

1A.2

Three-dimensional Semi-generalized Point Placement Method for Delaunay Mesh Refinement

Andrey N. Chernikov and Nikos P. Chrisochoides

Department of Computer Science
College of William and Mary
Williamsburg, VA 23185
{[ancher](mailto:ancher@cs.wm.edu),[nikos](mailto:nikos@cs.wm.edu)}@cs.wm.edu

Summary. A number of approaches have been suggested for the selection of the positions of Steiner points in Delaunay mesh refinement. In particular, one can define an entire region (called picking region or selection disk) inside the circumscribed sphere of a poor quality element such that any point can be chosen for insertion from this region. The two main results which accompany most of the point selection schemes, including those based on regions, are the proof of termination of the algorithm and the proof of good gradation of the elements in the final mesh. In this paper we show that in order to satisfy only the termination requirement, one can use larger selection disks and benefit from the additional flexibility in choosing the Steiner points. However, if one needs to keep the theoretical guarantees on good grading then the size of the selection disk needs to be smaller. We introduce two types of selection disks to satisfy each of these two goals and prove the corresponding results on termination and good grading first in two dimensions and then in three dimensions using the radius-edge ratio as a measure of element quality. We call the point placement method semi-generalized because the selection disks are defined only for mesh entities of the highest dimension (triangles in two dimensions and tetrahedra in three dimensions); we plan to extend these ideas to lower-dimensional entities in the future work. We implemented the use of both two- and three-dimensional selection disks into the available Delaunay refinement libraries and present one example (out of many choices) of a point placement method; to the best of our knowledge, this is the first implementation of Delaunay refinement with point insertion at any point of the selection disks (picking regions).

1 Introduction

Delaunay mesh generation algorithms start with the construction of the initial mesh, which conforms to the input geometry, and then refine this mesh until the element quality constraints are met. The general idea of Delaunay refinement is to insert additional (Steiner) points inside the circumdisks of

poor quality elements, which causes these elements to be destroyed, until they are gradually eliminated and replaced by better quality elements. Traditionally, Steiner points are selected at the circumcenters of poor quality elements [3, 6, 14, 16, 17]. Ruppert [14] and Shewchuk [17] proved that if the radius-edge upper bound $\bar{\rho}$ is greater than or equal to $\sqrt{2}$ in two dimensions and $\bar{\rho} \geq 2$ in three dimensions, then Delaunay refinement with circumcenters terminates. If, furthermore, the inequalities are strict, then good grading can also be proven both in two and in three dimensions. In two dimensions, in addition to good grading, one can also prove size-optimality which refers to the fact that the number of triangles in the resulting mesh is within a constant multiple of the smallest possible number of triangles satisfying the input constraints.

Recently, Üngör [20] introduced a new type of Steiner points called *off-centers*. The idea is based on the observation that sometimes the elements created as a result of inserting circumcenters of skinny triangles are also skinny and require further refinement. This technique combines the advantages of advancing front methods, which can produce very well-shaped elements in practice, and Delaunay methods, which offer theoretical guarantees. Üngör showed that the use of off-centers allows to significantly decrease the size of the final mesh in practice. While eliminating additional point insertions, this strategy creates triangles with the longest possible edges, such that one can prove the termination of the algorithm and still produce a graded mesh.

Chew [4] chooses Steiner points randomly inside a *picking sphere* to avoid slivers. The meshes in [4] have constant density; therefore, Chew proves the termination, but not the good grading.

Li and Teng [9, 10] extended the work in [4] by defining a picking sphere with a radius which is a constant multiple of the circumradius of the element. They use two different rules for eliminating the elements with large radius-edge ratio and for eliminating the slivers. In particular, in [9] the rules are defined as follows: “Add the circumcenter c_τ of any d -simplex with a large $\rho(\tau)$ ” and “For a sliver-simplex τ , add a good point $p \in \mathcal{P}(\tau)$ ”, where $\rho(\tau)$ is the radius-edge ratio, $\mathcal{P}(\tau)$ is the picking region of simplex τ , and the good point is found by a constant number of random probes. The authors in [9] prove that their algorithm terminates and produces a well graded mesh with good radius-edge ratio and without slivers.

In the present work, we define Type-2 selection disks similarly to the picking region in [9]. We extend the proofs in [9] to show that any point (not only the circumcenter) from the selection disk (picking region) can be used to eliminate the elements with large radius-edge ratios. We do not address the problem of sliver elimination, however, our work can be used in conjunction with the sliver removal procedure from [9] such that the Delaunay refinement algorithm can choose any points from the selection disks (picking regions) throughout both the stage of the construction of a good radius-edge ratio mesh (“almost good mesh” [9]) and the stage of sliver removal. Intuitively, the requirement of good grading has to be more restrictive than the requirement

of termination only, and, therefore, the definitions of the selection disk has to be different to satisfy each of these goals. The present paper improves upon our previous result [2] by decoupling the definitions of the selection disk used for the proof of termination (Type-1) and the selection disk used for the proof of good grading (Type-2). As can be seen further in the paper, the selection disk of Type-2 is always inside the selection disk of Type-1 of the same element, and as the radius-edge ratio ρ of an element approaches the upper bound $\bar{\rho}$, the Type-2 disk approaches the Type-1 disk.

The traditional proofs of termination and size optimality of Delaunay refinement algorithms [14, 17] explore the relationships between the insertion radius of a point and that of its parent. Stated shortly, the *insertion radius* of point p is the length of the shortest edge connected to p immediately after p is inserted into the mesh, and the *parent* is the vertex which is “responsible” for the insertion of p [17]. The proofs in [14, 17] rely on the assumption that the insertion radius of a Steiner point is equal to the circumradius of the poor quality element. This assumption holds when the Steiner point is chosen at the circumcenter of the element, since by Delaunay property the circumdisk of the element is empty, and, hence, there is no vertex closer to the circumcenter than its vertices. However, the above assumption does not hold if we pick an arbitrary point inside the selection disk of the element. For example, in Figure 2(right) the Steiner point p_i within the selection disk can be closer to the mesh vertex p_n than to any of the vertices p_k, p_l, p_m which define the skinny triangle. Therefore, we need to extend the existing theory in the new context, i.e., Steiner points can be inserted anywhere within the selection disks of poor quality elements.

One of the important applications of the flexibility offered by the use of selection disks is in conforming the mesh to the boundary between different materials. The advantages are especially pronounced in medical imaging, when the boundaries between different tissues are blurred, see Figure 1(left). In this case, after the image is segmented, instead of a clear separation, we have a boundary zone of some none-negligible width, see Figure 1(right). Then the goal of the mesh generation step is to avoid creating edges that would intersect the boundary, which can be achieved by inserting Steiner points inside the boundary zone. Another use for the selection disks would be to put vertices on a point lattice to reduce the occurrence of slivers, see [8].

In [2] we presented an example of a two-dimensional optimization-based method which allows to decrease the size of the final mesh in practice. Here we present a preliminary evaluation of an analogous three-dimensional method. The three-dimensional results are not yet as good as two-dimensional and require additional development.




Fig. 1. The use of the selection disk for the construction of boundary conformal meshes. **(Left)** An MRI scan showing a cross-section of a body. **(Right)** A zoom-in of the selected area containing an artery: the inside is white, the outside has different shades of gray and the black zone is an approximate boundary between these regions. The standard Delaunay refinement algorithm would insert the circumcenter c . However, in order to construct a mesh which conforms to the boundary, another point (p) would be a better choice.

2 Delaunay Refinement Background

We assume that the input domain Ω is described by a Planar Straight Line Graph (PSLG) in two dimensions, or a Planar Linear Complex (PLC) in three dimensions [14–17]. A PSLG (PLC) \mathcal{X} consists of a set of vertices, a set of straight line segments, and (in three dimensions) a set of planar facets. Each element of \mathcal{X} is considered *constrained* and must be preserved during the construction of the mesh, although it can be subdivided into smaller elements. The vertices of \mathcal{X} must be a subset of the final set of vertices in the mesh.

Let the mesh $\mathcal{M}_{\mathcal{X}}$ for the given PSLG (PLC) \mathcal{X} consist of a set $V = \{p_i\}$ of vertices and a set $T = \{t_i\}$ of elements which connect vertices from V . The elements are either triangles in two dimensions or tetrahedra in three dimensions. We will denote the triangle with vertices p_u , p_v , and p_w as $\triangle(p_u p_v p_w)$ and the tetrahedron with vertices p_k , p_l , p_m , and p_n as $\tau(p_k p_l p_m p_n)$. We will use the symbol $e(p_i p_j)$ to represent the edge of the mesh which connects points p_i and p_j .

As a measure of the quality of elements we will use the circumradius-to-shortest edge ratio specified by an upper bound $\bar{\rho}$, which in two dimensions is equivalent to a lower bound on a minimal angle [11, 15] since for a triangle with the circumradius-to-shortest edge ratio ρ and minimal angle A , $\rho = 1/(2 \sin A)$. We will denote the circumradius-to-shortest edge ratio of element t as $\rho(t)$.

Let us call the open disk corresponding to a triangle’s circumscribed circle or to a tetrahedron’s circumscribed sphere its *circumdisk*. We will use symbols $\circ(t)$ and $r(t)$ to represent the circumdisk and the circumradius of t , respectively. A mesh is said to satisfy the *Delaunay property* if the circumdisk of every element does not contain any of the mesh vertices [5].




Fig. 2. (Left) Encroachment in three dimensions. (Right) Selection disks of Type-1 (light shade) and Type-2 (dark shade) of the skinny triangle $\Delta(p_k p_l p_m)$.

We will extensively use the notion of *cavity* [7] which is the set of elements in the mesh whose circumdisks include a given point p . We will denote $\mathcal{C}_{\mathcal{M}}(p)$ to be the cavity of p with respect to mesh \mathcal{M} and $\partial\mathcal{C}_{\mathcal{M}}(p)$ to be the set of boundary edges in two dimensions (or triangles in three dimensions) of the cavity, i.e., the edges or the triangles which belong to only one element in $\mathcal{C}_{\mathcal{M}}(p)$. When \mathcal{M} is clear from the context, we will omit the subscript. For our analysis, we will use the Bowyer-Watson (B-W) point insertion algorithm [1, 22], which can be written shortly as follows:

$$\begin{aligned} V^{n+1} &\leftarrow V^n \cup \{p\}, \\ T^{n+1} &\leftarrow T^n \setminus \mathcal{C}_{\mathcal{M}^n}(p) \cup \{(p\xi) \mid \xi \in \partial\mathcal{C}_{\mathcal{M}^n}(p)\}, \end{aligned} \quad (1)$$

where ξ is an edge in two dimensions or a triangle in three dimensions, while $\mathcal{M}^{n+1} = (V^{n+1}, T^{n+1})$ and $\mathcal{M}^n = (V^n, T^n)$ represent the mesh before and after the insertion of p , respectively.

In order not to create skinny elements close to the constrained segments and faces, sequential Delaunay algorithms observe special *encroachment* rules [14–17]. In particular, if a Steiner point p is considered for insertion but it lies within the open equatorial disk of a constrained subfacet f , p is not inserted but the circumcenter of f is inserted instead. Similarly, if p is inside the open diametral circle of a constrained subsegment s , then the midpoint of s is inserted instead. Consider the example in Figure 2(left). The new point p_i is inside the three-dimensional equatorial disk of a constrained face $\Delta(p_k p_l p_m)$. In this case, p_i is rejected and the algorithm attempts to insert the circumcenter p_i' of $\Delta(p_k p_l p_m)$. If p_i' does not encroach upon any constrained segments, it is inserted into the mesh. If, however, it encroaches upon a constrained segment, which is $e(p_k p_l)$ in our example, p_i' is also rejected and the midpoint p_i'' of the constrained edge is inserted.

Definition 1 (Local feature size [14, 16, 17]). *The local feature size function $\text{lfs}(p)$ for a given point p is equal to the radius of the smallest disk centered at p that intersects two non-incident elements of the PSLG (PLC).*

$\text{lfs}(p)$ satisfies the Lipschitz condition:

Lemma 1 (Lemma 1 in [14], Lemma 2 in [17], Lemma 2 in [16]). *Given any PSLG (PLC) and any two points p_i and p_j , the following inequality holds:*

$$\text{lfs}(p_i) \leq \text{lfs}(p_j) + \|p_i - p_j\|. \quad (2)$$

Here and in the rest of the paper we use the standard Euclidean norm $\|\cdot\|$.

The following definitions of insertion radius and parent of a Steiner point play a central role in the analysis in [14, 16, 17] and we will adopt them for our analysis, too.

Definition 2 (Insertion radius [16, 17]). *The insertion radius $R(p)$ of point p is the length of the shortest edge which would be connected to p if p is inserted into the mesh, immediately after it is inserted. If p is an input vertex, then $R(p)$ is the Euclidean distance between p and the nearest input vertex visible from p .*

Here a vertex is called visible from another vertex if the straight line segment connecting both vertices does not intersect any of the constrained segments.

Remark 1. If p is a midpoint of an encroached subsegment or subfacet, then $R(p)$ is the distance between p and the nearest encroaching mesh vertex; if the encroaching mesh vertex was rejected for insertion, then $R(p)$ is the radius of the diametral circle of the subsegment or of the equatorial sphere of the subfacet [16, 17].

Remark 2. As shown in [16, 17], if p is an input vertex, then $R(p) \geq \text{lfs}(p)$. Indeed, from the definition of $\text{lfs}(p)$, the second feature (in addition to p) which intersects the disk centered at p is either a constrained segment, a constrained facet, or the nearest input vertex visible from p .

The following definition of a parent vertex generalizes the corresponding definition in [16, 17]. In our analysis, even though the child is not necessarily the circumcenter, the parent is still defined to be the same vertex.

Definition 3 (Parent of a Steiner point). *The parent \hat{p} of point p is the vertex which is defined as follows: (i) If p is an input vertex, it has no parent. (ii) If p is a midpoint of an encroached subsegment or subfacet, then \hat{p} is the encroaching point (possibly rejected for insertion). (iii) If p is inserted inside the circumdisk of a poor quality element (triangle or tetrahedron), \hat{p} is the most recently inserted vertex of the shortest edge of this element.*

The quantity $D(p)$ is defined as the ratio of $\text{lfs}(p)$ over $R(p)$ [16, 17]:

$$D(p) = \frac{\text{lfs}(p)}{R(p)}. \quad (3)$$

It reflects the density of vertices near p at the time p is inserted, weighted by the local feature size. To achieve good mesh grading we would like this density to be as small as possible. If the density is bounded from above by a constant, the mesh is said to have a *good grading* property.

3 Two-Dimensional Generalized Delaunay Refinement

In this section we introduce two types of selection disks which can be used for the insertion of Steiner points in two dimensions. The two-dimensional results are given here for completeness and prepare the mindset for the three-dimensional results in the following section. The Type-1 selection disk corresponds to the definition of the selection disk we used in [2] which we briefly review. Then we define the Type-2 disk and prove good grading and size-optimality.

Definition 4 (Selection disk of Type-1 in 2D). *If t is a poor quality triangle with circumcenter c , shortest edge length l , circumradius r , and circumradius-to-shortest edge ratio $\rho = r/l > \bar{\rho} \geq \sqrt{2}$, then the Type-1 selection disk for the insertion of a Steiner point that would eliminate t is the open disk with center c and radius $r - \sqrt{2}l$.*

For example, in Figure 2(right), $e(p_l p_m)$ is the shortest edge of a skinny triangle $\triangle(p_k p_l p_m)$ and c is its circumcenter. The selection disk of Type-1 is the lightly shaded disk with center c and radius $r(\triangle(p_k p_l p_m)) - \sqrt{2} \|p_l - p_m\|$.

Remark 3. As $\rho(\triangle(p_k p_l p_m))$ approaches $\sqrt{2}$, the Type-1 selection disk shrinks to the circumcenter c of the triangle. If, furthermore, $\rho(\triangle(p_k p_l p_m)) \leq \sqrt{2}$, the selection disk vanishes, which coincides with the fact that the triangle $\triangle(p_k p_l p_m)$ cannot be considered skinny.

In [2] we proved that any point inside the Type-1 selection disk of a triangle can be chosen for the elimination of the triangle, and that the generalized Delaunay refinement algorithm which chooses Steiner points inside Type-1 selection disks terminates.

Using the definition of the selection disk, in [2] we suggested an example of an optimization-based method which allows to improve the size of the mesh by up to 20% over the circumcenter insertion method and up to 5% over the off-center insertion method, for small values of the minimal angle bound. The underlying idea of our method is that, by choosing a point within the selection disk we can vary the set of triangles in its cavity, we simultaneously minimize the number of deleted good quality triangles and maximize the number of deleted poor quality triangles.

3.1 Proof of Good Grading and Size Optimality with Selection Disks of Type II

Definition 5 (Selection disk of Type-2 in 2D). *If t is a poor quality triangle with circumcenter c , shortest edge length l , circumradius r , and circumradius-to-shortest edge ratio $\rho = r/l > \bar{\rho} \geq \sqrt{2}$, then the Type-2 selection disk for the insertion of a Steiner point that would eliminate t is the open disk with center c and radius $r(1 - \frac{\sqrt{2}}{\bar{\rho}})$.*

For example, in Figure 2(right), $e(p_l p_m)$ is the shortest edge of a skinny triangle $\triangle(p_k p_l p_m)$ and c is its circumcenter. The Type-2 selection disk for this triangle is the darkly shaded disk with center c and radius $r(\triangle(p_k p_l p_m))(1 - \frac{\sqrt{2}}{\bar{\rho}})$.

Remark 4. Note from Definitions 4 and 5 that the radius of the Type-2 selection disk is always smaller than the radius of the Type-1 selection disk of the same skinny triangle because $r(1 - \frac{\sqrt{2}}{\bar{\rho}}) = r - \frac{r}{\bar{\rho}}\sqrt{2}l$ and $\rho > \bar{\rho}$.

Remark 5. As $\bar{\rho}$ approaches $\sqrt{2}$ the radius of the Type-2 selection disk approaches zero, which means that the selection disk shrinks to the circumcenter point.

As can be seen further below, the price which we pay for the gain in the flexibility in choosing points is the increase of the constants which bound the size of the mesh. To classify the Delaunay refinement algorithms with respect to the theoretical bounds on mesh grading we need the following definition.

Definition 6 (δ -graded Delaunay refinement algorithm in 2D). *If for every triangle t with circumcenter c , circumradius r , shortest edge length l , and circumradius-to-shortest edge length ratio $\rho = r/l > \bar{\rho} \geq \sqrt{2}$ a Delaunay refinement algorithm selects a Steiner point p_i within the Type-2 selection disk such that $\|p_i - c\| < r(1 - \delta)$, where*

$$\frac{\sqrt{2}}{\bar{\rho}} \leq \delta \leq 1,$$

we say that this Delaunay refinement algorithm is δ -graded.

Lemma 2. *If p_i is a vertex of the mesh produced by a δ -graded Delaunay refinement algorithm then the following inequality holds:*

$$R(p_i) \geq C \cdot R(\hat{p}_i), \tag{4}$$

where we distinguish among the following cases:

- (i) $C = \delta\bar{\rho}$ if p_i is a Steiner point chosen within the Type-2 selection disk of a skinny triangle;

Otherwise, let p_i be the midpoint of subsegment s . Then

- (ii) $C = \frac{1}{\sqrt{2}}$ if \hat{p}_i is a Steiner point which encroaches upon s , chosen within the Type-2 selection disk of a skinny triangle;
- (iii) $C = \frac{1}{2\cos\alpha}$ if p_i and \hat{p}_i lie on incident subsegments separated by an angle of α (with \hat{p}_i encroaching upon s), where $45^\circ \leq \alpha \leq 90^\circ$;
- (iv) $C = \sin\alpha$ if p_i and \hat{p}_i lie on incident segments separated by an angle of $\alpha \leq 45^\circ$.

If p_i is an input vertex or if p_i and \hat{p}_i are on non-incident features, then

$$R(p_i) \geq \text{lfs}(p_i).$$

Proof. We need to present a new proof only for case (i) since the proof for case (ii) is the same as in the corresponding lemma for the Type-1 selection disk [2], and the proofs for all other cases are independent of the choice of the point within the selection disk and are given in [17].

Case (i) By the definition of a parent vertex, \hat{p}_i is the most recently inserted endpoint of the shortest edge of the triangle; without loss of generality let $\hat{p}_i = p_l$ and $e(p_l p_m)$ be the shortest edge of the skinny triangle $\Delta(p_k p_l p_m)$, see Figure 2(right). If $e(p_l p_m)$ was the shortest edge among the edges incident upon p_l at the time p_l was inserted into the mesh, then $\|p_l - p_m\| = R(p_l)$ by the definition of the insertion radius; otherwise, $\|p_l - p_m\| \geq R(p_l)$. In either case,

$$\|p_l - p_m\| \geq R(p_l). \quad (5)$$

Then

$$\begin{aligned} R(p_i) &> \delta r && \text{(from Delaunay property and Definition 6)} \\ &= \delta \rho \|p_l - p_m\| && \text{(since } \rho = \frac{r}{\|p_l - p_m\|} \text{)} \\ &> \delta \bar{\rho} \|p_l - p_m\| && \text{(since } \rho > \bar{\rho} \text{)} \\ &\geq \delta \bar{\rho} R(p_l) && \text{(from (5)).} \end{aligned}$$

Hence, $R(p_i) > \delta \bar{\rho} R(\hat{p}_i)$; choose $C = \delta \bar{\rho}$. \square

Remark 6. We have proven Inequality (4) only for the case which involves the “free” Steiner points, i.e., the points which are chosen from selection disks and which do not lie on constrained segments. The proofs for the other types of Steiner points do not involve the definition of the selection disk and are applicable without change from [17].

Lemma 3. *If p is a vertex of the mesh produced by a δ -graded Delaunay refinement algorithm and C is the constant specified by Lemma 2, then the following inequality holds:*

$$D(p) \leq \frac{2-\delta}{\delta} + \frac{D(\hat{p})}{C}. \quad (6)$$

Proof. If p is inside a Type-2 selection disk of a skinny triangle with circum-radius r , then

$$\begin{aligned} \|p - \hat{p}\| &< 2r - \delta r && \text{(from the definition of } \hat{p} \text{ and Def. 6)} \\ &= (2 - \delta)r \\ &= \frac{2-\delta}{\delta} \delta r \\ &< \frac{2-\delta}{\delta} R(p) && \text{(from Delaunay property and Def. 6).} \end{aligned}$$

If p is an input vertex or lies on an encroached segment, then

$$\begin{aligned} \|p - \hat{p}\| &\leq R(p) && \text{(by definitions of } \hat{p} \text{ and } R(p)) \\ &\leq \frac{2-\delta}{\delta} R(p) && \text{(since from Def. 6, } \delta \leq 1). \end{aligned}$$

In all cases,

$$\|p - \hat{p}\| \leq \frac{2-\delta}{\delta} R(p). \quad (7)$$

Then

$$\begin{aligned} \text{lfs}(p) &\leq \text{lfs}(\hat{p}) + \|p - \hat{p}\| && \text{(from Lemma 1)} \\ &\leq \text{lfs}(\hat{p}) + \frac{2-\delta}{\delta} R(p) && \text{(from (7))} \\ &= D(\hat{p}) R(\hat{p}) + \frac{2-\delta}{\delta} R(p) && \text{(from (3))} \\ &\leq D(\hat{p}) \frac{R(p)}{C} + \frac{2-\delta}{\delta} R(p) && \text{(from Lemma 2).} \end{aligned}$$

The result follows from the division of both sides by $R(p)$. \square

Lemma 4 (Extension of Lemma 7 in [17] and Lemma 2 in [14]). *Suppose that $\bar{\rho} > \sqrt{2}$ and the smallest angle in the input PSLG is strictly greater than 60° . There exist fixed constants C_T and C_S such that, for any vertex p inserted (or considered for insertion and rejected) by a δ -graded Delaunay refinement algorithm, $D(p) \leq C_T$, and for any vertex p inserted at the midpoint of an encroached subsegment, $D(p) \leq C_S$. Hence, the insertion radius of a vertex has a lower bound proportional to its local feature size.*

Proof. The proof is by induction and is similar to the proof of Lemma 7 in [17]. The base case covers the input vertices, and the inductive step covers the other two types of vertices: free vertices and subsegment midpoints.

Base case: The lemma is true if p is an input vertex, because in this case, by Remark 2, $D(p) = \text{lfs}(p) / R(p) \leq 1$.

Inductive hypothesis: Assume that the lemma is true for \hat{p} , i.e., $D(\hat{p}) \leq \max\{C_T, C_S\}$.

Inductive step: There are two cases:

(i) If p is in the Type-2 selection disk of a skinny triangle, then

$$\begin{aligned} D(p) &\leq \frac{2-\delta}{\delta} + \frac{D(\hat{p})}{C} && \text{(from Lemma 3)} \\ &= \frac{2-\delta}{\delta} + \frac{D(\hat{p})}{\delta \bar{\rho}} && \text{(from Lemma 2)} \\ &\leq \frac{2-\delta}{\delta} + \frac{\max\{C_T, C_S\}}{\delta \bar{\rho}} && \text{(by the inductive hypothesis).} \end{aligned}$$

It follows that one can prove that $D(p) \leq C_T$ if C_T is chosen so that

$$\frac{2 - \delta}{\delta} + \frac{\max\{C_T, C_S\}}{\delta\bar{\rho}} \leq C_T. \quad (8)$$

(ii) If p is the midpoint of a subsegment s such that \hat{p} encroaches upon s , then we have 3 sub-cases:

(ii-a) If \hat{p} is an input vertex, then the disk centered at p and touching \hat{p} has radius less than the radius of the diametral disk of s and therefore $\text{lfs}(p) < R(p)$. Thus, $D(p) < 1$ and the lemma holds.

(ii-b) If \hat{p} is a rejected point from the Type-2 selection disk of a skinny triangle or lies on a segment not incident to s , then

$$\begin{aligned} D(p) &\leq \frac{2-\delta}{\delta} + \frac{D(\hat{p})}{C} && \text{(from Lemma 3)} \\ &= \frac{2-\delta}{\delta} + \sqrt{2}D(\hat{p}) && \text{(from Lemma 2)} \\ &\leq \frac{2-\delta}{\delta} + \sqrt{2}C_T && \text{(by the inductive hypothesis).} \end{aligned}$$

(ii-c) If \hat{p} lies on a segment incident to s , then

$$\begin{aligned} D(p) &\leq \frac{2-\delta}{\delta} + \frac{D(\hat{p})}{C} && \text{(from Lemma 3)} \\ &= \frac{2-\delta}{\delta} + 2 \cos \alpha D(\hat{p}) && \text{(from Lemma 2)} \\ &\leq \frac{2-\delta}{\delta} + 2C_S \cos \alpha && \text{(by the inductive hypothesis).} \end{aligned}$$

It follows that one can prove that $D(p) \leq C_S$ if C_S is chosen so that both of the following relations (9) and (10) are satisfied:

$$\frac{2 - \delta}{\delta} + \sqrt{2}C_T \leq C_S, \quad (9)$$

$$\frac{2 - \delta}{\delta} + 2C_S \cos \alpha \leq C_S. \quad (10)$$

If $\delta\bar{\rho} > \sqrt{2}$, relations (8) and (9) can be simultaneously satisfied by choosing

$$C_T = \frac{(2 - \delta)(\bar{\rho} + \delta)}{\delta\bar{\rho} - \sqrt{2}}, \quad \text{and} \quad C_S = \frac{(2 - \delta)\bar{\rho}(1 + \sqrt{2})}{\delta\bar{\rho} - \sqrt{2}}.$$

If the smallest input angle $\alpha_{\min} > 60^\circ$, relations (8) and (10) can be simultaneously satisfied by choosing

$$C_T = \frac{2 - \delta}{\delta} + \frac{C_S}{\delta\bar{\rho}}, \quad \text{and} \quad C_S = \frac{2 - \delta}{\delta(1 - 2 \cos \alpha_{\min})}.$$

□

The analysis in Lemma 4 assumes that all angles in the input PSLG are greater than 60° . Although such geometries are rare, in practice a modification of the algorithm with a concentric circular shell splitting [14, 17] allows to guarantee the termination of the algorithm even though the small angles adjacent to the segments of the input PSLG cannot be improved.

Theorem 1 (Theorem 8 in [17], Theorem 1 in [14]). *For any vertex p of the output mesh, the distance to its nearest neighbor is at least $\frac{\text{lfs}(p)}{C_S+1}$.*

The proof in [17] relies only on Lemmata 1 and 4 here and, therefore, holds for the arbitrary points chosen within selection disks of skinny triangles.

Theorem 1 is used in the proof of the following theorem.

Theorem 2 (Theorem 10 in [17], Theorem 14 in [12], Theorem 3 in [14]). *Let $\text{lfs}_{\mathcal{M}}(p)$ be the local feature size at p with respect to a mesh \mathcal{M} (treating \mathcal{M} as a PSLG), whereas $\text{lfs}(p)$ remains the local feature size at p with respect to the input PSLG. Suppose a mesh \mathcal{M} with smallest angle θ has the property that there is some constant $k_1 \geq 1$, such that for every point p , $k_1 \text{lfs}_{\mathcal{M}}(p) \geq \text{lfs}(p)$. Then the cardinality of \mathcal{M} is less than k_2 times the cardinality of any other mesh of the input PSLG with smallest angle θ , where $k_2 = \mathcal{O}(k_1^2/\theta)$.*

Smaller values of δ offer more flexibility to a δ -graded Delaunay refinement algorithm in choosing Steiner points. However, from Lemma 4 it follows that as $\delta\bar{\rho}$ approach $\sqrt{2}$, C_T and C_S approach infinity, which leads to the worsening of the good grading guarantees. Therefore, along with satisfying application-specific requirements, the point insertion schemes should try to place Steiner points as close to circumcenters as possible.

4 Three-Dimensional Generalized Delaunay Refinement

In this section we introduce again two types of selection disks which can be used for the insertion of Steiner points. First, we prove the termination of a Delaunay refinement algorithm with the Type-1 selection disks. Then we give an example of an optimization based strategy for the insertion of Steiner points from the Type-1 selection disks. Finally, we introduce the Type-2 selection disk (which is always inside the Type-1 selection disk of the same skinny tetrahedron) and prove the good grading. As in [16], the proofs require that all incident segments and faces in the input geometry are separated by angles of at least 90° . According to our own experience and that of other authors [16, 18, 19], the Delaunay refinement algorithm works for much smaller angles in practice, although skinny tetrahedra adjacent to the small input angles cannot be removed. To prevent the algorithm from creating edges of ever diminishing length (for details, see the proofs below and the dataflow diagram in Figure 4), one can compute the insertion radii of candidate Steiner points explicitly and forbid the insertion of points which lead to the introduction of very short edges [16].

4.1 Proof of Termination with Selection Disks of Type I

Definition 7 (Selection disk of Type-1 in 3D). *If τ is a poor quality tetrahedron with circumcenter c , shortest edge length l , circumradius r , and*

circumradius-to-shortest edge ratio $\rho = r/l > \bar{\rho} \geq 2$, then the selection disk of Type-1 for the insertion of a Steiner point that would eliminate τ is the open disk with center c and radius $r - 2l$.

Following [16], the analysis below requires that the Delaunay refinement algorithm prioritize the splitting of encroached faces. In particular, when a Steiner point p encroaches upon several constrained faces, the encroached face which contains the projection of p is split. The projection of a point p onto a plane is the point in the plane which is closest to p . This requirement allows to achieve better bounds on the circumradius-to-shortest edge ratios in the final mesh.

Lemma 5 (Projection Lemma [16]). *Let f be a subfacet of the Delaunay triangulated facet F . Suppose that f is encroached upon by some vertex p , but p does not encroach upon any subsegment of F . Then $\text{proj}_F(p)$ lies in the facet F , and p encroaches upon a subfacet of F that contains $\text{proj}_F(p)$.*

Now we can prove the following lemma which establishes the relationship between the insertion radius of a point and its parent.

Lemma 6. *If p_i is a vertex of the mesh produced by a Delaunay refinement algorithm which chooses points within Type-1 selection disks of tetrahedra with circumradius-to-shortest edge ratios greater than $\bar{\rho} \geq 2$, then the following inequality holds:*

$$R(p_i) \geq C \cdot R(\hat{p}_i), \quad (11)$$

where we distinguish among the following cases:

- (i) $C = 2$ if p_i is a Steiner point chosen within the Type-1 selection disk of a skinny tetrahedron;
- (ii) $C = \frac{1}{\sqrt{2}}$ if p_i is a circumcenter of an encroached constrained face;
- (iii) $C = \frac{1}{\sqrt{2}}$ if p_i is a midpoint of an encroached constrained segment.

If p_i is an input vertex or if p_i and \hat{p}_i are on non-incident features, then

$$R(p_i) \geq \text{lfs}(p_i).$$

Proof. We need to prove only cases (i) and (ii) since the proofs of all other cases are independent of the choice of the point within the selection disk and are given in [16].

Case (i) Without the loss of generality, let $e(p_l p_m)$ be the shortest edge of the skinny tetrahedron $\tau(p_k p_l p_m p_n)$ and $\hat{p}_i = p_l$. Then

$$\begin{aligned} R(p_i) &> 2\|p_l - p_m\| && \text{(from Delaunay property and Definition 7)} \\ &\geq 2R(p_l) && \text{(from Definition 2 of insertion radius)} \\ &= 2R(\hat{p}_i); \end{aligned}$$

choose $C = 2$.




Fig. 3. (Left) An illustration of the relationship between the insertion radii of Steiner points, in the case of encroachment in three dimensions. (Right) A tetrahedron with high circumradius-to-shortest edge ratio.

Case (ii) Consider Figure 3(left). Let \hat{p}_i be inside the selection disk (the smaller shaded circle) of some skinny tetrahedron (not shown), such that \hat{p}_i encroaches upon the constrained face $\triangle(p_k p_l p_m)$, and p_i is the circumcenter of $\triangle(p_k p_l p_m)$. According to the projection requirement, let $\triangle(p_k p_l p_m)$ include $\text{proj}_F(\hat{p}_i)$, where F is the facet of the input PLC containing $\triangle(p_k p_l p_m)$. Without the loss of generality, suppose p_m is the point closest to $\text{proj}_F(\hat{p}_i)$ among the vertices of $\triangle(p_k p_l p_m)$. Because $\text{proj}_F(\hat{p}_i)$ lies inside the triangle $\triangle(p_k p_l p_m)$, it cannot be farther away from p_m than the circumcenter of $\triangle(p_k p_l p_m)$, i.e.,

$$\|\text{proj}_F(\hat{p}_i) - p_m\| \leq r(\triangle(p_k p_l p_m)). \quad (12)$$

Furthermore, because \hat{p}_i is inside the equatorial disk of $\triangle(p_k p_l p_m)$,

$$\|\text{proj}_F(\hat{p}_i) - \hat{p}_i\| < r(\triangle(p_k p_l p_m)). \quad (13)$$

From (12) and (13), as well as the fact that the triangle with vertices \hat{p}_i , $\text{proj}_F(\hat{p}_i)$, and p_m has a right angle at $\text{proj}_F(\hat{p}_i)$, we have:

$$\|\hat{p}_i - p_m\| < \sqrt{2}r(\triangle(p_k p_l p_m)). \quad (14)$$

Since p_m has to be in the mesh by the time \hat{p}_i is considered for insertion, $R(\hat{p}_i) \leq \|\hat{p}_i - p_m\|$. Then, using (14), we obtain:

$$R(\hat{p}_i) < \sqrt{2}r(\triangle(p_k p_l p_m)) = \sqrt{2}R(p_i);$$

choose $C = \frac{1}{\sqrt{2}}$. \square

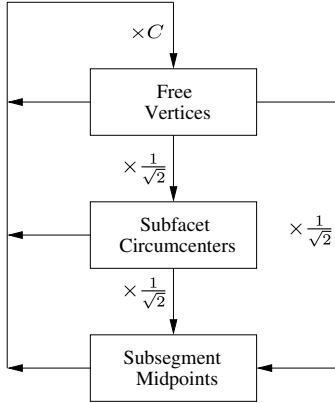


Fig. 4. Flow diagram from [16] illustrating the relationship between the insertion radius of a vertex and that of its parent in three dimensions. If no cycle has a product smaller than one, the algorithm will terminate. Input vertices are not shown since they do not participate in cycles. In [16] the constant $C = \bar{\rho} \geq 2$. In our case, with the use of Type-1 selection disks $C = 2$, and with the use of Type-2 selection disks $C = \delta\bar{\rho} \geq 2$.

Figure 4 shows the relationship between the insertion radii of mesh vertices and the insertion radii of their parents. We can see that if Inequality (11) is satisfied then no new edge will be created whose length is smaller than half of the length of some existing edge and the algorithm will eventually terminate because it will run out of places to insert new vertices.

4.2 An Example of a Point Selection Strategy

In two dimensions, by selecting the new Steiner point, we can construct only one new triangle which will be incident upon the shortest edge of the existing skinny triangle and which will have the required circumradius-to-shortest edge ratio. The three-dimensional case, however, is complicated by the fact that several new tetrahedra may be incident upon the shortest edge of the existing skinny tetrahedron. The example in Figure 3(right) shows a skinny tetrahedron $\tau(p_k p_l p_m p_n)$ with two new tetrahedra $\tau(p_i p_k p_l p_n)$ and $\tau(p_i p_k p_l p_m)$ that are incident upon the shortest edge $e(p_k p_l)$. By having to deal with multiple incident tetrahedra we face a multi-constrained optimization problem which to the best of our knowledge has not been analyzed with respect to the existence of the optimal solution and the construction of this solution.

A heuristic approach was suggested by Üngör [21] who proposed two types of off-center points in three dimensions. Based on the experimental data, he observes the following: “Use of both types of off-centers or the use of TYPE II off-centers alone outperforms the use of TYPE I off-centers alone, which in turn outperforms the use of circumcenters.” If a is the midpoint of the shortest edge $e(p_k p_l)$ of the tetrahedron and c is its circumcenter, than the TYPE II




Fig. 5. (Left) A wireframe model of a flying bat. (Right) The bat inside a box.

Table 1. The number of tetrahedra produced with the use of different point selection methods for three models in three dimensions. The first method (CC) always inserts circumcenters of skinny tetrahedra, the second method (OC) always inserts off-centers, and the third method (OPT) is optimization based. The minimal values in each cell are highlighted.

Model	Point position	$\bar{\rho}$								
		2.0	4.0	6.0	8.0	10.0	12.0	14.0	16.0	18.0
Points in cube	CC	1126	560	527	482	457	429	432	417	417
	OC	777	298	190	184	236	157	180	181	169
	OPT	1723	1619	174	174	174	174	174	174	174
Rocket	CC	1017	629	566	567	562	542	542	542	542
	OC	1060	729	540	673	679	660	660	660	660
	OPT	937	610	628	678	598	542	542	542	542
Bat	CC	24552	16561	15427	15226	15083	14923	14921	14923	14894
	OC	24985	21019	20781	20820	19764	19247	21058	17816	18301
	OPT	24628	17599	15970	15533	15267	15053	15074	15084	14939

off-center b is computed in [21] on the segment $\mathcal{L}(ac)$ such that $|\mathcal{L}(ab)| = \alpha_3 \left(\bar{\rho} + \sqrt{\bar{\rho}^2 - 1/4} \right)$, where α_3 is the perturbation factor. From experimental evidence in [21] the author suggests that a good choice for α_3 is 0.6. The insertion of TYPE II off-centers was implemented by Hang Si in **Tetgen** version 1.4.0 along with the circumcenter insertion method [18, 19]. We also added the implementation of an optimization-based point insertion method which with every insertion of a Steiner point within a Type-1 selection disk of a skinny tetrahedron minimizes the signed difference between the number of deleted good quality tetrahedra and the number of deleted poor quality tetrahedra. As opposed to the two-dimensional case, we did not restrict the position of the

Steiner point to a specific arc, but instead sampled 1,000 points uniformly in spherical coordinates of the Type-1 selection disk and chose the best one. For our experiments, we used the following three geometries. The “Points in cube” model is the unit cube with two additional points close to its center at 10^{-4} distance from each other. The “Rocket” model is the example PLC supplied with the `Tetgen` distribution. The “Bat” model [13] is shown in Figure 5. Table 1 summarizes the results of our experiments. From our point of view, these results do not provide a basis for conclusions, and further research is required.

4.3 Proof of Good Grading with Selection Disks of Type II

Definition 8 (Selection disk of Type-2 in 3D). *If τ is a poor quality tetrahedron with circumcenter c , shortest edge length l , circumradius r , and circumradius-to-shortest edge ratio $\rho = r/l > \bar{\rho} \geq 2$, then the Type-2 selection disk for the insertion of a Steiner point that would eliminate τ is the open disk with center c and radius $r(1 - \frac{2}{\rho})$.*

Remark 7. Note from definitions 7 and 8 that the radius of the Type-2 selection disk is always smaller than the radius of the Type-1 selection disk of the same skinny tetrahedron because $r(1 - \frac{2}{\rho}) = r - \frac{2l}{\rho}$ and $\rho > \bar{\rho}$.

Definition 9 (δ -graded Delaunay refinement algorithm in 3D). *If for every tetrahedron τ with circumcenter c , circumradius r , shortest edge length l , and circumradius-to-shortest edge length ratio $\rho = r/l > \bar{\rho} \geq 2$ a Delaunay refinement algorithm selects a Steiner point p_i within the Type-2 selection disk such that $\|p_i - c\| < r(1 - \delta)$, where $2/\bar{\rho} \leq \delta \leq 1$, we say that this Delaunay refinement algorithm is δ -graded.*

Lemma 7. *If p_i is a vertex of the mesh produced by a δ -graded Delaunay refinement algorithm then the following inequality holds:*

$$R(p_i) \geq C \cdot R(\hat{p}_i), \quad (15)$$

where we distinguish among the following cases:

- (i) $C = \delta \bar{\rho}$ if p_i is a Steiner point chosen within the Type-2 selection disk of a skinny tetrahedron;
- (ii) $C = \frac{1}{\sqrt{2}}$ if p_i is a circumcenter of an encroached constrained face;
- (iii) $C = \frac{1}{\sqrt{2}}$ if p_i is a midpoint of an encroached constrained segment.

If p_i is an input vertex or if p_i and \hat{p}_i are on non-incident features, then

$$R(p_i) \geq \text{lfs}(p_i).$$

Proof. We need to present a new proof only for case (i) since the proof for case (ii) is the same as in Lemma 6, and the proofs for all other cases are

independent of the choice of the point within the selection disk and are given in [16].

Case (i) Without the loss of generality, let $e(p_l p_m)$ be the shortest edge of the skinny tetrahedron $\tau(p_k p_l p_m p_n)$ and $\hat{p}_i = p_l$. Then

$$\begin{aligned} R(p_i) &> \delta r(\tau(p_k p_l p_m p_n)) && \text{(from Delaunay property and Def. 9)} \\ &= \delta \rho(\tau(p_k p_l p_m p_n)) \|p_l - p_m\| \\ &> \delta \bar{\rho} \|p_l - p_m\| && \text{(since the tetrahedron is skinny)} \\ &\geq \delta \bar{\rho} R(\hat{p}_i); \end{aligned}$$

choose $C = \delta \bar{\rho}$. \square

Lemma 8 (Extension of Lemma 3 to 3D). *If p is a vertex of the mesh produced by a δ -graded Delaunay refinement algorithm and C is the constant specified by Lemma 7, then the following inequality holds:*

$$D(p) \leq \frac{2 - \delta}{\delta} + \frac{D(\hat{p})}{C}. \quad (16)$$

Proof. The proof is formally equivalent to the proof of Lemma 3. \square

Lemma 9 (Extension of Theorem 5 in [16]). *Suppose that $\bar{\rho} > 2$ and the input PLC satisfies the projection condition. Then there exist fixed constants $C_T \geq 1$, $C_F \geq 1$, and $C_S \geq 1$ such that, for any vertex p inserted or rejected by a δ -graded Delaunay refinement algorithm, the following relations hold:*

- (i) $D(p) \leq C_T$ if p is chosen from the selection disk of a skinny tetrahedron;
- (ii) $D(p) \leq C_F$ if p is the circumcenter of an encroached constrained face;
- (iii) $D(p) \leq C_S$ if p is the midpoint of an encroached constrained segment.

Therefore, the insertion radius of every vertex has a lower bound proportional to its local feature size.

Proof. The proof is by induction and is similar to the proof of Lemma 7 in [16]. The base case covers the input vertices, and the inductive step covers the other three types of vertices: free vertices, the circumcenters of constrained faces, and the midpoints of constrained subsegments. The constants are chosen as follows:

$$C_T = C_0 \frac{C_1 + 1 + \sqrt{2}}{C_1 - 2}, \quad C_F = C_0 \frac{(1 + \sqrt{2})C_1 + \sqrt{2}}{C_1 - 2}, \quad C_S = C_0 \frac{(3 + \sqrt{2})C_1}{C_1 - 2},$$

where $C_0 = (2 - \delta)/\delta$ and $C_1 = \delta \bar{\rho}$. As δ approaches $2/\bar{\rho}$, the constants approach infinity. \square

Theorem 3 (Theorem 6 in [16]). *For any vertex p of the output mesh, the distance to its nearest neighbor is at least $\frac{\text{fs}(p)}{C_S + 1}$.*

The proof of this theorem does not depend on the specific position of the Steiner point inside the selection disk. Hence, the theorem holds in the context of three-dimensional generalized Delaunay refinement.

5 Conclusions

In this paper we presented the first to our knowledge decoupled definitions of two types of selection disks along with the corresponding proofs of termination and size optimality, both in two and in three dimensions. The primary goal which we pursued while developing the selection disk ideas was the design of parallel Delaunay refinement algorithms. A number of techniques have been developed by multiple authors for the selection of Steiner points for sequential Delaunay refinement. The development of the corresponding parallel algorithm for each of the sequential point selection techniques would have been cumbersome and time consuming. Instead, we have shown that there exists an entire region for the sequential selection of Steiner points. Therefore, if a parallel algorithm framework is designed with the assumption that the points are chosen from this region, all sequential point placement techniques can be used as a black box.

The additional flexibility in choosing the Steiner points can be used to satisfy a number of diverse mesh improvement goals. For example, one can construct the difference of the selection disk with the forbidden regions which create slivers to find the region which does not create slivers. Then, instead of using randomized point selection for sliver elimination like in [4, 9, 10], a deterministic insertion of a single point is sufficient. As another example, if a selection disk intersects a boundary between different materials, one would like to insert the point along the boundary with the goal of constructing a boundary conformal mesh.

Delaunay refinement algorithms can also use selection disks for the splitting of constrained segments and faces; we plan to address it in the future work along with the complete generalization of point insertion in all dimensions.

Acknowledgments

We thank Gary Miller for helpful references and George Karniadakis for the bat model. This work was supported (in part) by the following NSF grants: ACI-0312980, CNS-0521381, as well as a John Simon Guggenheim Award. We also thank the anonymous reviewers for insightful comments and constructive suggestions.

References

1. A. Bowyer. Computing Dirichlet tessellations. *Computer Journal*, 24:162–166, 1981.
2. A. N. Chernikov and N. P. Chrisochoides. Generalized Delaunay mesh refinement: From scalar to parallel. In *Proceedings of the 15th International Meshing Roundtable*, pages 563–580, Birmingham, AL, Sept. 2006. Springer.
3. L. P. Chew. Guaranteed quality mesh generation for curved surfaces. In *Proceedings of the 9th ACM Symposium on Computational Geometry*, pages 274–280, San Diego, CA, 1993.

4. L. P. Chew. Guaranteed-quality Delaunay meshing in 3D. In *Proceedings of the 13th ACM Symposium on Computational Geometry*, pages 391–393, Nice, France, 1997.
5. B. N. Delaunay. Sur la sphere vide. *Izvestia Akademia Nauk SSSR, VII Seria, Otdelenie Matematika i Estestvennyka Nauk*, 7:793–800, 1934.
6. W. H. Frey. Selective refinement: A new strategy for automatic node placement in graded triangular meshes. *International Journal for Numerical Methods in Engineering*, 24(11):2183–2200, 1987.
7. P.-L. George and H. Borouchaki. *Delaunay Triangulation and Meshing. Application to Finite Elements*. HERMES, 1998.
8. F. Labelle. Sliver removal by lattice refinement. In *Proceedings of the 21nd Symposium on Computational Geometry*, pages 347–356, New York, NY, USA, 2006. ACM Press.
9. X.-Y. Li. Generating well-shaped d-dimensional Delaunay meshes. *Theoretical Computer Science*, 296(1):145–165, 2003.
10. X.-Y. Li and S.-H. Teng. Generating well-shaped Delaunay meshes in 3D. In *Proceedings of the 12th annual ACM-SIAM symposium on Discrete algorithms*, pages 28–37, Washington, D.C., 2001.
11. G. L. Miller, D. Talmor, S.-H. Teng, and N. Walkington. A Delaunay based numerical method for three dimensions: Generation, formulation, and partition. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing*, pages 683–692, Las Vegas, NV, May 1995.
12. S. A. Mitchell. Cardinality bounds for triangulations with bounded minimum angle. In *Proceedings of the 6th Canadian Conference on Computational Geometry*, pages 326–331, Saskatoon, Saskatchewan, Canada, Aug. 1994.
13. I. Pivkin, E. Hueso, R. Weinstein, D. Laidlaw, S. Swartz, and G. Karniadakis. Simulation and visualization of air flow around bat wings during flight. In *Proceedings of the International Conference on Computational Science*, pages 689–694, Atlanta, GA, 2005.
14. J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of Algorithms*, 18(3):548–585, 1995.
15. J. R. Shewchuk. *Delaunay Refinement Mesh Generation*. PhD thesis, Carnegie Mellon University, 1997.
16. J. R. Shewchuk. Tetrahedral mesh generation by Delaunay refinement. In *Proceedings of the 14th ACM Symposium on Computational Geometry*, pages 86–95, Minneapolis, MN, 1998.
17. J. R. Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry: Theory and Applications*, 22(1–3):21–74, May 2002.
18. H. Si. On refinement of constrained Delaunay tetrahedralizations. In *Proceedings of the 15th International Meshing Roundtable*, pages 509–528, Birmingham, AL, Sept. 2006. Springer.
19. H. Si and K. Gaertner. Meshing piecewise linear complexes by constrained Delaunay tetrahedralizations. In *Proceedings of the 14th International Meshing Roundtable*, pages 147–163, San Diego, CA, Sept. 2005. Springer.
20. A. Üngör. Off-centers: A new type of Steiner points for computing size-optimal guaranteed-quality Delaunay triangulations. In *Proceedings of LATIN*, pages 152–161, Buenos Aires, Argentina, Apr. 2004.
21. A. Üngör. Quality triangulations made smaller. In *Proceedings of the European Workshop on Computational Geometry*, Technische Universiteit Eindhoven, Mar. 2005. Accessed online at <http://www.win.tue.nl/EWCG2005/Proceedings/2.pdf> on April 28, 2007.
22. D. F. Watson. Computing the n-dimensional Delaunay tessellation with application to Voronoi polytopes. *Computer Journal*, 24:167–172, 1981.