

# Size Competitive Meshing without Large Angles

★

Gary L. Miller   Todd Phillips   Donald Sheehy

Computer Science Department,  
Carnegie Mellon University, Pittsburgh, PA 15213  
{glmiller, tp517, dsheehy}@cs.cmu.edu

**Abstract.** We present a new meshing algorithm for the plane, Overlay Stitch Meshing (OSM), accepting as input an arbitrary Planar Straight Line Graph and producing a triangulation with all angles smaller than  $170^\circ$ . The output triangulation has competitive size with any optimal size mesh having equally bounded largest angle. The competitive ratio is  $O(\log(L/s))$  where  $L$  and  $s$  are respectively the largest and smallest features in the input. OSM runs in  $O(n \log(L/s) + m)$  time/work where  $n$  is the input size and  $m$  is the output size. The algorithm first uses Sparse Voronoi Refinement to compute a quality overlay mesh of the input points alone. This triangulation is then combined with the input edges to give the final mesh.

## 1 Introduction

The meshing problem is to take as input a domain containing a collection of features and return a triangulation of the domain. In 2D, the features are simply points and non-crossing edges; a planar straight line graph (PSLG). The design of a meshing algorithm involves analysis of four fundamental properties of the algorithm and its outputs. First, the output mesh must be **conforming**, so that all the vertices and edges should appear as a union of simplices in the final mesh. Secondly, all the triangular elements should have some guarantee of **element quality**. Third, the number of output triangles should be asymptotically competitive with any optimal triangulation that is conforming and good quality, so that we have an output size **competitive algorithm**. Finally the algorithm should be **work efficient**. This paper will only be concerned with the 2D meshing problem.

The 2D meshing problem was first posed by Bern, Eppstein, and Gilbert[1] who proposed a quadtree algorithm. Ruppert[2] gave an  $O(n^2)$  time constant size competitive algorithm for the meshing problem using Delaunay refinement. Mitchell and Vavasis[3] extended the quadtree algorithm to 3D and proved that the Bern, Eppstein, and Gilbert algorithm was in fact also a constant size competitive algorithm. In order for these algorithms to be constant factor competitive algorithms, two critical assumptions are made. First, element quality is

---

\* This work was supported in part by the National Science Foundation under grants CCR-0122581.

defined by the absence of arbitrarily small angles in any triangle. Secondly, an input PSLG with small angles is not permitted. Our main goal is to develop algorithms and techniques that will allow an arbitrary PSLG as input, relaxing the latter assumption. Note that this also forces us to abandon the former assumption, since any small input angle will require new small output angles to conform [4].

An alternative setting employs a definition of element quality that allows small angles, but bounds all angles strictly away from  $180^\circ$ , prohibiting arbitrarily large angles. Many meshing applications require only this weaker quality guarantee. Babuška and Aziz showed that only large angles affect interpolation error [5, 6], while Boman, Hendrickson, and Vavasis showed that only large angles affect their reduction of an elliptic problem to the solution of a Laplacian problem [7, 8]. An algorithm operating under weaker quality constraints is accordingly able to produce meshes with fewer added vertices (in theory and practice).

In this work, we present and analyze the Overlay Stitch Meshing algorithm (OSM), a new no-large-angle meshing algorithm for conforming to an arbitrary PSLG in 2D. For practical inputs, OSM achieves the best known competitive guarantees on output size relative to the optimal no-large-angle conforming mesh. The algorithm outputs a mesh with no angle larger than  $170^\circ$ . In addition to having good theoretical guarantees, the algorithm is remarkably straightforward and runs in time  $O(n \log(L/s) + m)$ .

## 1.1 Preliminaries

We will use a few standard definitions in our analysis.

**Definition 1.** The **spread**,  $L/s$ , of a PSLG is the ratio of the input diameter ( $L$ ) to the smallest distance between any two disjoint input features ( $s$ ).

**Definition 2.** The **gap-ratio**,  $\Gamma_M(x)$ , at a location  $x$  relative to a point set  $M$ .  $\Gamma_M(x) = R(x)/r(x)$ , where  $R(x)$  is the radius of the largest disc containing  $x$  but not intersecting  $M \setminus x$ , and  $r(x)$  is the nearest neighbor of  $x$  in  $M \setminus x$ . At the boundary of our geometric domain, we require that  $x$  and the center of the largest disc be contained in the convex closure of  $M$ . For shorthand, we say a mesh  $M$  has gap-ratio quality  $\Gamma$  if  $\forall x, \Gamma_M(x) \leq \Gamma$ .

Note that if some mesh  $M$  has some constant gap-ratio quality  $\Gamma$ , then it has no arbitrarily small or large angles.

**Definition 3.** The **local feature size**,  $\text{lfs}_M(x)$ , at a point  $x$  relative to a PSLG  $M$  is given by the radius of the smallest disc centered at  $x$  that intersects two disjoint features (segments or vertices) of  $M$ . When  $M$  is suppressed,  $\text{lfs}(x)$  will be with reference to the input PSLG. Additionally,  $\text{lfs}_0(x)$  is the radius of an identically defined disc intersecting two vertices of the input, ignoring segments.

Given these definitions, we state our main result:

**Theorem 1.** *On input a PSLG Algorithm Overlay Stitch Meshing generates a conforming triangulation  $\mathcal{T}$  with no angle greater than  $170^\circ$ . The size of  $|\mathcal{T}| \leq c \log(L/s) \cdot OPT$  where  $OPT$  is the size of an optimal conforming triangulation with all angle bound away from  $180^\circ$  and  $c$  is some fixed constant.*

## 1.2 Related Work

The present work derives from two disjoint lines of past research, well-graded meshes and no-large-angle triangulation. Both lines of research trace their lineage back to the same motivating problem, that of producing quality meshes for finite-element simulations. However, the methods, guarantees, and final output provided by each field are drastically different.

Well-graded meshing algorithms attempt to produce meshes in which the length of each edge is proportional to the lfs at its end points. These algorithms generally fall into two categories, structured and unstructured, as typified by Quadtree methods [1] and Delaunay refinement methods [2, 4] respectively.

One significant property of most well-graded meshing algorithms is that they provide guarantees regarding both the largest and the smallest angles in the resulting mesh. A bound on the smallest angle is a sufficient but not necessary condition for a mesh to have no large angles. However, many Delaunay refinement algorithms are only guaranteed to terminate for inputs where all input angles are at least  $60^\circ$ . Several algorithms avoid this restriction by severely weakening the smallest output angle guarantee in order to guarantee termination [9].

Significant research has been done on how to extend well-graded meshing algorithms to efficiently handle small input angles, most of which is based on extending the “concentric shelling” method proposed by Ruppert in [2]. The main idea requires that if two edges share a vertex, then the splitting of these edges should be identical, effectively “protecting” the region inside the small angle. Extensions of the protective region approach to 3D have so far been quite involved, and output size guarantees fairly weak [10, 11, 12]. The immediate consequence of concentric shelling is that the size and grading of the output mesh depends on the smallest angle in the input PSLG.

Even in the absence of small input angles, well-graded meshing techniques often produce meshes whose size is linearly dependent on the spread ( $L/s$ ) of the underlying vertex input set. Generally, the spread is assumed to be some polynomial in the size of the input.

Because of the blowup in size associated with well-graded meshing, much research has focused on the problem of eliminating large angles. Bern, Dobkin, and Eppstein give an algorithm that produces a no-large-angle mesh of a simple polygon using  $O(n \log n)$  triangles; or  $O(n^{3/2})$  triangles for polygons with holes [13]. This result was later improved by Bern, Mitchell, and Ruppert who gave an algorithm that produces a nonobtuse triangulation of a polygon with holes using  $O(n)$  triangles [14]. For arbitrary planar straight line graphs, there is a lower bound attributed to Paterson that says  $\Omega(n^2)$  triangles may be necessary.

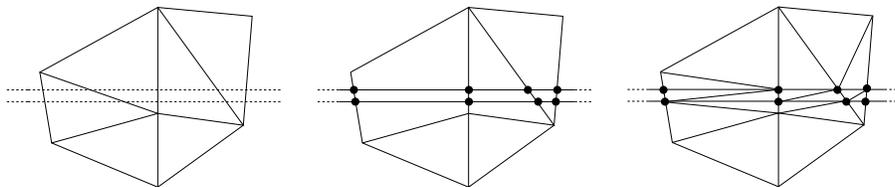
An algorithm of Mitchell [15] produces a triangulation with bounded angles and at most  $O(n^2 \log n)$  vertices. This was later improved by Tan [16] who im-

proved the angle bounds to  $\frac{11}{15}\pi$  and the size bounds to  $O(n^2)$  vertices. The latter result is within a constant factor of worst-case optimal. Both of these algorithms use the method of propagating paths in which a bad angle is subdivided, possibly creating a new large angle elsewhere. The methods and analysis depend on bounding the length of the paths of bad angles as they propagate through the triangulation. We present a much simpler algorithm that gives a triangulation with angles bounded below  $170^\circ$  of size  $O(\log L/s)$ -competitive with any triangulation achieving this angle bound. For inputs that admit a small no-large-angle triangulation, this improves on the worst-case optimal bound of  $O(n^2)$ . It is possible that propagating path algorithms also produce competitively sized meshes but a competitive analysis has proven difficult, even using the lower bound techniques of this paper.

The main technology used in OSM is a well-graded *overlay* mesh on just the input vertices. The input edges are then *stitched* in. The use of an axis-aligned overlay grid was introduced by Baker et al [17] and later extended by Dey [18]. The same paradigm can be traced in structured meshing algorithms that use a quadtree as an overlay and then stitch in the input edges [1]. In this paper, we use an unstructured triangular mesh as our overlay.

## 2 The OSM Algorithm

The algorithm has three phases: the Overlay mesh phase, the Stitching phase, and the Completion phase. In the first phase, a standard point set meshing algorithm is run on the input vertices to form the *overlay mesh*. In the second phase, the input segments are *stitched* into the mesh, carefully choosing to add vertices at some intersections of the overlay mesh and the input segments. The Stitching phase will leave some non-triangular faces. In the last phase, the Completion phase, these leftover faces are triangulated to minimize the largest angle. See Figure 1.



**Fig. 1.** The basic stages of overlay meshing: At left, an overlay mesh is generated that does not yet conform to input segments (shown dashed). Center, the input segments are stitched into the mesh by adding some intersections and discarding some overlay edges. At right, the triangulation is completed.

### 2.1 Phase 1: The Overlay Mesh

The overlay mesh is constructed on the input vertices using the Sparse Voronoi Refinement meshing algorithm, applied to point sets [19]. The output is a De-

launay mesh conforming to the input vertices. The overlay mesh has gap-ratio quality  $\Gamma$ , for a constant parameter  $\Gamma > 1$ . No angle in the overlay mesh will be smaller than  $\theta_1 = \arcsin(1/2\Gamma)$ . The size of the mesh (in terms of vertices) will be guaranteed to be  $O\left(\int_{\Omega} \frac{1}{(\text{fs}_0(z))^2} dz\right)$ .

The choice of SVR as a meshing algorithm for the overlay mesh is not necessarily critical for Overlay Stitch Meshing. SVR is used because of the strong runtime guarantees and the tightest analyzed bounds on  $\Gamma$  when applied to point sets. Any point set meshing algorithm could be used, so long as it outputs in a Delaunay mesh with guarantees on  $\Gamma$  and output size. The usual method is to set  $\Gamma = 1 + \varepsilon$  for some small constant  $\varepsilon$ , yielding a  $\theta_1$  slightly less than  $30^\circ$ .

## 2.2 Phase 2: Stitching in Edges

We now wish to begin conforming to the input segments. We can look at all the intersections between overlay edges and input segments, and we will classify every intersection as good or bad. A good intersection is approximately perpendicular, meaning that the segment and overlay edge meet at angles larger than  $\theta_1$ . A bad intersection is approximately parallel (angles smaller than  $\theta_1$ ).

If an overlay edge crosses no segment, or crosses any segment with a good intersection, it will be kept. Overlay edges that cross segments at solely bad intersections will be discarded. The intuition here is that if an edge of the overlay mesh intersects input segments only in a parallel fashion, then we can use the input segments instead, so we throw out the overlay edge.

We have kept some overlay edges that cross input segments. We will then add Steiner points that subdivide input segments where they intersect these crossing edges (at good or bad intersections). We also now add all the subdivided input segments to the mesh. The result is not, in general, a triangulation, although it does now conform to the input.

### Stitching in edges

```

1: for all input segments  $e$  do
2:   Compute intersections of  $e$  with overlay mesh edges
3:   if some overlay edge  $e'$  intersects  $e$  at an angle greater than  $\theta_1$  then
4:     mark  $e'$  as kept.
5:   /* Insert the 'good' intersections */
6:   for all intersections do
7:     if the overlay edge at the intersection is marked as kept then
8:       insert the intersection point into the mesh splitting corresponding input segment.
9:     else
10:      Remove the overlay edge at the intersection from the mesh.
11:    /* Recover the input */
12:   for all input subsegments  $e$  do
13:     insert  $e$  into the mesh

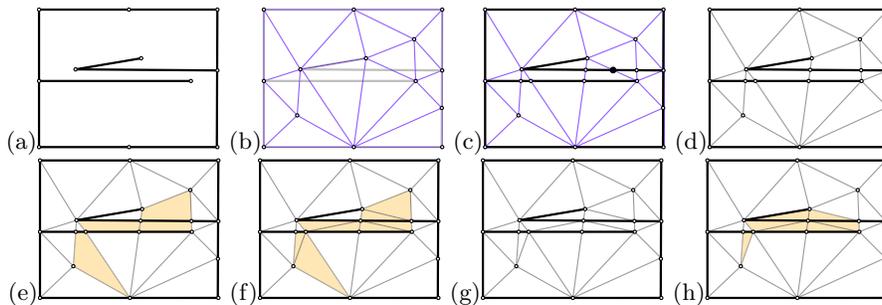
```

### 2.3 Phase 3: Completing the Mesh

After the edges are stitched in, all that remains is to add enough edges to get back to a triangulation. Each remaining non-triangular face is triangulated to minimize the maximum angle. Lemma 1 implies that this last step can naively be done efficiently because all of the non-triangular faces have at most six sides. Theorem 2 will show that these faces can always be triangulated to have no large angles.

**Lemma 1.** *After the stitching phase, all faces have at most 6 sides.*

A rigorous proof of Lemma 1 would proceed by first considering the case where no overlay edges are discarded. Then we can inductively remove the discarded edges, preserving the bound of six sides to every face.



**Fig. 2.** A sample run of OSM. In box (a), an input PSLG is given. The input contains both a small angle and a narrow pair of edges. In (b), the first phase of OSM constructs a quality overlay mesh on the input points, ignoring the input segments, and adding Steiner points where necessary. In (c), the input is intersected with the overlay mesh and one bad intersection is identified. Box (d) shows the (not yet triangular) mesh after the stitching phase. The remaining non-triangular faces highlighted in (e) are triangulated in the completion phase (f). The final mesh is shown in (g) and the remaining small angle triangles are shown in (h).

## 3 Output Angle Guarantees

The purpose of this section is to prove that the maximum angle in the output mesh is bounded from above by a constant  $\pi - \theta_2$  that is independent of the smallest input angle. Here,  $\theta_2$  is a constant depending only on  $\theta_1$ , the smallest angle in the overlay mesh. This result is stated formally below as Theorem 2. If our overlay mesh has  $\Gamma = 1 + \varepsilon$ , then we find that  $\theta_1 \geq 29.9^\circ$ , and  $\theta_2 = \arcsin(1/(4 + 2 \cos \theta_1 + O(\varepsilon))) \geq 10^\circ$ . First, we prove an important lemma regarding the angles at vertices inserted during the stitching phase.

**Lemma 2.** *If an overlay mesh edge  $e$  intersects an input segment  $f_1$  at an angle greater than  $\theta_1$  then no input segment  $f_2$  can intersect  $e$  at an angle less than  $\theta_2 = \arcsin\left(\frac{1}{2\Gamma(2\Gamma + \cos\theta_1)}\right)$ .*

*Proof.* Let  $F_1$  and  $F_2$  be the lines containing edges  $f_1$  and  $f_2$  respectively, and let  $v = F_1 \cap F_2$ . The statement is trivial if  $F_1, F_2$  are parallel or identical, so we may assume that  $v$  is a single point.

Without loss of generality, let us assume that the edge  $e$  is horizontal and  $v$  lies below the line containing  $e$ . Also, by symmetry, we may assume that the small angle intersections we are worried about are on the left.

Figure 3 (Left) shows the edge  $e$  with the circumcircle  $C_1$  of a triangle containing it. The gap ratio guarantees that there is an empty ball  $C_2$  centered at  $a$  with radius at least  $1/\Gamma$  times the radius of  $C_1$ . The point  $d$  is the lower intersection of these two circles. In the figure  $\theta_1 = \angle abd$ . The line  $L_2$  contains  $\overline{bd}$  and the line  $L_1$  is parallel to  $L_2$  passing through  $a$ . The point  $c$  is the lower intersection of  $L_1$  and  $C_2$ .

Since  $f_1$  passes through  $e$  at an angle greater than  $\theta_1$ ,  $v$  must lie below  $L_1$ . In order for  $f_2$  to intersect  $e$  at an angle less than  $\theta_1$ ,  $v$  must lie above  $L_2$ . Thus we see that the intersection  $v$  must lie somewhere in the shaded region. Therefore the smallest possible angle between  $f_2$  and  $e$  occurs if  $v = c$  and the angle is:

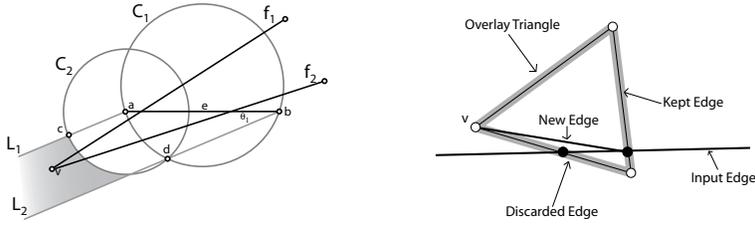
$$\angle abc \geq \arcsin\left(\frac{\sin\theta_1}{2\Gamma + \cos\theta_1}\right) = \arcsin\left(\frac{1}{2\Gamma(2\Gamma + \cos\theta_1)}\right)$$

The preceding Lemma ensures that if OSM chooses not to discard an edge of the overlay mesh, then that edge will not create any large angles. We can now proceed to the main theorem about the output angle guarantees.

**Theorem 2.** *All angles in the final mesh are bounded from above by  $\pi - \theta_2$ , where  $\theta_2 = \arcsin\left(\frac{1}{2\Gamma(2\Gamma + \cos\theta_1)}\right)$  is the lower bound on the angle of intersection between a kept edge of the overlay mesh and an input edge. For  $\Gamma = 1 + \varepsilon$ , this gives a largest angle of at most  $170^\circ$ .*

*Proof.* We need to consider the angles at two types of vertices, those vertices that appear in the overlay mesh and those that are added during the stitching phase. In the latter case, we were careful only to add vertices when the angles of intersection met this criterion. Thus, although we may have added more edges incident to such vertices in order to get a triangulation, the maximum angle at such vertices clearly achieves the desired bound. In the case of a vertex  $v$  from the overlay mesh, we have to be a little more careful. The overlay mesh itself satisfied this property so all large angles at overlay vertices must arise from discarded overlay edges. The completion phase of the algorithm returns the min-max angle triangulation of the non-triangular phases. It will suffice to show that there exists some triangulation that guarantees no large angles.

Observe that any input edge  $e$  crossing a triangle  $t$  of the overlay mesh can cause at most one edge of  $t$  to be discarded. This is because  $\theta_1 = \arcsin\left(\frac{1}{2\Gamma}\right)$  is



**Fig. 3.** **Left.** Circumball  $C_1$  and gap radius ball  $C_2$  for an edge  $e$  and its endpoint  $a$  respectively. The lines  $L_1$  and  $L_2$  are parallel. The point  $v$  must be in the shaded region. **Right.** The input edge causes exactly one edge of the overlay triangle to be discarded and replaced with a new edge whose angle with the discarded edge is less than  $\theta_1$ . The angles at  $v$  are changed by at most  $\theta_1$ .

a lower bound on the smallest angle in the overlay mesh, and thus the largest angle is at most  $\pi - 2\theta_1$ . It follows that at least one of these edges must intersect  $e$  at an angle greater than  $\theta_1$  and therefore will not be discarded. This fact implies that when an edge gets discarded, we can replace it with one that has been rotated by at most  $\theta_1$ . This is illustrated in Figure 3 (Right). The resulting largest angle at  $v$  is at most  $\pi - \theta_1$ . Now, if we look at the edges ordered radially around  $v$ , we see that no two adjacent edges can be rotated apart from each other. This is because if two edge of the triangle  $t$  are discarded then they both get replaced with edges that terminate on the third edge and thus lie entirely within the triangle. Thus the angle at  $v$  is strictly smaller in this case.

## 4 Size of the Triangulation

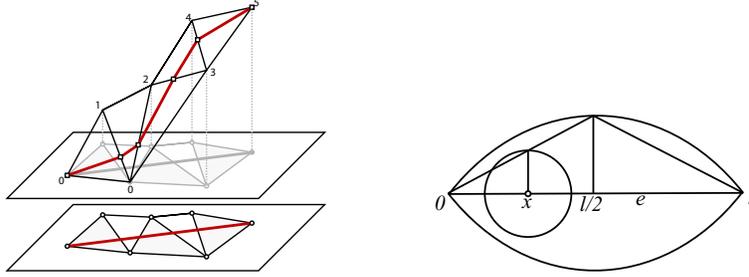
In this section we show that the size of the output mesh is determined only by the local feature size of the input PSLG.

**Lemma 3.** *For any input edge  $e$ , the number of triangles in the overlay mesh intersecting  $e$  is*

$$O\left(\int_{z \in e} \frac{1}{\text{fs}_0(z)} dz\right). \quad (1)$$

*Proof.* Let  $t_1, \dots, t_k$  be a minimal ordered sequence of adjacent triangles in the overlay mesh that covers  $e$ . Assign heights to the vertices of  $t_2, \dots, t_k$  so that the highest point of  $t_i$  is one more than the highest point of  $t_{i-1}$ . Set the heights for  $t_1$  to be 0, 0, 1. Let  $e_i = e \cap t_i$  be the subsegment of  $e$  contained in triangle  $t_i$ . We consider the lifted version of  $e$ , call it  $e^+$ , to be the polygonal chain in  $\mathbb{R}^3$  on the surface of these lifted triangles whose segments project down onto  $e$ .

Observe that the gradient of a lifted triangle  $t$  cannot be too steep because both the smallest angle of  $t$  and the maximum height difference between vertices of  $t$  are bounded by constants. The maximum difference between the height of the vertices of a triangle is bounded by the degree of vertices. Thus the gradient



**Fig. 4. Left.** An edge cuts through the overlay triangulation. Heights are assigned to the vertices so that the maximum height is exactly the number of triangles cut. **Right** An edge with an empty lens around it. The radius of the circle centered at  $x$  is the height of the triangle at  $x$ . This guarantees that the circle is contained entirely within the lens.

of  $t$  is bounded by  $\frac{\gamma}{r}$  where  $r$  is the radius of the circumcircle of  $t$  and  $\gamma > 0$  is a constant. Here, the constant  $\gamma$  depends only on the gap ratio  $\Gamma$  of the overlay mesh and the maximum degree of any vertex in the overlay mesh which is also a constant because the overlay mesh has small angle guarantees.

Let  $e_i$  be the subsegment of  $e$  lying in triangle  $t_i$ . The change in height along  $e_i^+$  is at most  $|e_i| \frac{\gamma}{r_i}$  where  $r_i$  is the circumradius of triangle  $t_i$ . Thus, the total change in height  $k$  along  $e^+$  is bounded as follows.

$$k \leq \sum_{i=1}^k |e_i| \frac{\gamma}{r_i} \leq \gamma \sum_{i=1}^k \int_{e_i} \frac{1}{r_i} dx \quad (2)$$

The triangles  $t_i$  are part of a well graded mesh and therefore,  $\text{lfs}_0(x) = cr_i$  for all  $x \in t_i$  and some constant  $c$ . So, we can rewrite the above inequality as follows to complete the proof.

$$k \leq \gamma \sum_{i=1}^k \int_{e_i} \frac{c}{\text{lfs}_0(x)} dx = c\gamma \int_e \frac{1}{\text{lfs}_0(x)} dx \quad (3)$$

**Theorem 3.** *The number of Steiner points added during the course of the algorithm is*

$$O\left(\int_{\Omega} \frac{1}{(\text{lfs}_0(z))^2} dz + \int_E \frac{1}{\text{lfs}_0(z)} dz\right). \quad (4)$$

where  $E$  is the set of input edges and  $\Omega$  is the input domain (i.e. the plane).

*Proof.* We look at the two phases of OSM where vertices are added. First, in the construction of the overlay mesh, the number of points added is  $O\left(\int_{\Omega} \frac{1}{(\text{lfs}_0(z))^2} dz\right)$ , as guaranteed by the SVR algorithm for point set meshing [19].

Second, in the stitching phase, we choose a subset of the intersections along each edge with the overlay mesh. It follows from Lemma 3 that the total number

of intersections is  $O\left(\int_E \frac{1}{\text{lfs}_0(z)} dz\right)$ . Therefore, it follows that the subset of these that we keep also achieves this bound. The statement of the theorem follows directly from summing the Steiner points added in each phase.

#### 4.1 Competitive results

The  $\alpha$ -lens is the main tool we will use to analyze the optimal mesh for a given maximum angle guarantee. Recall that an  $\alpha$ -lens on a line segment  $\overline{xy}$  is the set of all points  $z$  such that  $\angle xzy \geq \alpha$ . A lens is the intersection of two disks with the same radius. We note one important fact about  $\alpha$ -lenses: If  $\overline{ab}$  is a subsegment of  $\overline{xy}$  then the  $\alpha$ -lens around  $\overline{ab}$  is strictly contained in the  $\alpha$ -lens around  $\overline{xy}$ .

**Theorem 4.** *The output of OSM is a mesh with no large angles that is at most  $O(\log(L/s))$  times the size of any mesh achieving the same maximum angle guarantee for some fixed constant  $c$ .*

*Proof.* Just as in Theorem 3, we will consider the Steiner points added during the Overlay phase separate from those added during the Stitching phase. The Overlay phase only adds  $O(n \log(L/s))$  and any mesh conforming to the input has size  $\Omega(n)$  so we need only worry about Steiner points added during the Stitching phase.

Suppose we have an optimal size mesh  $M_{OPT}$  with the property that no angle is greater than some constant  $\alpha$ . For any edge  $e$  in  $M_{OPT}$ , the  $\alpha$ -lens around  $e$  contains no other vertices of  $M_{OPT}$ . This is just another way of stating the no large angle property. In particular, the edges  $e$  that are subsegments of input edges have  $\alpha$ -lenses that contain no input vertices.

In order to prove that OSM is  $\log(L/s)$ -competitive, it will suffice to prove that OSM stitches at most  $\log(L/s)$  Steiner points on any input subsegment of  $M_{OPT}$ . Recall that  $\text{lfs}_0(x)$  is the radius of the smallest circle centered at  $x$  containing two input vertices. In general,  $\text{lfs}_0$  is lower bounded by  $s$ , the distance between the two closest vertices in the mesh. We get a better lower bound on  $\text{lfs}_0$  near input edges from the fact that the lenses around the input subsegments of  $M_{OPT}$  contain no input vertices. We will use both of these lower bounds on  $\text{lfs}_0$  to upper bound the integral from Theorem 3. Recall that this integral bounds the number of stitch vertices added on any edge.

For a particular input subsegment  $e$  in  $M_{OPT}$ , we parameterize  $e$  on the interval  $[0, l]$  where  $l$  is the length of  $e$  as in Figure 4. To compute the lower bound on  $\text{lfs}_0(x)$  for  $x \in [0, l/2]$  it suffices to show that there is a circle centered at  $x$  that contains no input vertices. We first inscribe an isosceles triangle into the top half of the  $\alpha$ -lens around  $e$ . The altitude of the triangle at  $x$  is  $\tan^{-1}(\alpha/2)x$ .

Consider the circle  $C$  centered at  $x$  with radius  $\tan^{-1}(\alpha/2)x$ . Observe that the top edge of the inscribed triangle cuts off some half  $\alpha'$ -lens from the circle  $C$  and some half  $\alpha''$ -lens from the original lens around  $e$ . Observe that  $\alpha' = \alpha''$  and thus the smaller lens is entirely contained in the larger. It follows that the circle  $C$  is contained entirely within the  $\alpha$ -lens around  $e$  and thus,  $C$  contains no input vertices. Therefore,  $\text{lfs}_0(x) \geq \tan^{-1}(\alpha/2)x$  for all  $x \in e$ .

We can now use our bound on  $\text{lfs}_0$  to bound the size integral from Theorem 3 as follows.

$$\int_0^l \frac{1}{\text{lfs}_0(x)} dx \leq 2 \left( \int_0^s \frac{1}{s} dx + \int_s^{\frac{l}{2}} \frac{1}{\text{lfs}_0(x)} dx \right) \leq 2 + 2 \tan \alpha / 2 \int_s^{\frac{l}{2}} \frac{1}{x} dx \in O(\log \frac{L}{s})$$

## 5 Conclusions

### 5.1 Size bounds in terms of $n$

To better understand the size guarantees of OSM in relation to previous results that only analyze worst case performance, we can compute two coarse upper bounds on the mesh size.

First, we see that the mesh size is  $O(n^2 \log(L/s))$ . This follows from the fact that the overlay mesh is of size  $O(n \log(L/s))$ . There are at most  $O(n)$  edges so the stitching phase adds at most  $O(n^2 \log(L/s))$  points. When  $L/s \in O(\text{poly}(n))$ , this bound exactly matches the  $O(n^2 \log n)$  of Mitchell on triangulating with no large angles from [15].

Alternatively, one could recompute the sizing integral from Theorem 3 using  $s$  as a lower bound on  $\text{lfs}_0$ . The result is an  $O(n(L/s))$  upper bound on the mesh size. So, when  $L/s \in O(n)$ , the output size is  $O(n^2)$ , matching the bound of Tan [16]. Certain pathological examples such as Paterson’s example (see [13]) requiring  $\Omega(n^2)$  Steiner points can be drawn so that  $L/s \in O(n)$ . Thus, OSM is worst case optimal when the input vertices have linear spread.

For reasonable inputs where  $L/s \in O(\text{poly}(n))$ , Theorem 4 implies that the output of OSM is  $O(\log n)$ -competitive. For inputs that admit  $O(n)$  no-large-angle meshes, this is a factor of  $n/\log n$  better than the Tan guarantee [16].

### 5.2 Work Efficiency

Overlay Stitch Meshing can be easily implemented to run in time and space  $O(n \log(L/s) + m)$ . The runtime of the SVR algorithm used in the overlay phase was analyzed in [19]. This stage is the majority of the work. Simple arguments can show that the stitching and completion phases can be implemented as  $O(m)$  post-processes. In the case of input with  $O(\text{poly}(n))$  spread, this is asymptotically optimal work.

### 5.3 Extensions

The most obvious extension is into three and higher dimensions. Obtaining good guarantees on output size for meshing algorithms in three dimensions remains an interesting open problem [10, 11, 12]. For the general input case, unstructured meshing research has focused on concentric shelling techniques that have complicated implementations and have yielded no strong sizing results.

## References

- [1] Bern, M., Eppstein, D., Gilbert, J.: Provably good mesh generation. *J. Comput. Syst. Sci.* **48**(3) (1994) 384–409 [1](#), [3](#), [4](#)
- [2] Ruppert, J.: A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms* **18**(3) (1995) 548–585 [1](#), [3](#)
- [3] Mitchell, S., Vavasis, S.: Quality mesh generation in three dimensions. In: *Proc. 8th ACM Symp. Comp. Geom.* (1992) 212–221 [1](#)
- [4] Shewchuk, J.R.: Delaunay Refinement Algorithms for Triangular Mesh Generation. *Computational Geometry: Theory and Applications* **22**(1–3) (2002) 21–74 [2](#), [3](#)
- [5] Babuška, I., Aziz, A.K.: On the Angle Condition in the Finite Element Method. *SIAM Journal on Numerical Analysis* **13**(2) (1976) 214–226 [2](#)
- [6] Guattery, S., Miller, G.L., Walkington, N.: Estimating interpolation error: A combinatorial approach. In: *Tenth Annual ACM-SIAM Symposium on Discrete Algorithms*, Baltimore, ACM and SIAM (1999) 406–413 [2](#)
- [7] Boman, E., Hendrickson, B., Vavasis, S.: Solving elliptic finite element systems in near-linear time with support preconditioners (2004) [2](#)
- [8] Miller, G., Vavasis, S.: Only large angles mater. Private communications (2005) [2](#)
- [9] Pav, S.E.: Delaunay Refinement Algorithms. PhD thesis, Department of Mathematics, Carnegie Mellon University, Pittsburgh, Pennsylvania (2003) [3](#)
- [10] Cheng, S.W., Poon, S.H.: Graded Conforming Delaunay Tetrahedralization with Bounded Radius-Edge Ratio. In: *Proceedings of the Fourteenth Annual Symposium on Discrete Algorithms*, Baltimore, Maryland, Society for Industrial and Applied Mathematics (2003) 295–304 [3](#), [11](#)
- [11] Pav, S.E., Walkington, N.J.: Robust Three Dimensional Delaunay Refinement. In: *Thirteenth International Meshing Roundtable*, Williamsburg, Virginia, Sandia National Laboratories (2004) 145–156 [3](#), [11](#)
- [12] Cheng, S.W., Dey, T.K., Ramos, E.A., Ray, T.: Quality Meshing for Polyhedra with Small Angles. In: *Proceedings of the Twentieth Annual Symposium on Computational Geometry*, Brooklyn, New York, Association for Computing Machinery (2004) 290–299 [3](#), [11](#)
- [13] Bern, M.W., Dobkin, D.P., Eppstein, D.: Triangulating polygons without large angles. *Int. J. Comput. Geometry Appl* **5** (1995) 171–192 [3](#), [11](#)
- [14] Bern, M.W., Mitchell, S.A., Ruppert, J.: Linear-size nonobtuse triangulation of polygons. *Discrete & Computational Geometry* **14**(4) (1995) 411–428 [3](#)
- [15] Mitchell, S.A.: Refining a triangulation of a planar straight-line graph to eliminate large angles. In: *34th Annual Symposium on Foundations of Computer Science*, Palo Alto, California, IEEE (1993) 583–591 [3](#), [11](#)
- [16] Tan, T.S.: An optimal bound for high-quality conforming triangulations. *GEOMETRY: Discrete and Computational Geometry* **15** (1996) [3](#), [11](#)
- [17] Baker, B.S., Grosse, E., Rafferty, C.S.: Nonobtuse triangulation of polygons. *Discrete Comput. Geom.* **3**(2) (1988) 147–168 [4](#)
- [18] Dey, T.K.: Good triangulations in the plane. In: *CCCG: Canadian Conference in Computational Geometry*. (1990) [4](#)
- [19] Hudson, B., Miller, G., Phillips, T.: Sparse Voronoi Refinement. Technical Report CMU-CS-06-132, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania (2006) [4](#), [9](#), [11](#)