notice that here p and v may not be inserted due to encroaching.

Lemma 6.1 [INSERTION EDGE LENGTH] *Let v be a vertex in the final mesh and let $p = p(v)$ be the vertex responsible for the insertion of 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*

(1) $C = \varrho_0$ if v is from the picking region of a d-simplex μ with $\rho(\mu) \geq \varrho_0$;
(2) $C = \frac{1-\delta}{2}$ if v is selected from the picking region of an original sliver-simplex;
(3) $C = \frac{1-\delta}{1+\delta} b$ if v is selected from the picking region of a created sliver-simplex containing a sliver τ and $f(\tau)$ is also a sliver-simplex;
(4) $C = \frac{1-\delta}{2}$ if v is selected from the picking region of a created sliver-simplex containing sliver τ and $f(\tau)$ has a radius-edge ratio $\rho(f(\tau)) \geq \varrho_0$;
(5) $C = \frac{1-\delta}{1+\delta} b$ if v is selected from $\mathcal{P}(\tau)$ of a created sliver-simplex τ and $f(\tau)$ is a boundary k-simplex encroached by a sliver-simplex;
(6) $C = \frac{1-\delta}{2}$ if v is selected from $\mathcal{P}(\tau)$ of a created sliver-simplex τ and $f(\tau)$ is a boundary k-simplex encroached by a d-simplex with large radius-edge ratio;
(7) $C = \frac{1-\delta}{\sqrt{2}(1+\delta)}$ if v is selected from the picking region of a boundary k-simplex encroached directly or indirectly by a sliver-simplex.
(8) $C = \frac{1}{\sqrt{2}}$ if v is selected from the picking region of a boundary k-simplex encroached directly or indirectly by a d-simplex with a large radius-edge ratio.

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, $e_p \le \|p-q\| = L_\mu \le \frac{R_\mu}{\varrho_0}$. Thus $e_v = R_\mu \ge \varrho_0 \cdot e_p$.

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 $e_p \le \|p-q\| = L_\tau$. Notice that $R_\mu \ge R_\tau \ge L_\tau/2$. Thus $e_v \ge (1-\delta)R_\mu \ge \frac{1-\delta}{2} \cdot e_p$.

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, $R_\tau \ge bR_{f(\tau)}$, where $R_{f(\tau)}$ is the circumradius of $f(\tau)$. Notice that $e_p \le (1+\delta)R_{f(\tau)}$. Thus, $e_v \ge (1-\delta)bR_{f(\tau)} \ge \frac{1-\delta}{1+\delta}b \cdot e_p$.

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

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 $R_\tau \ge bR_{f(\tau)}$. Notice that $e_p \le (1+\delta)R_{f(\tau)}$. Thus, $e_v \ge (1-\delta)bR_{f(\tau)} \ge \frac{1-\delta}{1+\delta}b \cdot e_p$. 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 $e_p \le \|p-q\|$. We also have $R_\tau \ge \|p-q\|/2$. Thus, $e_v \ge (1-\delta)R_\mu \ge (1-\delta)R_\tau \ge \frac{1-\delta}{2} \cdot e_p$.

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 $R_\mu \le \sqrt{2}R_\chi$ and $e_p \le (1+\delta)R_\mu$. Thus, $e_v \ge (1-\delta)R_\chi \ge \frac{1-\delta}{\sqrt{2}(1+\delta)} \cdot e_p$.

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 $R_\mu \le \sqrt{2}R_\chi$ and $e_p = R_\mu$. Thus, $e_v = R_\chi \ge R_\mu\sqrt{2} \ge \frac{\sqrt{2}}{2} \cdot e_p$. ◻

The previous lemma 6.1 is concerned about the relationship between the insertion edge length of a child and its parent, if there is any. For a vertex v, as [4], 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. Notice that D_v is defined just immediately after v is introduced to the mesh; it is not defined based on the final mesh. The next lemma will discuss the relationship between D_v and D_p.

Lemma 6.2 [Density Ratio Relations] *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 because of eliminating a d-simplex with a large radius-edge ratio, then $D_v \leq 1 + \frac{D_p}{C}$.*

Proof. We prove the lemma by analyzing different scenarios for defining parent vertex $p(v)$ of v. If v is inserted inside the picking region of a bad simplex μ, p is then on the circumsphere of μ. Thus $e_v \geq (1-\delta)R_\mu$, and $\|v-p\| \leq (1+\delta)R_\mu$, where R_μ is the circumradius of μ. If v is inserted inside the picking region of an encroached boundary simplex μ, then p is inside the circumsphere of μ. Thus $e_v > (1-\delta)R_\mu$, and $\|v-p\| > (1+\delta)R_\mu$, where R_μ is the circumradius of μ. In both cases, we have $\|v-p\| \leq \frac{1+\delta}{1-\delta} e_v$.

From the 1-Lipschitz condition of the local feature size function *lfs*(), we have

$$\begin{aligned} lfs(v) &\leq lfs(p) + \|v-p\| \\ &\leq D_p \cdot e_p + \frac{1+\delta}{1-\delta} e_v \\ &\leq \frac{D_p}{C} e_v + \frac{1+\delta}{1-\delta} e_v. \end{aligned}$$

The lemma follows by dividing both side by e_v. Hereafter, we denote $\alpha = \frac{1+\delta}{1-\delta}$.

Notice that when the point v is inserted beause of eliminating d-simplex with large radius-edge ratio, then $D_v \leq 1 + \frac{D_p}{C}$. ◻

Theorem 6.3 [Bounded Density] *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:*

$$\begin{aligned} D_d &\geq max\{\frac{\varrho_0 A + \sqrt{2}^{d+1} B_{d-2}}{\varrho_0(1-\delta) - 2^d \alpha^{d-1}}, \frac{\alpha b + \alpha^2 + \alpha B_{d-2}}{b - \sqrt{2}^{d-1}\alpha^d}, \alpha + \frac{2}{1-\delta}\}, \\ D_{d-k} &= B_{k-1} + \sqrt{2}^k \alpha^k D_d, \end{aligned}$$

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* $D_v \le D$ *for all mesh vertex* v.

PROOF. We prove the theorem by induction. First consider any original input vertex p, the length e_p of the shortest edge connected to p is at least $lfs(p)$ from the definition of $lfs(p)$. Thus $D_p = \frac{lfs(p)}{e_p} \le 1$. Then assume that the lemma is true for the parent vertex p of vertex v. Hereafter, let $\alpha = \frac{1+\delta}{1-\delta}$.

If v is selected from the picking region of a d-simplex μ with $R_\mu/L_\mu \ge \varrho_0$, then $e_v \ge \varrho_0 \cdot e_p$. Therefore, by above Lemma 6.3, we have $D_v \le 1 + D_p/\varrho_0$. Notice that here we have $D_p \le D$. Thus a sufficient condition that one can prove that $D_v \le D_d$ is

$$1 + \frac{D}{\varrho_0} \le D_d \tag{1}$$

If v is selected from the picking region of a sliver-simplex μ containing an original k-sliver τ, then $e_v \ge \frac{1-\delta}{2} \cdot e_p$. Here p is an original input vertex, i.e., $D_p \le 1$. By above Lemma 6.3, we have

$$D_v \le \alpha + \frac{D_p}{\frac{1-\delta}{2}}.$$

Thus a sufficient condition that one can prove that $D_v \le D_d$ is

$$\alpha + \frac{2}{1-\delta} \le D_d. \tag{2}$$

Then consider that v is selected from the picking region of a created sliver-simplex μ containing a created k-sliver τ. There are three cases about the parent element $f(\tau)$ of τ. $f(\tau)$ is a d-dimensional sliver-simplex, $f(\tau)$ is a d-simplex with large radius-edge ratio, and $f(\tau)$ is an encroached m-dimensional boundary simplex.

If parent element $f(\tau)$ is a sliver-simplex, then $e_v \ge \frac{1-\delta}{1+\delta}b \cdot e_p$. Thus, we have $D_v \le \frac{1+\delta}{1-\delta} + \frac{D_p}{\frac{1-\delta}{1+\delta}b}$. Here, $D_p \le D_d$ by induction. Thus a sufficient condition that one can prove that $D_v \le D_m$ is

$$\alpha + \alpha\frac{D_d}{b} \le D_d. \tag{3}$$

If parent element $f(\tau)$ is a d-simplex with a large radius-edge ratio, then $e_v \ge \frac{1-\delta}{2} \cdot e_p$. The fact that p is inserted from $f(\tau)$ implies that $e_p \ge \varrho_0 \cdot e_{p(p)}$, where $p(p)$ is the parent vertex of p. Then Lemma 6.2 implies that $D_v \le \frac{1+\delta}{1-\delta} + \frac{D_p}{\frac{1-\delta}{2}}$, and $D_p \leq 1 + \frac{D_{p(p)}}{\varrho_0}$. Notice that $D_{p(p)} \leq D$. Thus a sufficient condition to prove $D_v \leq D_d$ is

$$\alpha + \frac{2}{1-\delta} + \frac{2}{1-\delta}\frac{D}{\varrho_0} \leq D_d. \tag{4}$$

We then study that the parent element $f(\tau)$ is an encroached boundary k-simplex. We first consider the scenario that $f(\tau)$ is encroached by a sliver-simplex directly or indirectly. Thus we have $e_v \geq \frac{1-\delta}{1+\delta} b \cdot e_p$. Therefore, we have $D_v \leq \frac{1+\delta}{1-\delta} + \frac{D_p}{\frac{1-\delta}{1+\delta} b}$. Notice that parent vertex p is from a k-dimensional boundary simplex. Then a sufficient condition to prove $D_v \leq D_d$ is, for all $1 \leq k \leq d$,

$$\alpha + \alpha\frac{D_k}{b} \leq D_d. \tag{5}$$

We then consider the scenario that $f(\tau)$ is encroached by a d-simplex μ with large radius-edge ratio directly or indirectly. Thus we have $e_v \geq \frac{1-\delta}{2} \cdot e_p$. If $f(\tau)$ is encroached directly by μ, then $e_p \geq \frac{1}{\sqrt{2}} \cdot e_{c_\mu}$, where c_μ is the circumcenter of μ, i.e., the parent vertex of p. If $f(\tau)$ is encroached by μ indirectly, then there is a sequence of simplices μ_0, μ_1, ..., μ_t in different dimensions such that μ_i encroaches simplex μ_{i+1} and $\mu_0 = \mu$, $\mu_t = f(\tau)$. Let p_i be the circumcenter of simplex μ_i, $i = 0, 1, \ldots, t$. Then we have $e_{p_i} \geq \frac{1}{\sqrt{2}} \cdot e_{p_{i-1}}$. Because μ has a large radius-edge ratio, we have $e_{p_0} \geq \varrho_0 e_{p(p_0)}$. By applying Lemma 6.2, we have $D_v \leq \frac{1+\delta}{1-\delta} + \frac{D_{p_t}}{\frac{1-\delta}{2}}$, $D_{p_i} \leq 1 + \frac{D_{p_{i-1}}}{\frac{1}{\sqrt{2}}}$ for $1 \leq i \leq t \leq d-1$, and $D_{p_0} \leq 1 + \frac{D_{p(p_0)}}{\varrho_0}$. Notice that vertex $p_0 = c_\mu$ and its parent $p(p_0)$ could be from interior or on the boundary. Then a sufficient condition to prove $D_v \leq D_d$ is

$$\alpha + \frac{\sqrt{2}^{d+2} - 2}{(\sqrt{2}-1)(1-\delta)} + \frac{\sqrt{2}^{d+1}}{1-\delta}\frac{D}{\varrho_0} \leq D_d. \tag{6}$$

If v is selected from the picking region of a boundary m-simplex χ encroached by a sliver directly or indirectly, then $e_v \geq \frac{1-\delta}{\sqrt{2}(1+\delta)} \cdot e_p$. Here parent p is the circumcenter of a k-dimensional simplex μ, where $k > m$. Thus a sufficient condition that one can prove that $D_v \leq D_m$ is, for all $1 \leq m < k \leq d$,

$$\alpha + \sqrt{2}\alpha D_k \leq D_m. \tag{7}$$

Then consider that v is selected from the picking region of a boundary m-simplex χ encroached by a d-simplex with a large radius-edge ratio directly or indirectly. Assume that τ is encroached by a k-simplex μ, where $k > m$. We have $e_v \geq \frac{1}{\sqrt{2}} \cdot e_p$. Here parent p is the circumcenter of μ. Thus a sufficient condition that one can prove that $D_v \le D_m$ is, for all $1 \le m < k \le d$

$$1 + \sqrt{2} D_k \le D_m. \tag{8}$$

Notice that some inequalities are satisfied if other inequalities were satisfied. One can show that the above inequalities are simultaneously satisfied if we choose

$$\begin{aligned}
D_d &\ge \frac{\varrho_0 A + \sqrt{2}^{d+1} B_{d-2}}{\varrho_0(1-\delta) - 2^d \alpha^{d-1}} \\
D_d &\ge \frac{\alpha b + \alpha^2 + \alpha B_{d-2}}{b - \sqrt{2}^{d-1} \alpha^d}, \\
D_d &\ge \alpha + \frac{2}{1-\delta}, \\
D_{d-k} &= B_{k-1} + \sqrt{2}^k \alpha^k D_d,
\end{aligned}$$

where $1 \le k \le d-1$, $A = \alpha(1-\delta) + \frac{\sqrt{2}^{d+1}-2}{\sqrt{2}-1}$, and $B_i = \sum_{j=1}^{i}(\sqrt{2}^j \alpha^{j+1})$. ◻

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 $lfs(v)$. Ruppert and Shewchuk had similar theorems for classic Delaunay refinement methods; see [3,4].

Theorem 6.4 [Good Grading] *For any mesh vertex v generated by refinement algorithm, the distance connected to its nearest neighbor vertex u is at least $\frac{lfs(v)}{D+1}$.*

Proof. Similarly result for 3-dimensional mesh has been shown by various reseachers. From previous Lemma 6.3, we know that $D_v = \frac{lfs(v)}{e_v} \le D$ is true for any mesh vertex v. There are two cases about the order of inserting v and its nearest neighbor vertex u. If u is inserted before v, the shortest edge connected to v is not changed after inserting v. Thus

$$\|v - u\| = e_v \ge \frac{lfs(v)}{D}.$$

Otherwise, u is inserted after v, then the shortest edge connected to v is decreased after inserting v. We then study vertex u. First $D_u \le D$ implies that $e_u \ge \frac{lfs(u)}{D}$. In addition, $\|v - u\| \ge e_u$ from the definition of e_u. From the 1-Lipschitz condition of $lfs()$, we have $lfs(v) \le lfs(u) + \|v - u\|$. So

$$\begin{aligned}\|v-u\| &\geq \frac{lfs(u)}{D} \\ &\geq \frac{lfs(v)-\|v-u\|}{D}.\end{aligned}$$

It follows that $\|v-u\| \geq \frac{lfs(v)}{1+D}$. $\square$

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 6.4 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 6.5 [Linear Size] *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] G. L. Miller, D. Talmor, S. H. Teng, N. Walkington, 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.

[2] L. P. Chew, Constrained Delaunay triangulations, Algorithmica 4 (1989) 97–108.

[3] J. Ruppert, A new and simple algorithm for quality 2-dimensional mesh generation, in: Third Annual ACM-SIAM Symposium on Discrete Algorithms, 1992, pp. 83–92.

[4] J. R. Shewchuk, Delaunay refinement mesh generation, Ph.D. thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania (1997).

[5] M. Bern, P. Chew, D. Eppstein, J. Ruppert, Dihedral bounds for mesh generation in high dimensions, in: Proc. 6th ACM-SIAM Symposium on Discrete Algorithms, ACM, 1995, pp. 189–196.

[6] L. P. Chew, Guaranteed-quality delaunay meshing in 3d (short version), in: 13th ACM Sym. on Comp. Geometry, 1997, pp. 391–393.

[7] X. Y. Li, Sliver-free three dimensional delaunay mesh generation, Ph.D. thesis, University of Illinois at Urbana-Champaign (2000).

[8] J. C. Cavendish, D. A. Field, W. H. Frey, An approach to automatic three-dimensional finite element mesh generation, Internat. J. Numer. Methods Engrg. 21.

[9] H. Edelsbrunner, N. R. Shah, Incremental topological flipping works for regular triangulations, Algorithmica 15.

[10] H. Edelsbrunner, X. Y. Li, G. Miller, A. Stathopoulos, D. Talmor, S. H. Teng, A. ÜNgÖR", N. Walkington, Smoothing and cleaning up slivers, in: ACM Symposium on Theory of Computing (STOC00), 2000.

[11] X. Y. Li, S. H. Teng, Sliver-free three dimensional delaunay mesh generation, in: Symposium on Discrete Algorithms (SODA), 2001.

[12] S. W. Cheng, T. K. Dey, H. Edelsbrunner, M. A. Facello, S. H. Teng, Silver exudation, in: Proc. 15th ACM Symposium on Computational Geometry, 1999, pp. 1–13.

[13] D. Talmor, Well-spaced points for numerical methods, Ph.D. thesis, Carnegie Mellon University (1997).

[14] J. R. Shewchuk, Tetrahedral mesh generation by delaunay refinement, in: 14th Annual ACM Symposium on Computational Geometry, 1998, pp. 86–95.