Where and How Chew's Second Delaunay Refinement Algorithm Works

Alexander Rand*

Abstract

Chew's second Delaunay refinement algorithm with off-center Steiner vertices leads to practical improvement over Ruppert's algorithm for quality mesh generation, but the most thorough theoretical analysis is known only for Ruppert's algorithm. A detailed analysis of Chew's second Delaunay refinement algorithm with off-centers is given, improving the guarantee of well-graded output for any minimum angle threshold $\alpha^* \leq 28.60°$.

1 Introduction

Ruppert's algorithm for quality triangular mesh generation [10] has a number of theoretical and practical advantages making it the prototypical Delaunay refinement setting: it is relatively simple to state, implement, and analyze. For non-acute input and a minimum angle threshold of about 20.70°, the algorithm is guaranteed to terminate and produce a mesh of optimal size up to a constant factor. Over the past 15 years, this elegant theory has been adjusted and refined to produce better and better meshes. From a theoretical standpoint, Miller, Pav, and Walkington gave an improved analysis of Ruppert's algorithm demonstrating that, under mild assumptions on the input, termination is guaranteed for a minimum angle threshold as high as 26.45° [7]. Off-center Steiner vertices provide an alternative to circumcenter insertion, reducing the mesh sizes produced in practice. Üngör introduced this concept and demonstrated its success with Ruppert's algorithm [13].

Chew's second Delaunay refinement algorithm [3] was originally studied for meshing surfaces embedded in 3D, but the restriction of this algorithm to the standard 2D mesh generation problem yields two specific advantages over Ruppert's algorithm: the algorithm is theoretically guaranteed to terminate for a larger minimum angle threshold (26.57°) and in practice the resulting meshes have fewer vertices [12]. Most of the improvements to Ruppert's algorithm have been applied to Chew's second Delaunay refinement algorithm and are similarly successful in practice; in fact, the default quality mesh generation algorithm in Triangle [11] is Chew's second Delaunay refinement algorithm with off-centers.

We improve the analysis of Chew's second Delaunay refinement algorithm with off-center vertices. By extending the Miller-Pav-Walkington analysis, we prove the termination of Chew's second Delaunay refinement algorithm for any minimum angle threshold less than 28.60°, and this guarantee holds not only for circumcenters but also for off-center Steiner vertices. Moreover, we generalize the Üngör off-center to a larger class of Steiner vertices characterized by a target angle and note that in some cases these vertices are outside existing selection discs. Finally, a simple example demonstrates the impact of the target angle parameter.

*Institute of Computational Engineering and Sciences, The University of Texas at Austin, arand@ices.utexas.edu

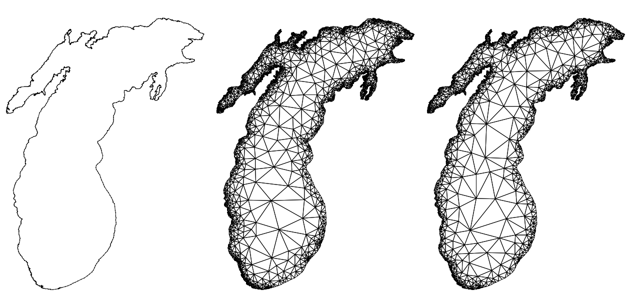

Figure 1: Using the boundary of Lake Michigan as input (left, 1537 vertices) and a minimum angle threshold of 25°, the results of Ruppert's algorithm (center, 3707 vertices) and Chew's second Delaunay refinement algorithm with off-centers (right, 2960 vertices) are shown.

2 Preliminaries

The input to a 2D mesh generator is a consistent collection of straight segments and vertices. The goal of the mesh generator is to add vertices so that a triangulation (in this paper, the constrained Delaunay triangulation) of the final vertex set both conforms to the input segments and contains only high quality triangles.

Formally we follow [7]: a **planar straight-line graph** (PSLG), $\mathcal{G} = (\mathcal{P}, \mathcal{S})$, is a pair of sets of vertices $\mathcal{P}$ and segments $\mathcal{S}$, such that the endpoints of each segment of $\mathcal{S}$ are contained in $\mathcal{P}$ and the intersection of any two segments of $\mathcal{S}$ is also contained in $\mathcal{P}$. A PSLG $\mathcal{G}' = (\mathcal{P}', \mathcal{S}')$ is a **refinement** of the PSLG $\mathcal{G}$ if $\mathcal{P} \subset \mathcal{P}'$ and each segment in $\mathcal{S}$ is the union of segments in $\mathcal{S}'$.

Problem Statement. Given an input PSLG $\mathcal{G}$ and a minimum angle threshold α^* compute a refinement $\mathcal{G}'$ such that all angles of all triangles of the constrained Delaunay triangulation of $\mathcal{G}'$ are larger than α^*.

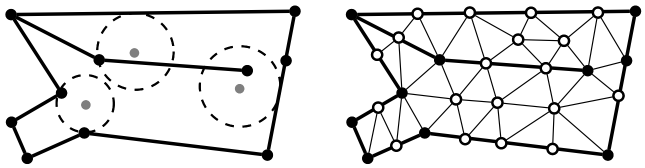

Figure 2: A PSLG (left) with local feature size indicated at several points (gray) and a refinement (right) of the PSLG that gives a quality, conforming triangulation.

The **local feature size** at point $\mathbf{x}$ with respect to PSLG $\mathcal{G}$, $\text{lfs}(\mathbf{x})$, is the radius of the smallest closed disk centered at $\mathbf{x}$ which intersects two *disjoint* features of $\mathcal{G}$. Most Delaunay refinement algorithm analysis is based on relating the mesh size to the local feature size of the input PSLG. Throughout this paper, local feature size is always considered with respect to the input PSLG. Moreover, local feature size is 1-Lipschitz: $\text{lfs}(\mathbf{x}) \leq \text{lfs}(\mathbf{y}) + |\mathbf{x} - \mathbf{y}|$.

Before stating and analyzing Chew's second Delaunay refinement algorithm, we state one fact about constrained Delaunay triangulations which satisfy an empty circumdisk property with respect to *visible* vertices; for a complete definition see [2].

Proposition 1 *Let $\mathcal{T}$ be a constrained Delaunay triangulation of PLSG $(\mathcal{P}, \mathcal{S})$. Suppose that triangle $T \in \mathcal{T}$ has circumcenter $\mathbf{c}$ and that $\mathbf{c}$ is not visible to T. Then T lies inside the diametral disk of the constrained segment $S \in \mathcal{S}$ nearest to T that prevents visibility.*

3 Chew's Second Delaunay Refinement Algorithm

Stated carefully as Algorithm 1, Chew's second Delaunay refinement algorithm has a few key differences from Ruppert's algorithm. The final constrained Delaunay triangulation is generated from three types of vertices, classified by why they were inserted into the mesh: input vertices, midpoints, and circumcenters.

Algorithm 1 Chew's second Delaunay refinement

Require: PSLG $\mathcal{G}$ and angle threshold α^*.
  Compute constrained Delaunay triangulation $\mathcal{T}$ of $\mathcal{G}$.
  while $\mathcal{T}$ contains a poor quality triangle T **do**
    if T encroaches a segments S **then**
      Remove circumcenters from diametral disk of S.
      Split S by adding its midpoint to $\mathcal{T}$.
    else
      Insert the circumcenter of T into $\mathcal{T}$.
    end if
  end while

Two particular steps above must be made precise.

Encroachment. A segment S is encroached if there is a poor quality triangle T in the current triangulation such that T and the circumcenter of T lie on opposite sides of S, and T is visible to S. Note the "converse": if T and its circumcenter lie on the opposite sides of S, then some segment (but possibly not S) is encroached.

Vertex Removal. When adding the midpoint $\mathbf{m}$ of a segment S, Chew's algorithm removes circumcenter vertices which lie in the diametral disk of S. In this treatment, we slightly relax this operation and fully specify a procedure for removing vertices. After inserting $\mathbf{m}$, the nearest visible neighbor to $\mathbf{m}$ is removed if it is a circumcenter, and this is repeated until the nearest visible neighbor is not a circumcenter. Some circumcenters may remain in the diametral disk of S.

The termination of Chew's second Delaunay refinement algorithm and good grading of the resulting mesh follow from a proof that no two vertices are placed too close together. The **insertion radius** $r_{\mathbf{q}}$ of vertex $\mathbf{q}$ is the distance from $\mathbf{q}$ to the nearest visible vertex in the mesh immediately following the insertion of $\mathbf{q}$. We call a mesh **well-graded** if there exists C depending only upon α^* such that for all vertices $\mathbf{q}$ inserted by the algorithm, $\text{lfs}(\mathbf{q}) \leq C r_{\mathbf{q}}$. This is a natural measure of success of a mesh generation algorithm: it guarantees termination and that the size of the triangles in the mesh are proportional to the underlying size of the input geometry. Proof that a mesh generation algorithm produces a well-graded mesh is usually performed via induction using an appropriate previously inserted vertex (called the parent vertex) on which to base the estimate.

The **parent** of a vertex $\mathbf{q}$, denoted $\mathbf{p}(\mathbf{q})$, is defined to be a specific vertex near $\mathbf{q}$ following insertion:

(1) If $\mathbf{q}$ is a circumcenter, then $\mathbf{p}(\mathbf{q})$ is the newest vertex on the shortest edge of triangle T of which $\mathbf{q}$ is the circumcenter.

(2) If $\mathbf{q}$ is a midpoint and the nearest visible neighbor to $\mathbf{q}$ is not contained in the input segment containing $\mathbf{q}$, then $\mathbf{p}(\mathbf{q})$ is this nearest visible neighbor.

(3) If $\mathbf{q}$ is a midpoint and after deletion of some vertices no vertices remain in the diametral disk of S, let $\mathcal{P}_r$ be the set containing all removed circumcenters. If either endpoint of S is newer than than any vertex in $\mathcal{P}_r$, the most recently inserted endpoint of S is the $\mathbf{p}(\mathbf{q})$. Otherwise, $\mathbf{p}(\mathbf{q})$ is the vertex in $\mathcal{P}_r$ with the smallest insertion radius.

Define $\mathbf{p}_2(\mathbf{q}) := \mathbf{p}(\mathbf{p}(\mathbf{q}))$, $\mathbf{p}_3(\mathbf{q}) := \mathbf{p}(\mathbf{p}(\mathbf{p}(\mathbf{q})))$, etc. Next we prove Chew's second Delaunay refinement algorithm succeeds for non-acute input.

Theorem 2 ([12]) *For $\alpha^* < \tan^{-1}(1/2) \approx 26.6°$ and non-acute input, Chew's second Delaunay refinement algorithm terminates producing a well-graded, quality mesh.*

This proof follows the argument in [12] using the slightly relaxed vertex removal procedure mentioned

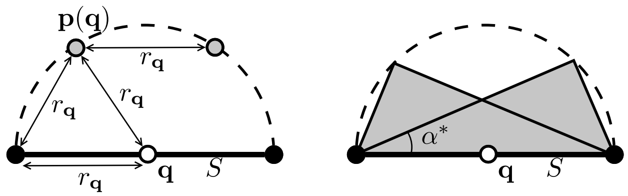

Figure 3: Subcases 3b (left) and 3c (right) in Theorem 2.

previously. The cases are carefully enumerated so the proof can be augmented in later sections to provide an improved analysis and accept variants of the algorithm.

Proof. To prove that the resulting mesh is well-graded, we inductively find two constants $0 < C_c < C_m < \infty$ such that $\mathrm{lfs}(\mathbf{q}) < C_c r_{\mathbf{q}}$ for any circumcenter and $\mathrm{lfs}(\mathbf{q}) < C_c r_{\mathbf{q}}$ for any midpoint. We consider three cases corresponding to the definition of the parent vertex.
Case 1. $\mathbf{q}$ is a circumcenter. Then,

$$\begin{aligned}\mathrm{lfs}(\mathbf{q}) &\leq |\mathbf{q} - \mathbf{p}(\mathbf{q})| + \mathrm{lfs}(\mathbf{p}(\mathbf{q})) \leq r_{\mathbf{q}} + C_m r_{\mathbf{p}(\mathbf{q})} \\ &\leq (1 + 2C_m \sin\alpha)\, r_{\mathbf{q}}.\end{aligned} \tag{1}$$

Case 2. $\mathbf{q}$ is a midpoint and a vertex other than $\mathbf{q}$ remains in the diametral disk of the segment which was split. Then $\mathbf{p}(\mathbf{q})$ must be an input vertex or midpoint. Then since the input is non-acute, this vertex belongs to an input feature which is disjoint from the input segment containing $\mathbf{q}$ and thus

$$\mathrm{lfs}(\mathbf{q}) \leq |\mathbf{q} - \mathbf{p}(\mathbf{q})| = r_{\mathbf{q}}. \tag{2}$$

Case 3. $\mathbf{q}$ is a midpoint and the diametral disk of the newly split segment is empty (other than $\mathbf{q}$). Recalling Proposition 1, all the vertices of the encroaching triangle must lie inside the diametral disk of the segment containing $\mathbf{q}$.
Subcase 3a. $\mathbf{p}(\mathbf{q})$ is a midpoint. The assumption of non-acute input and the parent vertex definition imply that $\mathbf{p}(\mathbf{q})$ is an endpoint of the segment S. Recalling Proposition 1, let $\mathbf{c}$ be a circumcenter that is older than $\mathbf{p}(\mathbf{q})$ and was removed from the diametral disk of S. Since $\mathbf{c}$ was not removed when $\mathbf{p}(\mathbf{q})$ was inserted, the diametral disk of $\mathbf{p}(\mathbf{q})$ was not completely emptied and thus Case 2 applies to $\mathbf{p}(\mathbf{q})$. So $\mathrm{lfs}(\mathbf{p}(\mathbf{q})) \leq r_{\mathbf{p}(\mathbf{q})} \leq |\mathbf{p}(\mathbf{q}) - \mathbf{c}|$, and thus

$$\mathrm{lfs}(\mathbf{q}) \leq |\mathbf{q} - \mathbf{p}(\mathbf{q})| + \mathrm{lfs}(\mathbf{p}(\mathbf{q})) \leq r_{\mathbf{q}} + r_{\mathbf{p}(\mathbf{q})} \leq 3r_{\mathbf{q}}. \tag{3}$$

Subcase 3b. $\mathbf{p}(\mathbf{q})$ is a circumcenter and at least two circumcenters were removed from the half of the diametral disk of S which is visible to $\mathbf{p}(\mathbf{q})$. Since all of these circumcenters were inserted after the endpoints of S (by the definition of the parent vertex), one of these vertices must have an insertion radius no larger than $r_{\mathbf{q}}$; see Figure 3(left). Then,

$$\mathrm{lfs}(\mathbf{q}) \leq |\mathbf{q} - \mathbf{p}(\mathbf{q})| + \mathrm{lfs}(\mathbf{p}(\mathbf{q})) \leq (1 + C_c) r_{\mathbf{q}}. \tag{4}$$

Subcase 3c. $\mathbf{p}(\mathbf{q})$ is a circumcenter and $\mathbf{p}(\mathbf{q})$ was the only circumcenter removed from the half of the diametral disk of S visible to $\mathbf{p}(\mathbf{q})$. Then to form a skinny triangle with circumcenter on the opposite side of S, $\mathbf{p}(\mathbf{q})$ must belong to the shaded area in Figure 3(right). Then $r_{\mathbf{p}(\mathbf{q})} \leq r_{\mathbf{q}}/\cos\alpha^*$ and thus,

$$\begin{aligned}\mathrm{lfs}(\mathbf{q}) &\leq |\mathbf{q} - \mathbf{p}(\mathbf{q})| + \mathrm{lfs}(\mathbf{p}(\mathbf{q})) \leq r_{\mathbf{q}} + C_c r_{\mathbf{p}(\mathbf{q})} \\ &\leq \left(1 + \frac{C_c}{\cos\alpha}\right) r_{\mathbf{q}}.\end{aligned} \tag{5}$$

The requirements from the various cases (1)-(5) can be summarized by three conditions: $C_c \geq 1 + 2C_m \sin\alpha$, $C_m \geq 3$, and $C_m \geq 1 + \frac{C_c}{\cos\alpha}$. Suitable constants exist only if $\tan\alpha^* < 1/2$. □

4 Off-Centers

Off-center Steiner vertices were developed as an alternative to circumcenter insertion to reduce the number of vertices inserted by Delaunay refinement algorithms [13]. We use the term off-center (or Υ-off-center to identify the parameter described below) to refer to the special class of Steiner points described by Üngör as opposed to the more general selection disks [1, 5] or selection regions [4, 6] in the literature.

If triangle T has a smallest angle less than $\alpha^*/2$, then inserting its circumcenter is guaranteed to create another poor-quality triangle since the newly inserted circumcenter and the shortest edge of T form a poor quality triangle. Üngör recognized that by selecting an alternative Steiner point, the mesh generator can control the quality of this particular newly formed triangle and, in practice, produce a smaller mesh.

First, we define the class of Υ-off-centers and remark how they generalize Üngör's definition. Let T be a poor quality constrained Delaunay triangle (i.e., the smallest angle of T, denoted α_T, is less than α^*), let $\overline{\mathbf{q}_1\mathbf{q}_2}$ be the shortest edge of T, and let $\mathbf{c}$ denote the circumcenter of T. The Υ-off-center $\mathbf{c}'$ is an attempt to create a new triangle with smallest angle Υ. If $\mathbf{q}_1$ and $\mathbf{q}_2$ are the endpoints of the shortest edge of triangle T, the Υ-off-center $\mathbf{c}'$ is defined as the unique point such that (i) $|\mathbf{q}_1 - \mathbf{c}'| = |\mathbf{q}_2 - \mathbf{c}'|$, (ii) $\angle\mathbf{q}_1\mathbf{c}\mathbf{q}_2 = \Upsilon$, and (iii) $(\mathbf{c} - \mathbf{q}_1) \cdot (\mathbf{c}' - \mathbf{q}_1) > 0$. See Figure 4 for a depiction of the Υ-off-center region.

Üngör's original work suggested using $\Upsilon_T := \max(2\alpha_T, \alpha^*)$ which separates the points as much as possible without creating a poor quality triangle between the new off-center and the shortest edge of the split triangle. In this setting, the algorithm the was shown to terminate and produce a well-graded mesh.

Theorem 3 (Ungor [13]) *Let minimum angle parameter $\alpha^* < \arcsin(1/(2\sqrt{2}))$ be given. Then Ruppert's algorithm with Υ-off-centers and $\Upsilon_T = \max(2\alpha_T, \alpha^*)$ terminates producing a well-graded, quality mesh.*

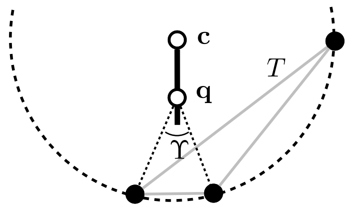

Figure 4: For a poor quality triangle T, the set of admissible Υ-off-centers is shown with the triangle circumcenter $\mathbf{c}$ and a typical Υ-off-center $\mathbf{q}$.

In Triangle [11], slightly larger values of Υ_T (about 5%) are used. In practice this makes "bunches" of nearly minimal quality triangles likely to appear near input edges and yields a mesh with fewer vertices. The proof of Theorem 3 can be extended to admit any $\Upsilon_T \in [2\alpha_T, \alpha^*]$ and (recalling Proposition 1) Chew's second Delaunay refinement algorithm. We provide a more detailed analysis which admits larger values of Υ_T.

Theorem 4 *If* $\alpha^* < \tan^{-1}(1/2)$ *and* $\Upsilon_T \in [2\alpha_T, 2\sin^{-1}(\cos(\alpha^*/2)))$*, Chew's second Delaunay refinement algorithm with* Υ*-off-centers terminates producing a well-graded, quality mesh.*

Proof. We will verify that the general structure of the proof of Chew's algorithm still applies, albeit with a few additional cases. Estimates on the insertion radii of Υ-off-centers must be revisited.
Case 1. Let $\mathbf{q}$ denote an Υ-off-center associated with poor quality triangle T with shortest edge $\overline{\mathbf{v}_1\mathbf{v}_2}$ and $\mathbf{v}_1$ is more recently inserted than $\mathbf{v}_2$. Since the nearest vertex to $\mathbf{q}$ may not be a vertex of T, we must deal with two subcases. In one of these subcases, the parent vertex of $\mathbf{q}$ will be redefined.
Subcase 1a. $\mathbf{v}_1$ is the nearest vertex to $\mathbf{q}$. Then,

$$\begin{aligned}\mathrm{lfs}(\mathbf{q}) &\le |\mathbf{q}-\mathbf{v}_1| + \mathrm{lfs}(v_1) \le r_{\mathbf{q}} + C_m r_{\mathbf{v}_1} \\ &\le \left(1 + 2C_m \sin\frac{\Upsilon_T}{2}\right) r_{\mathbf{q}}. \end{aligned} \tag{6}$$

Subcase 1b. $\mathbf{u}_1 \neq \mathbf{v}_1$ is the nearest vertex to $\mathbf{q}$. The edge $\overline{\mathbf{q}\mathbf{u}_1}$ is shared by two new Delaunay triangles and let $\mathbf{u}_2$ denote the additional vertex of one of these triangles that is nearest to $\mathbf{u}_1$. Since $\mathbf{q}$ must be a Delaunay neighbor to $\mathbf{v}_1$ and $\mathbf{v}_2$, $\mathbf{u}_1$ and $\mathbf{u}_2$ must both live in the (closed) diametral disk of $\overline{\mathbf{v}_1\mathbf{v}_2}$, and thus $|\mathbf{u}_1 - \mathbf{u}_2| \le |\mathbf{v}_1 - \mathbf{v}_2|/\sqrt{2}$. Define the parent of $\mathbf{c}'$ to be the newest vertex in $\{\mathbf{u}_1, \mathbf{u}_2\}$. Then

$$\begin{aligned}\mathrm{lfs}(\mathbf{q}) &\le |\mathbf{q}-\mathbf{p}(\mathbf{q})| + \mathrm{lfs}(\mathbf{p}(\mathbf{q})) \le r_{\mathbf{q}} + C_m r_{\mathbf{p}(\mathbf{q})} \\ &\le \left(1 + \sqrt{2}C_m \sin\frac{\Upsilon_T}{2}\right) r_{\mathbf{q}}. \end{aligned} \tag{7}$$

Cases 2 and 3 of Theorem 2 are identical in the Υ-off-center algorithm. Now the worst case involves simultaneously satisfying Subcases 1a and 3c:

$$C_c \ge 1 + 2C_m \sin\frac{\Upsilon_T}{2}; \qquad C_m \ge 1 + \frac{C_c}{\cos\alpha^*}.$$

If $\Upsilon_T < 2\sin^{-1}(\cos(\alpha^*/2))$, C_c and C_m exist. □

Observation 1 *The region of admissible* Υ*-off-centers is not a subset of the selection disks in [1, 5]: the larger values of* Υ_T *lie outside the standard disk.*

5 The Three Circumcenter Lemma

The critical cases in the proofs of Theorems 2 and 4 occur when a segment midpoint is inserted following encroachment due to a circumcenter. Circumcenters always have larger insertion radii than their parent vertices, while midpoints can have slightly smaller radii. The improved analysis of Ruppert's algorithm by Miller, Pav, and Walkington [7] demonstrated that several circumcenters must lie between certain midpoints in a sequence of parent vertices and thus insertion radii gains from the extra circumcenters can be used to offset the insertion radii reduction of the final midpoint. The result improved the admissible minimum angle threshold of Ruppert's algorithm from 20.70° to 26.45°.

Let $\mathbf{q}$ be a midpoint inserted by a Delaunay refinement algorithm. The **circumcenter** (or **Υ-off-center**) **sequence** associated with $\mathbf{q}$ is the sequence of points $\{\mathbf{p}_i(\mathbf{q})\}_{i=0}^{n}$, where n is the smallest positive index such that $\mathbf{p}_n(\mathbf{q})$ lies on a feature of the input PSLG. $\mathbf{q} = \mathbf{p}_0(\mathbf{q})$ is called the final vertex in the sequence and $\mathbf{p}_n(\mathbf{q})$ is called the initial vertex in the sequence. The crux of the Miller-Pav-Walkington analysis relies on studying circumcenter sequences that begin and end on the same input segment.

Lemma 5 (Miller-Pav-Walkington [7]) *If a circumcenter sequence both (i) begins and ends on the same input segment and (ii) the insertion radius of the final vertex is no larger than that of the initial vertex, then the sequence contains at least three circumcenters.*

The only property of circumcenters that is used in the proof of Lemma 5 is that circumcenters lie on the boundary of the Voronoi cell of their parent vertex. Thus the lemma can be extended to Υ-off-centers as stated below. For technical reasons to be made clear in the upcoming proof define $A(\alpha) := 2\sin^{-1}((\cos(\alpha^*/2))^{1/3})$.

Corollary 6 *Let* $\Upsilon_T \in [2\alpha_T, A(\alpha_T))$*. If a* Υ*-off-center sequence (i) begins and ends on the same input segment, (ii) the insertion radius of the final vertex is no larger than that of the initial vertex, and (iii) contains only vertices handled by Theorem 4 Subcase 1a, then the sequence contains at least three* Υ*-off-centers.*

This section closes with a related technical lemma.

Lemma 7 *Let* $\Upsilon_T \in [2\alpha_T, A(\alpha_T))$ *and let* $\{\mathbf{p}_i(\mathbf{q})\}_{i=0}^{n}$ *be an* Υ*-off-center sequence. Then there exists* C_d *such that* $|\mathbf{q} - \mathbf{p}_n(\mathbf{q})| \le C_d r_{\mathbf{q}}$.

6 Restricted Input Class

The core of the argument is given by restricting attention to PSLGs with no adjacent input segments.

Theorem 8 *Suppose no segments in the input PSLG are adjacent. Let $\alpha \leq 28.60°$ and select Υ-off-centers such that $\Upsilon_T \in [2\alpha_T, A(\alpha_T))$. Chew's second Delaunay refinement algorithm terminates producing a well-graded, quality mesh.*

Proof. The proof involves considering the interaction between cases in Theorem 4. The estimates for (sub)cases 1a, 1b, 2, 3a, and 3b are used without any changes. Let $\mathbf{q}$ be a "subcase 3c"-vertex and let $\mathbf{P} = \{\mathbf{p}_i(\mathbf{q})\}_{i=0}^{n}$ be the associated Υ-off-center sequence.

Case A: There is at least one vertex $\mathbf{p}_j(\mathbf{q}) \in \mathbf{P}$ that is a "subcase 1b"-vertex. Since $\alpha^* < 30°$, $r_{\mathbf{p}_1(\mathbf{q})} > r_{\mathbf{p}_2(\mathbf{q})} > \ldots > r_{\mathbf{p}_n(\mathbf{q})}$. Then (applying Lemma 7),

$$\begin{aligned}\mathrm{lfs}(\mathbf{q}) &\leq |\mathbf{q} - \mathbf{p}_j(\mathbf{q})| + \mathrm{lfs}(\mathbf{p}_j(\mathbf{q})) \\ &\leq C_d r_{\mathbf{q}} + \left(1 + C_m\sqrt{2}\sin\frac{\Upsilon_T}{2}\right) r_{\mathbf{p}_{j+1}(\mathbf{q})} \\ &\leq \left(C_d + \left(1 + C_m\sqrt{2}\sin\frac{\Upsilon_T}{2}\right)\frac{1}{\cos\alpha^*}\right) r_{\mathbf{q}}. \quad (8)\end{aligned}$$

Case B: $\mathbf{P}$ contains no "subcase 1b" vertices and $r_{\mathbf{q}} > r_{\mathbf{p}_n(\mathbf{q})}$. Since all subsegments are derived by midpoint splits from an original segment, $r_{\mathbf{q}} \geq 2r_{\mathbf{p}_n(\mathbf{q})}$ and

$$\mathrm{lfs}(\mathbf{q}) \leq |\mathbf{q} - \mathbf{p}_n(\mathbf{q})| + \mathrm{lfs}(\mathbf{p}_n(\mathbf{q})) \leq (C_d + C_m/2) r_{\mathbf{q}}. \quad (9)$$

Case C: $\mathbf{P}$ contains no "subcase 1b" vertices and $r_{\mathbf{q}} \leq r_{\mathbf{p}_n(\mathbf{q})}$. By Corollary 6, $n \geq 4$. Υ-off-centers are constructed such that $2\sin(\Upsilon_T/2) r_{\mathbf{p}_i(\mathbf{q})} > r_{\mathbf{p}_{i+1}(\mathbf{q})}$ for $i \in \{1, 2, 3\}$. Then

$$\begin{aligned}\mathrm{lfs}(\mathbf{q}) &\leq |\mathbf{q} - \mathbf{p}_4(\mathbf{q})| + \mathrm{lfs}(\mathbf{p}_4(\mathbf{q})) \\ &\leq C_d r_{\mathbf{q}} + C_m 8\sin^3(\Upsilon_T/2)\, r_{\mathbf{p}_1(\mathbf{q})} \\ &\leq \left(C_d + C_m\frac{8\sin^3(\Upsilon_T/2)}{\cos\alpha^*}\right) r_{\mathbf{q}}. \quad (10)\end{aligned}$$

Requirement (10) is stronger than (8) and (9) so we focus our attention there. $A(\alpha)$ has been defined so that $8\sin^3(\Upsilon_T/2)/\cos\alpha^* < 1$ and thus a suitable constant C_m exists. The interval $[2\alpha_T, A(\alpha_T))$ is nonempty exactly when $8\sin^3\alpha^*/\cos\alpha^* < 1$ which is equivalent to our assumption $\alpha \leq 28.60°$. □

7 General Input

Acute angles between input segments pose a fundamental problem in Delaunay refinement and any application of the three circumcenter lemma requires some restrictions on the allowable adjacent input segments [7]. Perhaps the simplest protection strategy is to split adjacent segments at equal lengths proportional to the local feature size and disallow the resulting adjacent subsegments to be split by the algorithm; a description of this "collar" protection strategy can be found in [9]. The advantage of this approach is that following initial grooming there are no adjacent input segments that can be refined which ensures the analysis of Theorem 8 holds.

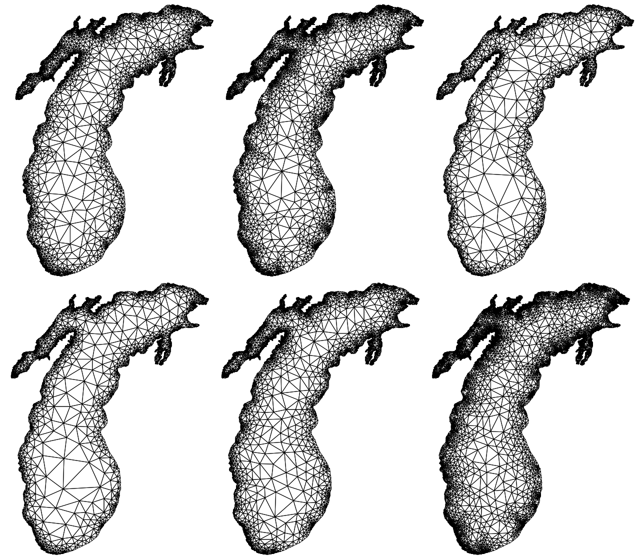

Figure 5: Meshes produced using Υ-off-centers, $\alpha^* = 28°$. (top left)$\Upsilon = 0$ (i.e., circumcenter insertion) gives 4975 vertices. (top center) $\Upsilon = 27 \Rightarrow 6475$ vertices. (top right) $\Upsilon = 29 \Rightarrow 3432$ vertices. (bottom left) $\Upsilon = 40 \Rightarrow 3955$ vertices. (bottom center) $\Upsilon = 50 \Rightarrow 5346$ vertices. (bottom right) $\Upsilon = 55 \Rightarrow 8617$ vertices.

Corollary 9 *Let $\alpha \leq 28.60°$ and select Υ-off-centers such that $\Upsilon_T \in [2\alpha_T, A(\alpha_T))$. Chew's second Delaunay refinement algorithm with "collar" vertex protection terminates producing a well-graded, quality mesh away from small input angles.*

Another strategy for protecting small input angles (which we call the "wedge" method) disallows the refinement of poor quality triangles which lie between adjacent input segments [7]. This approach is especially important because no large angles are created even in the presence of very small input angles. The complete analysis of this scheme is rather involved and only appears in [8], but the crux of the analysis is the three-circumcenter lemma. Thus we claim that this algorithm also succeeds in creating a well-graded mesh.

Claim Let $\alpha \leq 28.60°$ and select Υ-off-centers such that $\Upsilon_T \in [2\alpha_T, A(\alpha_T))$. Chew's second Delaunay refinement algorithm with "wedge" vertex protection terminates producing a well-graded, quality mesh away from small input angles.

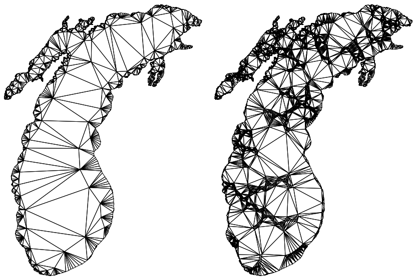

Figure 6: Meshes produced using $\alpha^* = 5°$ with $\Upsilon = 6°$ (left, 1660 vertices) and $\Upsilon = 59.5°$ (right, 6072 vertices).

8 Example

Using a 1537 vertex boundary of Lake Michigan as input, we give examples demonstrating the impact that Υ-off-centers have on the meshes generated. To denote a fixed target angle $\Upsilon = \gamma$ is used as a shorthand for the strategy $\Upsilon_T = \max(2\alpha_T, \gamma)$. Figures 5 and 6 contain meshes generated for the Lake Michigan example using various values of Υ and α^*. Figure 7 contains histograms of the smallest angles of all the triangles in meshes resulting from different Υ values and Figure 8 plots the number of mesh vertices as a function of Υ.

Acknowledgment

The author acknowledges useful discussions with Noel Walkington and Todd Phillips. All example meshes were generated using Jonathan Shewchuk's Triangle program (modified slightly).

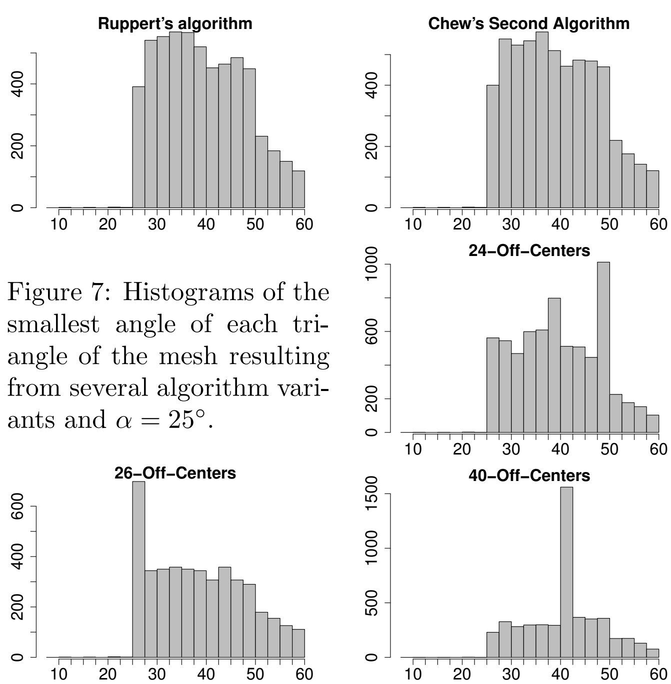

Figure 7: Histograms of the smallest angle of each triangle of the mesh resulting from several algorithm variants and $\alpha = 25°$.

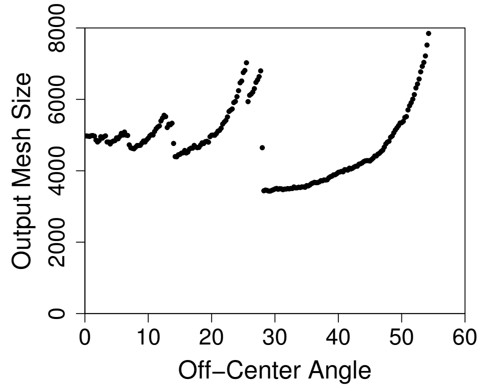

Figure 8: The number of vertices in the resulting mesh using Υ-off-centers for various values of Υ and $\alpha = 25°$.

References

[1] A. Chernikov and N. Chrisochoides. Generalized two-dimensional Delaunay mesh refinement. *SIAM J. Sci. Comput.*, 31(5):3387–3403, 2009.

[2] L. P. Chew. Constrained Delaunay triangulations. In *Proc. 3rd Symp. Comput. Geom.*, pages 215–222, 1987.

[3] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. In *Proc. 9th Symp. Comput. Geom.*, pages 274–280, 1993.

[4] H. Erten and A. Üngör. Quality triangulations with locally optimal Steiner points. *SIAM J. Sci. Comput.*, 31:2103–2130, 2009.

[5] P. Foteinos, A. Chernikov, and N. Chrisochoides. Fully generalized two-dimensional constrained Delaunay mesh refinement. *SIAM J. Sci. Comput.*, 32(5):2659–2686, 2010.

[6] B. Hudson. Safe Steiner points for Delaunay refinement. In *Res. Notes 17th Int. Meshing Roundtable*, 2008.

[7] G. L. Miller, S. E. Pav, and N. Walkington. When and why Delaunay refinement algorithms work. *Int. J. Comput. Geom. Appl.*, 15(1):25–54, 2005.

[8] S. E. Pav. *Delaunay Refinement Algorithms.* PhD thesis, Carnegie Mellon University, May 2003.

[9] A. Rand and N. Walkington. Collars and intestines: Practical conforming Delaunay refinement. *Proc. 18th Int. Meshing Roundtable*, pages 481–497, 2009.

[10] J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms*, 18(3):548–585, 1995.

[11] J. R. Shewchuk. Triangle: A 2D quality mesh generator and Delaunay triangulator. `http://www.cs.cmu.edu/~quake/triangle.html`.

[12] J. R. Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Comput. Geom. Theory Appl.*, 22(1–3):86–95, 2002.

[13] A. Üngör. Off-centers: A new type of Steiner points for computing size-optimal quality-guaranteed Delaunay triangulations. In *Proc. 6th Latin Amer. Symp. Theor. Inform.*, pages 152–161, 2004.