ROBUST THREE DIMENSIONAL DELAUNAY REFINEMENT

Steven E. Pav[1] Noel J. Walkington[2]

[1] *University of California at San Diego, La Jolla, CA. spav@ucsd.edu*
[2] *Carnegie Mellon University, Pittsburgh, PA. noelw@andrew.cmu.edu*

ABSTRACT

The Delaunay Refinement Algorithm for quality meshing is extended to three dimensions. The algorithm accepts input with arbitrarily small angles, and outputs a Conforming Delaunay Tetrahedralization where most tetrahedra have radius-to-shortest-edge ratio smaller than some user chosen $\mu > 2$. Those tets with poor quality are in well defined locations: their circumcenters are describably near input segments. Moreover, the output mesh is well graded to the input: short mesh edges only appear around close features of the input. The algorithm has the added advantage of not requiring *a priori* knowledge of the "local feature size," and only requires searching locally in the mesh.

Keywords: mesh generation, tetrahedral, unstructured, Delaunay, Ruppert's Algorithm.

1. INTRODUCTION

The method of Delaunay Refinement, advanced by Ruppert and Chew [1, 2], has been successfully applied to problems of mesh generation in two dimensions. The essence of the method is to maintain the Delaunay Triangulation of a set of points, use this data structure to efficiently locate problems in the mesh, then to resolve problematic triangles by adding their circumcenters to the mesh as Steiner Points. The method has been expanded to allow for curved domain boundaries [3], and its input and output conditions have been improved with some variations to the method [4, 5, 6].

Much work has gone into the extending the method to three (or more) dimensions. The early formulations suffered from the same problem as in Ruppert's original exposition, namely that planar facets meet at obtuse angles [7, 8]. This Draconian restriction was relaxed by Murphy *et al.* and Cohen-Steiner *et al.* [9, 10]. However, with both of these algorithms, only a Conforming Delaunay Tetrahedralization is produced, and no attempt is made to remove poor quality tetrahedra. In both of these methods, too, the output is not necessarily "well-graded," meaning that mesh edges can be produced which are much smaller than the "local feature size," which, roughly speaking, is a measure of the size of the input near a given point in space.

These deficiencies were addressed by Cheng and Poon [11]. Their algorithm accepts input with arbitrarily small facet angles, and outputs a Conforming Delaunay mesh which is well-graded to the input, with weak guarantees on the quality of tetrahedra of the mesh.

The algorithms of Cohen-Steiner *et al.* and Cheng and Poon can be similarly characterized: they both start by using the local feature size of the input as a black box function to generate a protecting region or buffer zone; the structure of this region insures conformality of the mesh to the input; outside this region, the obtuse angle condition applies, so regular Delaunay Refinement suffices. This simplification does neither algorithm justice, but illustrates the distinctions between them and traditional 2D Delaunay Refinement. In particular note that 2D methods do not require calculation of the local feature size, rather, to the contrary, one can view 2D Delaunay Refinement

Supported in part by National Science Foundation Grants DMS–9973285 and CCR–9902091. This work was also supported by the NSF through the Center for Nonlinear Analysis.

as a means of *calculating* the local feature size [5].

This leads to another similarity between these algorithms: no mention is made of how the local feature size is to be calculated. The reader is left to imagine a $\mathcal{O}\left(n^2\right)$ brute force search was intended. Omission of these details lead to the present work, along with a desire to both improve the quality guarantees of Cheng and Poon's algorithm, and to extend the "adaptive" approach to poor quality triangles to 3D [8, 5].

Our solution builds up knowledge of the local feature size as it constructs the mesh. It starts with an estimate of local feature size based only on location of input points, then recovers the segments of the input, updating its estimate. A protecting region is then constructed, after which the algorithm attempts to recover input facets. New knowledge about facets causes the estimate of local feature size to be updated, which can cause shrinkage of the protecting region. When facets are recovered, poor quality tets are attacked.

This strategy of adaptively sizing the protecting region as the local feature size is found may actually be inferior to simply accepting a brute force $\mathcal{O}\left(n^2\right)$ search up front; we claim that there is an obvious (and simpler) variant of our algorithm which assumes the local feature size is known *a priori*. Having not yet implemented either variant, we cannot compare them.

Recent work by Cheng *et al.* addresses many of the deficiencies of earlier works. Their algorithm requires computation of the local feature size only at input vertices[1], has good output quality guarantees (similar to those of the present work), is simple to describe, and does not require explicit construction of a protecting region where input facets make small angles [12]. However, to implement the latter property, the algorithm does not enforce the "empty circumball property" of facets. For this reason, their algorithm will only accept as input a 2-manifold and cannot accept a general piecewise linear complex[2] (see Definition 2.1). This is a deficiency which is fundamental to their algorithm, one which cannot be eliminated without making major nontrivial changes to their method.

We must note that all the 3D Delaunay based methods bound only the circumradius-to-shortest-edge ratio of the output. These algorithms all leave behind "slivers," which have modest circumradius-edge ratios, but large circumradius-inradius ratios [7, 13]. A number of methods have been proffered to deal with slivers [14, 15, 16, 17].

[1]Here again, a brute force search is implied.

[2]Indeed, it will most certainly fail for some input sets with many facets sharing a single edge.

2. PROBLEM DEFINITION

We assume the input to the algorithm takes the form of a piecewise linear complex, which we define as follows:

Definition 2.1 (PLC). A set of points, a set of segments, and a set of faces, $(\mathcal{P}, \mathcal{S}, \mathcal{F})$ form a *Piecewise Linear Complex* (PLC) if the union $\mathcal{W} = \{\emptyset\} \cup \mathcal{P} \cup \mathcal{S} \cup \mathcal{F}$ has the following properties:

(i) for any simplex $S \in \mathcal{W}$, the boundary of S is the union of elements of $\mathcal{W}$. (Let the boundary of a point be the empty set, and the boundary of the empty set be the empty set).
(ii) given two simplices of $\mathcal{W}$, S_1, S_2, the intersection of the two simplices is the union of simplices of $\mathcal{W}$, and is a (possibly empty) subset of the boundaries of both S_1 and S_2.

Let $\theta^*_{1,1} \leq \pi/3$ be a lower bound on the minimum angle subtended by segments of a given PLC. Let $\theta^*_{2,2} \leq \pi/3$ be a lower bound on the minimum angle subtended by two faces of $\mathcal{F}$, and let $\theta^*_{1,2} \leq \pi/3$ be a lower bound on the minimum angle subtended by a segment of $\mathcal{S}$ and a face of $\mathcal{F}$; let $\theta^* = \theta^*_{1,1} \wedge \theta^*_{1,2} \wedge \theta^*_{2,2}$.

Care must be taken defining $\theta^*_{2,2}$. If f_1, f_2 are faces which share only an input point, z, then the angle between them is the minimum angle subtended at z by a segment in f_1 and another in f_2. When f_1, f_2 are two faces sharing an input segment s, then the angle between them is the angle subtended by the intersection of f_1, f_2 with the plane P perpendicular to s.

We assume that it is algorithmically 'easy' to detect if distinct features of $(\mathcal{P}, \mathcal{S}, \mathcal{F})$ are nondisjoint, and to calculate the intersection of a face of $\mathcal{F}$ with a sphere.

A mesher should "refine" the input PLC:

Definition 2.2 (PLC Refinement). A PLC $(\mathcal{P}', \mathcal{S}', \mathcal{F}')$ *refines* a PLC $(\mathcal{P}, \mathcal{S}, \mathcal{F})$ if

(i) $\mathcal{P} \subseteq \mathcal{P}'$,
(ii) for every $s \in \mathcal{S}$, s is the union of segments in $\mathcal{S}'$,
(iii) for every $f \in \mathcal{F}$, f is the union of faces in $\mathcal{F}'$.

Local feature size will be used to define the meshing problem, and to show termination of our algorithm. Our definition is the classical one:

Definition 2.3 (Local Feature Size). When a given PLC, $(\mathcal{P}, \mathcal{S}, \mathcal{F})$ is understood, define the *local feature size,* $\text{lfs}_i(x)$, as the distance from x to two features of the PLC which are of dimension no greater than i. Thus, for example, $\text{lfs}_0(x)$ is the distance from x to the second nearest point of $\mathcal{P}$. Let $\text{lfs}(x) = \text{lfs}_2(x)$.

Note 2.4. Note the following facts about this function: (i) For any x, $\text{lfs}_2(x) \leq \text{lfs}_1(x) \leq \text{lfs}_0(x)$, (ii) $\text{lfs}_i(x)$ is a Lipschitz function with constant 1, (iii) $\text{lfs}_i(x)$ has a positive minimum value on $\mathbb{R}^3$.

The Meshing Problem is then as follows: given a PLC, $\mathcal{W}$, construct a set of points $\mathcal{P}'$ such that: (i) the Delaunay Tetrahedralization of $\mathcal{P}'$, considered as a PLC, refines $\mathcal{W}$, (ii) for any $p \in \mathcal{P}'$, $\text{lfs}(p)$ bounds the nearest neighbor distance of p in $\mathcal{P}'$, (iii) any tet in the Delaunay Tetrahedralization of $\mathcal{P}'$ has "good quality."

3. ALGORITHM OVERVIEW

First the input is fed to a 'groomer' (see Theorem 4.2), which adds Steiner Points to the segments of $\mathcal{S}$. This groomer does not require knowledge of the local feature size, and all operations are based on 'local' information.

Arcs are then constructed in the input facets, based on subsegments output from the groomer. The union of these arcs will be a number of closed curves. The closed curves divide each facet into a "free" area, and the "collar" area, with the latter being also bounded by the boundary of the facet. The properties of the groomer are tailored so that adjacent arcs meet at obtuse angles and the local feature size after the collar construction is no smaller than a constant times the local feature size of the input.

The algorithm finishes with the refinement stage. A Conforming Delaunay Tetrahedralization is constructed by refining the facets outside the collar regions; the collar regions are designed so that they do not require refinement. This is essentially the strategy of Murphy *et al.* and Cohen-Steiner *et al.* [9, 10]. Note, however, that because we do not have a local feature size oracle, the algorithm may have to further refine the 1-skeleton during the refinement stage, thus redefining the arcs. When this occurs, we will remove arcs from $\mathcal{A}$ and replace them with others.

In the refinement stage, poor quality tets are removed by adding their circumcenters. We employ the usual strategy of rejecting circumcenters if they encroach lower dimensional simplices, and splitting such a simplex instead. The arcs form a protective layer around the collar, and may be split by the algorithm if encroached. This is the "SOS" strategy of Cohen-Steiner *et al.* [10]. Note, however, that since adjacent arcs meet at obtuse angles, we need not split arcs on concentric shells as Cohen-Steiner *et al.* do.

Some poor quality tets may be left behind in the mesh if their circumcenters would disrupt the collar region. This is essentially the adaptive approach to poor quality triangles generalized to 3D [8, 5]. Note this is also the technique employed in the recent work by Cheng *et al.* [12].

4. GROOMING

4.1 A Grooming Theorem

We first claim the existence of a grooming procedure which will refine the input. The proof of the theorem is constructive: we describe the algorithm. The following definition simplifies notation in the theorem.

Definition 4.1 (end segment). Let $(\mathcal{P}', \mathcal{S}', \mathcal{F}')$ be a refinement of the PLC $(\mathcal{P}, \mathcal{S}, \mathcal{F})$. We say that a segment $s \in \mathcal{S}'$ is an *end segment* if one of its endpoints is a point of $\mathcal{P}$. Otherwise s is a *non-end segment* .

Theorem 4.2 (Well Graded Groomer). *There is a groomer that takes PLC,* $\mathcal{W} = (\mathcal{P}, \mathcal{S}, \mathcal{F})$, *parameters* $(\beta, \gamma_p, \gamma_e, \gamma_n)$, *with* $2 < \beta$, $0 < \gamma_p, \gamma_e, \gamma_n$, *and* $5.53 \leq \gamma_p^2 < \gamma_e^2 + 1$, *and outputs a PLC,* $\mathcal{W}' = (\mathcal{P}', \mathcal{S}', \mathcal{F})$ *which refines the input and such that the following hold:*

(i) **Point Location:** *If* $p \in \mathcal{P}' \setminus \mathcal{P}$, *then there is some segment* $s \in \mathcal{S}$ *containing* p.

(ii) **Grading:** *There are constants* C_0, C_1, *with* $\gamma_p \leq C_0 \leq 2\gamma_p < C_1$, *depending only on* $\beta, \gamma_p, \gamma_e, \gamma_n, \theta^*_{1,1}$, *such that if* $s \in \mathcal{S}'$ *has endpoint* $p \in \mathcal{P}'$, *then*

$$\text{lfs}_1(p) \leq \begin{cases} C_0 |s| & \text{if } p \in \mathcal{P}, \\ C_1 |s| & \text{if } p \notin \mathcal{P}, \end{cases}$$

where lfs_1 *is defined in terms of the input PLC.*

(iii) β**-Balance:** *If* $s_1, s_2 \in \mathcal{S}'$ *are non-end segments which share an endpoint, then*

$$\frac{1}{\beta} \leq \frac{|s_1|}{|s_2|} \leq \beta.$$

If $s_1, s_2 \in \mathcal{S}'$ *share an endpoint and* s_1 *is an end segment, while* s_2 *is a non-end segment, then*

$$|s_2| \leq |s_1| .$$

(iv) γ**-Isolation:** *For any* $p \in \mathcal{P}' \setminus \mathcal{P}$, *let* l *be the length of the shorter segment of which* p *is an endpoint. Then if* $q \in \mathcal{P}'$ *is* not *collinear to the segment containing* p, *then*

$$|p - q| \geq \begin{cases} \gamma_n\, l & \text{if } p \text{ between non-end segments}, \\ \gamma_e\, l & \text{otherwise}. \end{cases}$$

(v) **Input Point Isolation:** *If* $p \in \mathcal{P}$ *is the endpoint of some segment of* $\mathcal{S}$, *then all segments of* $\mathcal{S}'$ *with endpoint* p *have the same length, call it* $\mathrm{d}(p)$. *Moreover*

- *(a) if* q *is a point not on a segment with* p *as endpoint, then* $|p - q| \geq \gamma_p\, \mathrm{d}(p)$.
- *(b)* $\gamma_p\, \mathrm{d}(p) \leq \text{lfs}_0(p)$, *where* lfs_0 *is defined in terms of the input.*

Moreover, we make the claim that a groomer described in Theorem 4.2 can be implemented using 3D Delaunay code. The local feature size need not be known, and all searching is local.

Proof. First we construct the Delaunay Tetrahedralization of the point set $\mathcal{P}$. There is an edge in the Delaunay Tetrahedralization from each point to its nearest neighbor. Thus from this data structure we can compute $\operatorname{lfs}_0(p)$ for every $p \in \mathcal{P}$. We will need the tetrahedralization anyway, so there's no need to rely on fancier methods to find the local feature size.

Our groomer maintains a set of points and a set of segments, $\mathcal{P}', \mathcal{S}'$, respectively, which are initialized as the input. Herein, when we talk of a segment of $\mathcal{S}'$, call it (p, q), being "split," at a point m, we mean that (p, q) is removed from $\mathcal{S}'$, and $(p, m), (m, q)$ are added to $\mathcal{S}'$, and m is added to $\mathcal{P}'$. When the point m is not specified, it is the midpoint of the segment. Since we will only split segments, all points of $\mathcal{P}' \setminus \mathcal{P}$ are clearly on segments.

For each point $p \in \mathcal{P}$, the groomer splits each segment of $\mathcal{S}$ with p as endpoint by a point m at a distance $\operatorname{lfs}_0(p)/\gamma_p$ from p. Letting $\mathrm{d}(p)$ be the length of each segment with p as endpoint, this guarantees that $\gamma_p \mathrm{d}(p) = \operatorname{lfs}_0(p)$. Throughout the algorithm we may decrease $\mathrm{d}(p)$, but never increase it; thus we have shown the second half of item (v).

We make the assumption that $C_0 \geq \gamma_p$, and see that $\operatorname{lfs}_1(p) \leq \operatorname{lfs}_0(p) \leq C_0 |s|$, where s is a segment of $\mathcal{S}'$ with p as endpoint. We claim that this procedure is a λ-Feature Size Augmenter [5], for $\lambda = 1 + \gamma_p$. In particular we claim that if $p \in \mathcal{P}'$, and $(p, q) \in \mathcal{S}'$, then $\operatorname{lfs}_1(p) \leq (1 + \gamma_p) |p - q|$. We show this here:

- If $p \in \mathcal{P}$, then, as above, $\operatorname{lfs}_1(p) \leq \operatorname{lfs}_0(p) = \gamma_p |p - q|$, which suffices.
- Otherwise, if $q \in \mathcal{P}$, then $\operatorname{lfs}_1(p) \leq |p - q| + \operatorname{lfs}_1(q) \leq |p - q| + \operatorname{lfs}_0(q) = (1 + \gamma_p) |p - q|$, which suffices.
- Otherwise let (s, t) be the input segment containing (p, q). Then $\operatorname{lfs}_0(p) \leq \frac{1}{2} |s - t|$. Moreover, as $\operatorname{lfs}_0(s) \leq |s - t| \geq \operatorname{lfs}_0(t)$, $|p - q| \geq (1 - 2/\gamma_p) |s - t|$, and so $\operatorname{lfs}_1(p) \leq |p - q| / (2 - 4/\gamma_p)$. Assumptions on γ_p imply that $(1/\gamma_p) \leq \left(-3 + \sqrt{41}\right)/8$, and so $\operatorname{lfs}_1(p) \leq (1 + \gamma_p) |p - q|$.

We make the requirement that $C_1 \geq 1 + \gamma_p$. We inductively assume, throughout this proof, that item (ii) of the theorem is satisfied.

We now conceive of the groomer as a game: the groomer 'plays' any of the following 'moves' until it can play no more, then it outputs the augmented PLC. As we describe each move, we show that it preserves grading, *i.e.*, that item (ii) is satisfied under certain assumptions on C_0, C_1; this will prove termination of the process [1, 5].

1. Suppose p is an input point, and there is some point, q of $\mathcal{P}'$ such that q is not on a segment of $\mathcal{S}$ with p as endpoint, and $|p - q| < \gamma_p \mathrm{d}(p)$, where $\mathrm{d}(p)$ is the length of the segments with p as endpoint. Then split each segment of $\mathcal{S}'$ with p as endpoint at their midpoints. Note that this cuts $\mathrm{d}(p)$ in half.
   First note that $\operatorname{lfs}_1(p) \leq |p - q|$. Let (p, t) be a segment which is split in this way. Let m be its midpoint. Then
   $$\begin{aligned} \operatorname{lfs}_1(p) &\leq |p - q| < \gamma_p \mathrm{d}(p) = \gamma_p |p - t| \\ &\leq 2\gamma_p |p - m|. \end{aligned}$$
   So it suffices to take
   $$\boxed{2\gamma_p \leq C_0}$$
   Now note that
   $$\operatorname{lfs}_1(t) \leq |p - t| + \operatorname{lfs}_1(p) \leq 2|p - m| + 2\gamma_p |p - m|.$$
   Thus we require that
   $$\boxed{2 + 2C_0 \leq C_1}$$
   Under this assumption we could also show that $\operatorname{lfs}_1(m) \leq C_1 |p - m|$.
2. Suppose (p, s),(s, q) are non-end segments of $\mathcal{S}'$ such that $|p - s| > \beta |s - q|$. Then split (p, s) at its midpoint.
   Let m be the midpoint. Then
   $$\begin{aligned} \operatorname{lfs}_1(p) &\leq |p - s| + \operatorname{lfs}_1(q) \\ &\leq 2|p - m| + C_1 |s - q| \\ &< \left(2 + \frac{2C_1}{\beta}\right) |p - m|. \end{aligned}$$
   It suffices to take
   $$\boxed{\frac{2\beta}{\beta - 2} \leq C_1}$$
   Under this assumption we can also show $\operatorname{lfs}_1(m) \leq C_1 |m - p|$, and $\operatorname{lfs}_1(s) \leq C_1 |s - m|$.
3. Suppose p is an input point, (p, s) is an end segment, (s, q) is a non-end segment of $\mathcal{S}'$ and $|s - q| > |p - s|$. Then split (s, q) at its midpoint.
   Let m be the midpoint. Then
   $$\begin{aligned} \operatorname{lfs}_1(q) &\leq |q - p| + \operatorname{lfs}_1(p) \\ &\leq |q - s| + |s - p| + C_0 |s - p| \\ &< (2 + 2C_0) |p - m|. \end{aligned}$$
   We have already assumed that $2 + 2C_0 \leq C_1$. Laboring under this assumption we can also show $\operatorname{lfs}_1(m) \leq C_1 |m - p|$, and $\operatorname{lfs}_1(s) \leq C_1 |s - m|$.

4. Suppose (s,p), (p,t) are non-end segments of $\mathcal{S}'$ and q is a point of $\mathcal{P}'$ noncollinear to (s,p) such that

$$|p-q| < \gamma_n |p-s| \leq \gamma_n |p-t| .$$

Then split (p,t) at its midpoint.
Let m be the midpoint of (p,t). If q is (on) an input feature disjoint from the one containing (p,t), then $\operatorname{lfs}_1(p) \leq |p-q|$. If q is not on such a feature, then there is an input point x such that $\angle pxq \geq \theta^*_{1,1}$. Thus $|p-q| \geq |x-p| \sin\theta^*_{1,1}$. Now note that $C_0 |x-p| \geq \operatorname{lfs}_1(x)$. Then

$$\begin{aligned} \operatorname{lfs}_1(p) &\leq |x-p| + \operatorname{lfs}_1(x) \\ &\leq (1+C_0)|x-p| \leq \frac{1+C_0}{\sin\theta^*_{1,1}} |p-q| . \end{aligned}$$

Then

$$\begin{aligned} \operatorname{lfs}_1(t) &\leq |t-p| + \operatorname{lfs}_1(p) \\ &\leq 2|t-m| + \frac{1+C_0}{\sin\theta^*_{1,1}} |p-q| \\ &< \left[2 + 2\gamma_n \frac{1+C_0}{\sin\theta^*_{1,1}}\right] |t-m| . \end{aligned}$$

It suffices to take

$$\boxed{2 + 2\gamma_n \frac{1+C_0}{\sin\theta^*_{1,1}} \leq C_1}$$

Laboring under this assumption we can also show $\operatorname{lfs}_1(m) \leq C_1 |m-t|$, and $\operatorname{lfs}_1(p) \leq C_1 |p-m|$.

5. Suppose s is an input point, (s,p) an end segment, (p,t) a non-end segment of $\mathcal{S}'$ and q is a point of $\mathcal{P}'$ noncollinear to (p,s) such that

$$|p-q| < \gamma_e |p-t| .$$

Then split (p,t) at its midpoint.
Let m be the midpoint of (p,t). As above, it can be shown that

$$\operatorname{lfs}_1(p) \leq \frac{1+C_0}{\sin\theta^*_{1,1}} |p-q| .$$

Then

$$\begin{aligned} \operatorname{lfs}_1(t) &\leq |t-p| + \operatorname{lfs}_1(p) \\ &< \left[2 + 2\gamma_e \frac{1+C_0}{\sin\theta^*_{1,1}}\right] |t-m| . \end{aligned}$$

It suffices to take

$$\boxed{2 + 2\gamma_e \frac{1+C_0}{\sin\theta^*_{1,1}} \leq C_1}$$

Laboring under this assumption we can also show $\operatorname{lfs}_1(m) \leq C_1 |m-t|$, and $\operatorname{lfs}_1(p) \leq C_1 |p-m|$.

We see that we can take

$$C_0 = 2\gamma_p, \quad C_1 = \left[\frac{2\beta}{\beta-2}\right] \vee \left[2 + 2(\gamma_n \vee \gamma_e) \frac{1+C_0}{\sin\theta^*_{1,1}}\right] .$$

It should be clear that when the algorithm terminates, then item (iii) and item (iv) will be satisfied. Also the first part of item (v) will be satisfied; we already showed the second half. ☐

Regarding the claim that this algorithm can be performed with only local searching, we note that the ability to play any of the five moves of the game can be detected locally in a Delaunay Tetrahedralization of $\mathcal{P}'$, which we may assume the algorithm maintains at all times. Actually, detecting that the first move can be played may be a bit tricky. We discuss this briefly here.

Suppose that move 1 can be played, but not move 5. Let p be an input point for which there is some q with $|p-q| < \gamma_p \operatorname{d}(p)$. Let q be the closest such point. We know that q is not an input point. If the sphere with (p,q) is empty, then (p,q) is an edge in the Delaunay Tetrahedralization, so a local test suffices to find that q is 'too close.' If this sphere contains some points of $\mathcal{P}'$, by assumption that q is the closest such point, there must be some point t on a segment with p as endpoint that is inside the sphere. By convexity we can assume (p,t) is a segment of $\mathcal{S}'$. Since t is inside the sphere we know $\angle ptq > \pi/2$. Then by the cosine rule $|p-q|^2 \geq |p-t|^2 + |t-q|^2$. However, since item (iv) and item (iii) are satisfied we have

$$|p-q|^2 \geq \left(1+\gamma_e^2\right) |t-q|^2 \geq \gamma_p^2 |t-q|^2 .$$

This contradicts that move 1 can be played.

Lemma 4.3. *Let $(\mathcal{P}', \mathcal{S}', \mathcal{F})$ be the output of Theorem 4.2. Suppose that*

$$2\gamma_e \sin\phi \geq 1, \quad 2\gamma_n \sin\phi \geq \beta.$$

Then if $s \in \mathcal{S}'$ is a non-end segment, no point of $\mathcal{P}'$ subtends angle greater than 2ϕ with s. If s is an end segment, then no point of $\mathcal{P}'$ subtends angle greater than $\pi/2$ with s.

Proof. The case where s is an end segment is covered by item (v) of the theorem. Let s be a non-end segment, let q be a point not collinear to s, and let p be an endpoint of s.

If p is between s and an end segment, then by item (iii) and item (iv) of the theorem, $|p-q| \geq \gamma_e |s| \geq |s|/2\sin\phi$. If p is between s and a non-end segment, then by item (iii) and item (iv) of the theorem, $|p-q| \geq \gamma_n |s|/\beta \geq |s|/2\sin\phi$. Applying this fact to both endpoints of s (plus some simple geometry) shows that q cannot subtend angle greater than 2ϕ with s. ☐

4.2 Caps, Joins, Gates

We define a cap:

Definition 4.4 (α-cap). Given a line segment (p, q) with midpoint m, and $\alpha \in (0, \pi/2)$, the α-cap is the set of all points x such that $\angle xmp = \alpha$, or $\angle xmq = \alpha$. Such a point, x, is a α-cap point of (p, q). The α-cap consists of two circles perpendicular to (p, q). See Figure 1.

When the segment (p, q) is on the border of some understood facet F, we will unambiguously refer to the two points of intersection of the α-cap with F as *the* α-cap points.

When the segment (p, q) is inside the facet F, there may be four points of intersection; we will often consider the cap points pairwise depending on which side of the segment they are on.

Given two cap points inside a facet F, and on the same side of the segment, we let the *cap arc* be the arc of the diametral sphere of the segment between the cap points. See Figure 1(a).

Definition 4.5 (join sphere). Let $\alpha \in (0, \pi/2)$ be given. Let (p, q), (q, s) be two collinear segments sharing only the endpoint q. Then the *α-join sphere* of the two segments is the sphere centered at the midpoint of (p, s) and having radius

$$\frac{1}{2}\sqrt{|p-q|^2 + |q-s|^2 - 2\,|p-q|\,|q-s|\cos\alpha}.$$

When α is understood, the term join sphere is used alone.

The cosine rule proves that the α-join sphere of this pair of segments intersects the diametral spheres of the segments at their α-cap. When a facet F containing the two segments is understood, we call the intersection of the α-join sphere with F the *α-join arc.* This is an arc between cap points of (p, q) and (q, s). See Figure 1(b).

Claim 4.6. $\alpha \in (0, \pi/2)$ be given. Let j be a join arc between collinear segments (p, q) and (q, s), and let a_1, a_2 be cap arcs of (p, q), (q, s), respectively (as for example in Figure 1(b)). Without loss of generality, suppose

$$1 \le \frac{|p-q|}{|q-s|} \le \beta.$$

Then if $\alpha \ge \arccos \frac{1}{\beta}$ then a_1 and j meet at an obtuse angle, as do a_2 and j.

Lastly we make a simple but useful claim regarding intersection of spheres.

Claim 4.7. Let S_1 be a sphere and Let p be a point on the sphere S_1. Let S_2 be a sphere with p as center. Suppose S_1 and S_2 intersect. Then they meet at an obtuse angle.

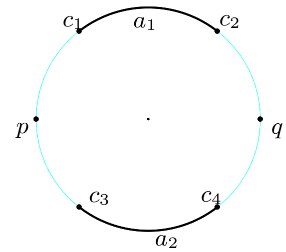

(a) cap points and cap arcs

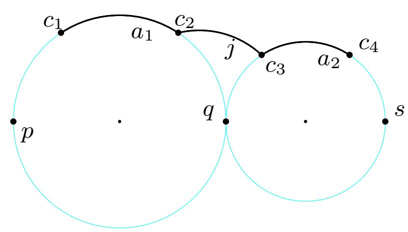

(b) cap arcs and join arc

Figure 1: In (a) the α-cap points of the segment (p, q) are shown, for some α. The cap points are labeled c_1, c_2, c_3, c_4. The arc labeled a_1 is the cap arc between c_1 and c_2, while a_2 is the cap arc between c_3 and c_4. In (b), the join arc between two adjacent segments, (p, q) and (q, s), is shown, labeled as j. In this figure a_1 is the cap arc between cap points c_1, c_2, and a_2 is the cap arc between c_3, c_4.

This will be useful because of how we deal with non-end segments.

We must deal with end segments differently than we deal with non-end segments.

Definition 4.8 (gate sphere). Let $\alpha \in (0, \pi/2)$ be given. Consider segment (q, s). Then the *gate sphere* of q is the sphere centered at q of radius $|q-s|\sin\frac{\alpha}{2}$. The gate sphere of q intersects the diametral sphere of (q, s) in the α-cap of (q, s).

Suppose (p, q) is an end segment of $\mathcal{S}'$, while (q, s) is a non-end segment. Let F be a facet containing both segments, and consider one side of the segments, in F. The *gate arc* of q is the part of the circle of intersection of the gate sphere of q that is between the α-cap of (q, s) and the sphere centered at p of radius $|p-q|$. See Figure 2.

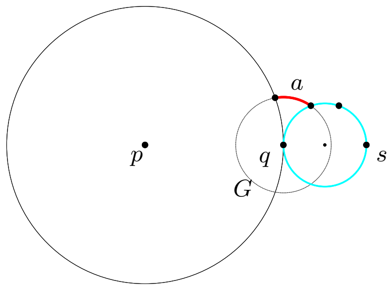

Figure 2: The gate sphere of q for the given segment (q, s) is shown as G. The gate arc of q is shown in bold as a.

4.3 A Collar

We propose the construction of a collar, much like the protecting region of Cohen-Steiner *et al.* [10], built on the groomed output of Theorem 4.2. We define a set of circular arcs which will cut up each face of the PLC into a free area, away from the boundary of the facet, and the collar area. We will show that a naïve Ruppert Algorithm applied to the free area, along with an adaptive poor-quality-tet test [5], will result in a well-graded conforming Delaunay mesh of the entire original input, with poor-quality tets in rigidly defined areas.

We will augment the set $\mathcal{P}'$ with some extra points while adding arcs to the originally empty set $\mathcal{A}$, as follows. Let F be a facet of the PLC. Then

1. If s is an end segment of $\mathcal{S}'$ that is contained in (the boundary of) F, then add the one or two cap arcs of s that intersect F to $\mathcal{A}$. Add the cap points of s that are in F to $\mathcal{P}'$.
2. If s_1, s_2 are two adjacent end segments of $\mathcal{S}'$ sharing an endpoint, and both are contained in (the boundary of) F, then add the one or two join arcs between them and which intersect F to $\mathcal{A}$.
3. If $(p, q), (q, s)$ are, respectively, an end segment and a non-end segment of $\mathcal{S}'$, with both in (the boundary of) F, then add the gate arc of q to $\mathcal{A}$, and add the endpoint of this gate arc which is not a cap point of (q, s) to $\mathcal{P}'$.
4. If p is a point of $\mathcal{P}$ in F which is the endpoint of some segments of $\mathcal{S}'$ which are in F, then consider the circle in F with center p and radius $\mathrm{d}(p)$. Some points on this circle will have been added to $\mathcal{P}'$ as endpoints of gate arcs. Consider the part of the circle between these points: cut this circle into arcs no larger than $2\pi/3$, if necessary, adding the endpoints to $\mathcal{S}'$, and putting the arcs in $\mathcal{A}$. We will call these arcs *point arcs*.
5. If p is a point of $\mathcal{P}$ in F which is the endpoint of some segment of $\mathcal{S}'$, but not the endpoint of a segment in F, then take the circle which is the intersection of F with the sphere centered at p of radius $\mathrm{d}(p)$, and cut it into (at least three) arcs smaller than $2\pi/3$. Add these arcs to $\mathcal{A}$, and add their endpoints to $\mathcal{P}'$. These arcs will also be called point arcs.

We illustrate this process by showing how it would work for some small part of an output from a groomer, see Figure 3. We stress that this illustration need not represent a typical output from the groomer (we expect the groomer would add more Steiner Points).

4.4 Lenses

We will 'protect' the arcs of $\mathcal{A}$ in a fashion similar to that of Boivin and Gooch [3]. Towards this end we will define the lens of a segment and of an arc.

Definition 4.9 (ϕ-lens). The *ϕ-lens* of a segment s is the locus of points which subtend angle $\pi - \phi$ with the endpoints of the segment. The endpoints of the segment are also agreed to be in the ϕ-lens, which makes it a closed set.

The tangents of the ϕ-lens of a segment subtend an angle of ϕ with the segment at its endpoints.

Definition 4.10 (Arc Lens). Given an arc a, of $\theta \leq \pi$ degrees of a circle, C, that has endpoints p, q, the *ϕ-arc-lens* of a is the ψ-lens of the segment (p, q), where $\psi = \phi + \frac{\theta}{2}$.

The tangents of the 'outside' part of the ϕ-arc-lens subtend an angle of ϕ with the tangents of a at its endpoints, as shown in Figure 4(a).

The *ϕ-half arc-lens* is the part of the ϕ-arc-lens that is on the 'convex' side of a; see Figure 4(a).

We will use the following facts regarding the arc-lens:

Claim 4.11. Let a be an arc on sphere S with endpoints p, q, which subtend angle $\psi < \pi$ with t, the center of the sphere. Let S^+ be the $\pi/2$-arc-lens of a. Let x be a point in S such that $\angle pxq \leq (\pi + \psi)/2$. Let S^- be the equatorial circumsphere of Δpxq, *i.e.*, the sphere with the circumcircle of Δpxq as equator. Then $S^- \subseteq S \cup S^+$. In particular this holds for $x = t$.

The proof of the previous claim uses Thales' theorem and the Laguerre Diagram [18]. The following claim is from basic geometry:

Claim 4.12. Let a be an arc of a circle C, with endpoints p, q. Let a_1 be a subarc of a. Then the ϕ-arc-lens of a_1 is contained in the ϕ-arc-lens of a.

Let x be a point in the $\pi/2$-arc-lens of an arc a with endpoints s, t then

$$\frac{|x - s| \wedge |x - t|}{|p - s|} \leq \sqrt{2}\left(\frac{\cos(\psi/4)}{\cos(\psi/4) - \sin(\psi/4)}\right),$$

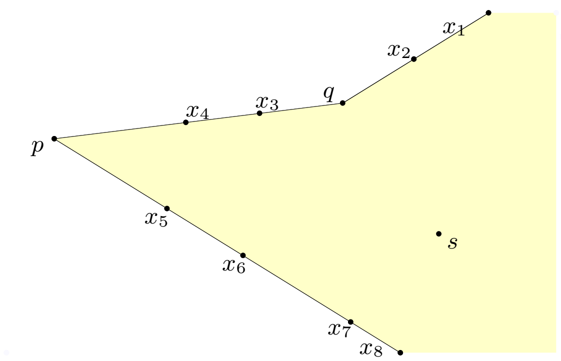

(a) groomer output

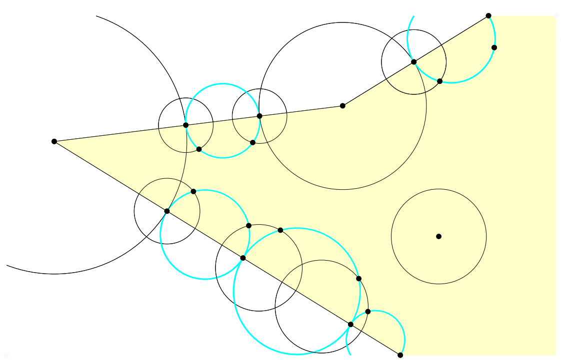

(b) diametrals and cap points

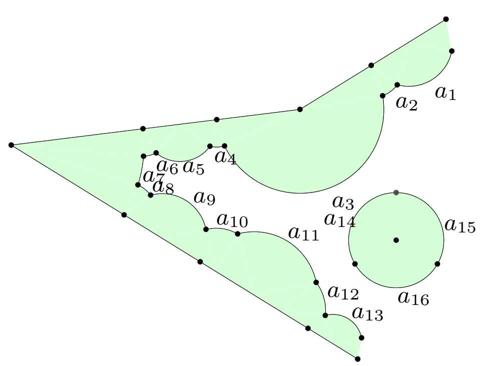

(c) collar and arcs

Figure 3: The construction of $\mathcal{A}$ from groomed output is shown; in (a), points p, q, s are in $\mathcal{P}$, the x_i were added by the groomer. The facet F containing the points is shaded. Input point s is the endpoint of a segment not in F. In (b) the diametral spheres of the non-end segments are shown, along with the join spheres, the gate spheres, and the $\mathrm{d}(\cdot)$ sized spheres around input points p, q, s. In (c), the arcs of $\mathcal{A}$ are shown: $a_1, a_5, a_9, a_{11}, a_{13}$ are cap arcs; a_2, a_4, a_6, a_8 are gate arcs; a_{10}, a_{12} are join arcs; $a_3, a_7, a_{14}, a_{15}, a_{16}$ are point arcs. The circle about s is split into 3 arcs of size $2\pi/3$. The collar is shown shaded.

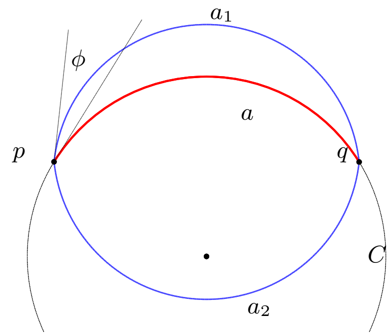

(a) arc-lens definition

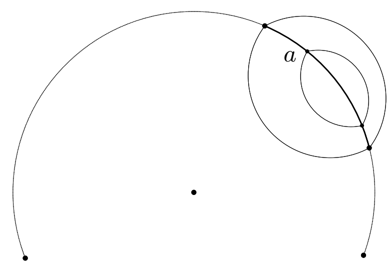

(b) subarc nesting

Figure 4: In (a) the ϕ-arc-lens of the arc a on circle C is shown as the closed set $a_1 \cup a_2$. The border of the arc-lens subtends angle ϕ with the arc. The half arc-lens is the closed arc a_1. In (b), it is shown that the ϕ-arc-lens of a subarc nests in the ϕ-arc-lens of a given arc.

where p is the midpoint of the arc, and ψ is the angle of the arc. If $\psi \leq 4 \arctan(\epsilon/(1+\epsilon))$ then the right hand side of the above bound is less than $\sqrt{2}(1+\epsilon)$.

Also

$$\frac{|x-p|}{|p-s|} \leq \frac{1}{\cos(\psi/4) - \sin(\psi/4)}$$

4.5 Collar Facts

We now look at some conditions on the grooming parameters $\alpha, \beta, \gamma_p, \gamma_e, \gamma_n$ which will ensure that the collar has desirable properties. Intuitively, it should be clear that by making $\gamma_p, \gamma_e, \gamma_n$ sufficiently large, the collar will be constructed arbitrarily close to the 1-skeleton of the input. Doing so will affect the constants C_0, C_1 of Theorem 4.2, *i.e.*, how well graded the output will be to the input local feature size. We claim that modest values of these parameters, however, suffice to construct a collar with desirable properties. We will consider the collection $(\mathcal{P}', \mathcal{S}', \mathcal{A}, \mathcal{F}')$, where $\mathcal{F}'$ consists of the faces of $\mathcal{F}$ cut into pieces by the arcs of $\mathcal{A}$. For example, the facet F of Figure 3 is divided into three pieces in Figure 3(c).

The following claim will be key in the refinement stage of the algorithm. Its proof relies on Claim 4.6, Claim 4.7, and a number of tedious geometric arguments which will not fit in the limited amount of space available here. This claim introduces the parameter η, which roughly controls the gap between collar arcs associated with disjoint input segments.

Claim 4.13 *(Collar Facts).* Let $(\mathcal{P}', \mathcal{S}', \mathcal{A}, \mathcal{F}')$ be the augmented points, the segments, the collar arcs, and the faces of $\mathcal{F}$ cut into pieces by the collar arcs. Let lfs' be defined in terms of this collection.

Let $\eta > 0$ be some chosen constant. Assume that $2 \leq \beta \leq 1 + \sqrt{2}$ and $\arccos\left(\frac{1}{\beta}\right) \leq \alpha \leq \arccos\left(\frac{1}{2\beta}\right)$, $\gamma_n \geq 0.5\beta\left(\eta + 4.163\right)$, $\gamma_e \geq 2.15 + \eta/2$, and $\gamma_p \geq \eta + 4.041$. Then

- (i) Two adjacent arcs of $\mathcal{S}'$ meet at obtuse angles.
- (ii) For any x, $\mathrm{lfs}\left(x\right) \leq C_4 \mathrm{lfs}'\left(x\right)$, where $\mathrm{lfs}\left(x\right)$ is in terms of the input PLC, and $C_4 = C_4\left(\alpha, \beta, \gamma_p, \gamma_e, \gamma_n, \eta, \theta^*\right)$.

5. REFINEMENT STAGE

The refinement is an iterative process. At each step, the algorithm can choose to play any of a number of moves until no move can be played, at which point the algorithm terminates. The algorithm maintains the collection $(\mathcal{P}', \mathcal{S}', \mathcal{A}, \mathcal{F}')$. At times the an arc of $\mathcal{A}$ will have to be split, that is, the arc will be replaced by two subarcs, and the common endpoint of these subarcs will be added to $\mathcal{P}'$. Arcs are split at their "midpoints," *i.e.*, they are "split on sphere" [10]. Since the initialization phase may have given segments which are too *large* compared to local feature size, $\mathrm{lfs}_2(\cdot)$, segments may have to be split as well. This will occur when it is discovered, for example, that a segment is very near a disjoint plane. In this case, the arcs associated with the segment will have to be altered. We employ the simple strategy of removing them altogether.

We describe how we will protect the collar region. By collar region, we mean the union of the cap spheres, join spheres, gate spheres, and the $\mathrm{d}(\cdot)$-sized spheres around input points. If a point p is proposed for addition to $\mathcal{P}'$, and p is strictly inside the collar region (but p is not on an input segment), then a segment will be split: if p encroaches a non-end segment, then split that segment; if p encroaches an end segment associated with input point p, then split all the end segments with p as endpoint, effectively halving $\mathrm{d}(p)$; if p is inside a join sphere, then split the larger segment associated with the sphere; if p is inside a gate sphere, then split the non-end segment associated with the gate sphere.

In the case where end segments are split, some of the newly created non-end segments will have to be split as well to ensure that arcs do not overlap one another. In particular, segments may have to be split to enforce the β-balance and γ-isolation conditions (see item (iii) and item (iv) of Theorem 4.2). In fact, we see the segment re-refinement process as a local execution of the algorithm described in the proof of Theorem 4.2. The γ-isolation conditions only need to be enforced with respect to points on the 1-skeleton of the input.

The arcs will have to be fixed to reflect the change in their generating segments. Thus, if an end segment is split, gate arcs and point arcs will be removed from $\mathcal{A}$, and new ones associated with the new end segments will be added, as well as join arcs and cap arcs associated with the new non-end segments. If a non-end segment is split, cap arcs and join arcs and possibly gate arcs will be removed from $\mathcal{A}$, and replaced by new ones. This has the effect of altering $\mathcal{F}'$. Additionally, if arcs are removed, then any points of $\mathcal{P}'$ on removed arcs will be removed from $\mathcal{P}'$.

We claim that if the splitting of the segments is justified, then we preserve local feature size in the updated collection $\mathcal{F}'$.

Claim 5.1 *(Updated Collar LFS).* Suppose that a segment is split because the collar region was cut by a plane disjoint from the input segment containing the split segment. Then $\mathrm{lfs}'\left(x\right) \leq C_4 \mathrm{lfs}\left(x\right)$, where $\mathrm{lfs}'\left(x\right)$ is with respect to the new collection $\mathcal{F}'$, after all segment splits and arc updates have been performed.

We now give the rules for the refinement stage. We believe that implementation difficulties will be minimized when the moves are played with preference given in the listed order, with the first move preferred over the second, preferred over the third, etc. This also simplifies their description to some extent.

1. If $p \in \mathcal{P}'$ encroaches the collar, then update the collar as described above. Note that if this move is to be played, then p is a point on a facet.
2. If $p \in \mathcal{P}'$ encroaches the $\pi/2$-half arc-lens of an arc in $\mathcal{A}$, split the arc.
3. If p, q, s are points in the same facet of $\mathcal{F}'$, and have the Delaunay Property with respect to all points of $\mathcal{P}'$ in this facet, and there is some point of $\mathcal{P}'$ in the equatorial circumsphere of the triangle Δpqs, then attempt to add the circumcenter of the triangle to the set $\mathcal{P}'$. If this circumcenter encroaches the $\pi/2$-half arc-lens of an arc in $\mathcal{A}$, then split that arc instead.
4. Suppose p, q, s, t are points of $\mathcal{P}'$, the tetrahedron $pqst$ is Delaunay with respect to $\mathcal{P}'$ and has circumradius to shortest edge ratio greater than μ. Then attempt to add the circumcenter of the tetrahedron, call it x, to $\mathcal{P}'$. However, if the addition of x would trigger move 3 or move 2, then abort the addition of x and play that other move instead. Moreover, if x is inside the collar region,

then abort the addition of x, and do not attempt to remove the tet again.

5.1 Good Grading

Theorem 5.2. *Suppose $\mu > 2$. Then there are constants G_1, G_2, G_3 depending only on C_4, θ^*, μ such that if point p is proposed for addition to the set $\mathcal{P}'$ then*

1. *if p is the center of an arc, $\text{lfs}(p) \leq G_1 |p-q|$, where q is the closest point to p in $\mathcal{P}'$.*
2. *if p is the circumcenter of a triangle, then $\text{lfs}(p) \leq G_2 |p-q|$, where q is the closest point to p in $\mathcal{P}'$.*
3. *if p is the circumcenter of a tet, then $\text{lfs}(p) \leq G_3 |p-q|$, where q is the closest point to p in $\mathcal{P}'$.*

By "proposed for addition," we explicitly exclude the case of the circumcenter of a tet which is rejected because it encroaches the collar region.

Proof. We take $0 < \epsilon < (\mu/2) - 1$.

1. Let p be the center of arc a, with endpoints t, s. Let q be a point in $\mathcal{P}'$ closest to p. If q is on a disjoint facet of $\mathcal{F}'$; then $\text{lfs}(p) \leq C_4 \text{lfs}'(p) \leq C_4 |p-q|$, and we simply take $C_4 \leq G_1$. By Claim 4.13, the arc cannot have been encroached by a q on an adjacent nondisjoint arc, because they meet at obtuse angles.
   If q is outside the half arc-lens of a, then p is added because of a tet or triangle circumcenter encroaching the half arc-lens. Thus there was some point p' such that, inductively, $\text{lfs}(p') \leq G_i (|p'-s| \wedge |p'-t|)$, for $i = 2$ or 3. We should note that in this case it must be that q is either s or t.
   Let a be an arc of angle ψ. If $\psi \geq \pi/3$, then a is an arc made by the collar construction. In this case $\text{lfs}(p) \leq C_4 \text{lfs}'(p) \leq C_4 |p-s|$, thus we can take $C_4 \leq G_1$.
   Similarly suppose that $\psi \geq 2\pi/(2^k 3)$, where k is the smallest integer such that
   $$\frac{2\pi}{2^k 3} \leq 4 \arctan \frac{\epsilon}{1+\epsilon}.$$
   Because the collar construction makes no arcs greater than $2\pi/3$, a must be the result of no more than k splits of an arc a' created by the collar construction. Let that arc have endpoints s', t'. Because no more than k splits have occurred,
   $$\text{lfs}'(p) \leq |p-s'| \vee |p-t'| \leq 2^k |p-s|.$$
   We now have
   $$\text{lfs}(p) \leq C_4 \text{lfs}'(p) \leq 2^k C_4 |p-s|.$$
   Thus we need to take $2^k C_4 \leq G_1$.
   Otherwise $\psi \leq 2\pi/(2^k 3) \leq 4 \arctan(\epsilon/(1+\epsilon))$. We will use Claim 4.12, and assume that $|p'-s| \leq |p'-t|$:
   $$\begin{aligned} \text{lfs}(p) &\leq |p-p'| + \text{lfs}(p') \\ &\leq \frac{1}{\sqrt{2}} |p-s| + G_i |p'-s| \\ &\leq \left(\frac{1}{\sqrt{2}} + \sqrt{2}(1+\epsilon) G_i\right) |p-s|. \end{aligned}$$
   So it suffices to take $(1/\sqrt{2}) + \sqrt{2}(1+\epsilon) G_i \leq G_1$, for $i = 2, 3$.
   It could be the case that q is inside the half arc-lens. This can only happen if a is newly created because the collar is resized. In this case, however, it can be shown that the above argument regarding p' can be applied to q. The only wrinkle is that s, t may not have been in $\mathcal{P}'$ when p' was proposed for addition to the mesh.
2. Let p be the circumcenter of a triangle that is split. Let q be the closest point to p. If q is inside the equatorial sphere of the triangle, then q must be on a nondisjoint feature. Thus $\text{lfs}(p) \leq C_4 \text{lfs}'(p) \leq C_4 |p-q|$. We can take $C_4 \leq G_2$. On the other hand, if q is not inside the sphere, then the triangle was split by a tet circumcenter. Using a standard argument regarding projections [7], we can claim that it suffices to take $1 + \sqrt{2} G_3 \leq G_2$.
3. Let p be the circumcenter of a tet. We use the standard argument concerning the identity of the last added point of the short edge of the offending tet. We claim it suffices to take
   $$1 + G_i/\mu \leq G_3, \quad \text{for } i = 1, 2, 3.$$

It suffices to take

$$G_1 = 2^k C_4 \vee \frac{4.13\mu}{\mu - 2(1+\epsilon)},$$
$$G_3 = 1 + G_1/\mu, \qquad G_2 = 1 + \sqrt{2} G_3,$$

where $k = \mathcal{O}(-\log \epsilon)$, and $0 < \epsilon < 1 \wedge (\mu/2) - 1$. ☐

Thus the algorithm can use any $\mu > 2$, but with smaller values result in larger grading constants, which is to be expected: if μ is small more tets will be split. This output guarantee is an improvement over Cheng and Poon's value of 16, and matches the standard value of 2 achievable for input satisfying the obtuse angle condition [11, 7]. This improvement is had due to the argument concerning the number of splits a collar arc has suffered: if the number is small, the arc is similar enough to the constructed collar to use the starting bounds; if the number is large, then the arc is very nearly a straight line segment, and its arc-lens is nearly a diametral sphere.

Note that the analysis is somewhat imprecise regarding where poor quality tets may be left in the output mesh: these tets must have circumcenters inside the collar region. It may be the case that they can be proven to be relatively near small angles of the input, and that their radius-edge ratio is bounded by the size of said small angle, as is the case in 2D [5]. However, the analysis appears difficult.

Care must be taken when implementing the algorithm to ensure that a poor quality tet, one which cannot be killed because its circumcenter is in the collar region, does not remain in the mesh if the collar region is subsequently refined such that the circumcenter *can* be added to the mesh. We believe this problem can be avoided if the moves of the refinement stage are played with preference in the given order; an extra argument is required for this, one which we have not yet completed.

5.2 Output Conditions

We now make the claim that the output from the refinement stage is a Conforming Delaunay Tetrahedralization which respects the input to the algorithm.

If s is a segment of $\mathcal{S}$, then it is represented as the union of several segments, s_i, in $\mathcal{S}'$. Because the collar region is empty, then the diametral sphere of each s_i is devoid of points of $\mathcal{P}'$. Thus each s_i is strongly Delaunay with respect to the output point set.

Let a be an arcs of $\mathcal{A}'$, with endpoints p, q. Because the half arc-lens is empty of points of $\mathcal{P}'$, and the collar region only contains specifically located segment endpoints, it can be shown that there is a sphere with (p, q) as chord that is empty of points of $\mathcal{P}'$, and so the edge is Delaunay. This relies on the fact that arcs of $\mathcal{A}'$ are no bigger than $2\pi/3$, and due to the special structure of join arcs.

Let F be a facet of $\mathcal{F}$; then F is represented by at least two facets of $\mathcal{F}'$, perhaps more if there are internal boundaries or included points. Only one of these is a non-collar subfacet. Let F_0 be the non-collar subfacet. Let Δpqs be a triangle with points in F_0, and Delaunay with respect to the points of $\mathcal{P}'$ in F_0. Then the triangle is strongly Delaunay with respect to $\mathcal{P}'$, as otherwise the algorithm could play move 3.

Now we consider F_1 a collar subfacet. Let a be an arc on the boundary of F_1. As per above, the segment associated with a, call it (p, q) has the Delaunay property. We claim that, by a straightforward yet tedious application of Claim 4.11, to the possible cases of the type of arc of $\mathcal{A}$ that a lies on, shows that the segments (p, s) and (q, s) have the Delaunay property, where s is the associated input point if a is on a point arc, the center of the arc if a is on a gate arc, and the obvious midpoint if a is on a join arc. When a is on a cap arc then (p, s),(q, s) are Delaunay when s is either endpoint of the associated segment; that is, there is a degeneracy, and some care must be taken when implementing the algorithm.

Moreover, we claim that by Claim 4.11, any triangle with corners in F_1 which is Delaunay with respect to the points of this facet is Delaunay with respect to the whole point set. And thus F_1 is represented as the union of Delaunay triangles in the output mesh.

5.3 Choice of Parameters

Based on their theoretical relationships to the grading constants, G_i, we propose the following choices for the user chosen parameters:

$$\eta = \sqrt{2} - 1, \quad \beta = 2.04, \quad \alpha = \arccos(1/2\beta)$$
$$\gamma_p = 4.46, \quad \gamma_e = 4.35, \quad \gamma_n = 4.67$$

For this choice of parameters and with $(\theta^*, \mu) = (\pi/40, 2.4)$, we can make the following bounds on the grading constants:

$$G_3 \leq G_2 \leq G_1 \leq 5.8 \times 10^5$$

When $(\theta^*, \mu) = (\pi/6, 2.25)$, G_1 drops to 2.9×10^4.

While we admit these constants are large, this is normally the case for grading constants [1, 19, 5]. We make the usual argument that the constants are based on several worst case assumptions, and in practice should be much smaller. However, it is not inconceivable that the practical grading constants may be large enough to make the algorithm impractical or useless.

6. FUTURE WORK

While a number of theoretical questions remain regarding this algorithm, the most pressing work, we reiterate, is to implement it. We believe that testing the algorithm will: guide the choices of the user-chosen parameters, η, β, etc; lead to simplifications of the algorithm itself; test the hypothesis that *a priori* calculation of local feature size is an inferior strategy; give some indication of the size of the practical grading constants; and, finally, indicate whether the algorithm is practical.

We also believe that, as with the algorithm of Cohen-Steiner *et al.*, some relaxation of the collar region may be possible at facets that meet at large (larger than $\pi/3$) angles [10]. This is a fruitful avenue of exploration, and may reduce the cardinality of meshes output by the algorithm. Alternatively, it may be possible to eliminate the collar altogether near segments which are part of only two input facets. That is, to employ the strategy of Cheng *et al.* where the input looks like a 2-manifold [12].

References

[1] Ruppert J. "A Delaunay refinement algorithm for quality 2-dimensional mesh generation." *J. Algorithms*, vol. 18, no. 3, 548–585, 1995. Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) (Austin, TX, 1993)

[2] Chew L.P. "Guaranteed-quality triangular meshes." Tech. Rep. 89-983, Computer Science Department, Cornell University, 1989

[3] Boivin C., Ollivier-Gooch C.F. "Guaranteed-Quality Triangular Mesh Generation For Domains with Curved Boundaries." *International Journal for Numerical Methods in Engineering*, vol. 55, no. 10, 1185–1213, 2002

[4] Shewchuk J.R. "Mesh generation for domains with small angles." *Proceedings of the Sixteenth Annual Symposium on Computational Geometry (Hong Kong, 2000)*, pp. 1–10 (electronic). ACM, New York, 2000

[5] Pav S.E. *Delaunay Refinement Algorithms*. Ph.D. thesis, Department of Mathematics, Carnegie Mellon University, Pittsburgh, Pennsylvania, May 2003. URL `http://www.andrew.cmu.edu/~spav/work`

[6] Miller G.L., Pav S.E., Walkington N.J. "When and Why Ruppert's Algorithm Works." *Proceedings of the 12th International Meshing Roundtable*, pp. 91–102. Sandia National Laboratory, September 2003

[7] Shewchuk J.R. *Delaunay Refinement Mesh Generation*. Ph.D. thesis, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, May 1997. Available as Technical Report CMU-CS-97-137

[8] Miller G.L., Pav S.E., Walkington N.J. "Fully Incremental 3D Delaunay Mesh Generation." *Proceedings of the 11th International Meshing Roundtable*, pp. 75–86. Sandia National Laboratory, September 2002

[9] Murphy M., Mount D.M., Gable C.W. "A point-placement strategy for conforming Delaunay tetrahedralization." *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, pp. 67–74. Society for Industrial and Applied Mathematics, 2000

[10] Cohen-Steiner D., Colin de Verdière E., Yvinec M. "Conforming Delaunay Triangulations in 3D." *Proc. 18th Annu. ACM Sympos. Comput. Geom.* 2002

[11] Cheng S.W., Poon S.H. "Graded conforming Delaunay tetrahedralization with bounded radius-edge ratio." *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 2003)*, pp. 295–304. ACM, New York, 2003

[12] Cheng S.W., Dey T.K., Ramos E.A., Ray T. "Quality meshing for polyhedra with small angles." *Proceedings of the Twentieth Annual Symposium on Computational Geometry (Brooklyn, New York)*. ACM, June 2004

[13] Miller G.L., Talmor D., Teng S., Walkington N.J. "On the Radius–Edge Condition in the Control Volume Method." *SIAM J. Numer. Anal.*, vol. 36, no. 6, 1690–1708, 1998

[14] Edelsbrunner H., Li X., Miller, Stathopoulos G., A. Talmor D., Teng S., Ungor A., Walkington N.J. "Smoothing and cleaning up slivers." *ACM Symposium on Theory of Computing*, pp. 273–277. 2000. URL `citeseer.nj.nec.com/edelsbrunner00smoothing.html`

[15] Cheng S.W., Dey T.K., Edelsbrunner H., Facello M.A., Teng S.H. "Sliver exudation." *Proceedings of the fifteenth annual symposium on Computational geometry*, pp. 1–13. ACM Press, 1999

[16] Li X.Y., Teng S.H. "Generating well-shaped Delaunay meshes in 3D." *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pp. 28–37. ACM Press, 2001. URL `citeseer.ist.psu.edu/li01generating.html`

[17] Cheng S.W., Dey T.K. "Quality meshing with weighted Delaunay refinement." *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 137–146. Society for Industrial and Applied Mathematics, 2002

[18] Okabe A., Boots B., Sugihara K., Chiu S.N. *Spatial tessellations: concepts and applications of Voronoi diagrams*. John Wiley & Sons Ltd., Chichester, second edn., 2000. With a foreword by D. G. Kendall

[19] Mitchell S.A. "Cardinality Bounds for Triangulations with Bounded Minimum Angle." *Sixth Canadian Conference on Computational Geometry*, pp. 326–331. 1994