On the Termination of Ruppert's Algorithm

Alexander Rand

University of Texas-Austin arand@ices.utexas.edu

For a non-acute planar straight-line graph, Ruppert's algorithm produces a conforming Delaunay triangulation composed of triangles containing no angles less than α. Ruppert proved the algorithm terminates for all $\alpha \lessapprox 20.7°$ [5], but this constraint has been seen to be overly conservative in practice. Ruppert observed that the minimum angle reaches 30° during typical runs of the algorithm. Further experimentation by Shewchuk [6] suggested that even higher values are admissible: "In practice, the algorithm generally halts with an angle constraint of 33.8°, but often fails [at] 33.9°."

The constraint on α can be improved to about 26.5° by splitting adjacent segments at equal length [3] or modifying the algorithm [1, 7]. Certain modifications of the vertex insertion procedure have also been demonstrated that this constraint can be improved to possibly 40° or more [2]. Many other variants of Ruppert's algorithm have been designed to relax or eliminate the requirement that the input be non-acute; see [4] for a comprehensive investigation. However, no improvement has been proved for the minimum angle constraint for Ruppert's original algorithm.

For each variant of the algorithm, theoretical requirement falls short of the observed behavior. We aim to study this gap for the original algorithm with the hope that a precise analysis of its behavior will lead to a better understanding of the subsequent improvements. There are three parts of this investigation. First, we give a counterexample demonstrating that Ruppert's algorithm can fail to terminate for $\alpha \approx 30.7°$. Second, several experiments of Ruppert's algorithm using random inputs are considered. Finally, we prove that Ruppert's algorithm terminates for all $\alpha \lessapprox 22.2°$.

Ruppert's Algorithm

A brief description of Ruppert's algorithm, the prototypical Delaunay refinement algorithm, is given in Table 1. Vertices inserted by Operations 1 and 2 are called *midpoints* and *circumcenters*, respectively. The *insertion radius* of a vertex $\mathbf{q}$, denoted $r_{\mathbf{q}}$, is the distance from the vertex to its nearest neighbor immediately after its insertion into the Delaunay triangulation.

Table 1. Ruppert's algorithm repeats Operations 1 and 2 until no encroached segments or Delaunay triangles with small angles exist.

Input	A non-acute planar straight-line graph (PSLG), $\mathcal{C} = (\mathcal{P}, \mathcal{S})$, consisting of input vertices $\mathcal{P}$ and input segments $\mathcal{S}$.
Parameter	A minimum output angle threshold α.
Preprocess	Compute the Delaunay triangulation of $\mathcal{P}$.
Operation 1	Insert the midpoint of a segment with non-empty diametral ball.
Operation 2	Insert the circumcenter of a Delaunay triangle with angle smaller than α, unless this point lies in the diametral ball of a segment. Otherwise, perform Operation 1 on the segment.
Output	Set of vertices $\mathcal{P}'$ with Delaunay triangulation $\mathcal{T}'$ which conforms to $\mathcal{C}$ and contains no triangles with angle smaller than α.

Each vertex $\mathbf{q}$ is associated with a parent vertex $\mathbf{p}(\mathbf{q})$. The parent of a midpoint $\mathbf{q}$ is the vertex encroaching the segment causing $\mathbf{q}$ to be inserted. The parent of a circumcenter $\mathbf{q}$ is the vertex which is the newer endpoint of the shortest edge of the poor-quality triangle with circumcenter $\mathbf{q}$. Define $\mathbf{p}_2(\mathbf{q}) := \mathbf{p}(\mathbf{p}(\mathbf{q}))$ and inductively $\mathbf{p}_k(\mathbf{q}) := \mathbf{p}(\mathbf{p}_{k-1}(\mathbf{q}))$.

The *local feature size* of a point $\mathbf{x}$, denoted $\mathrm{lfs}(\mathbf{x})$, is the radius of the smallest closed ball centered at $\mathbf{x}$ which intersects two disjoint features (segments or vertices) in the input. Local feature size is a 1-Lipschitz function; i.e., $\mathrm{lfs}(\mathbf{x}) \leq |\mathbf{x} - \mathbf{y}| + \mathrm{lfs}(\mathbf{y})$. Define

$$\alpha^* := \sup \{\alpha \mid \text{there exists a constant } C_\alpha \text{ such that } \mathrm{lfs}(\mathbf{q}) \leq C_\alpha r_{\mathbf{q}}\}.$$

Ruppert's original analysis is summarized in the following theorem.

Theorem 1 (Ruppert, [5]). $\alpha^* \geq \arcsin 2^{-3/2} \approx 20.7°$.

Counterexample

Consider a non-acute input containing four adjacent segments of lengths 2, $2^{3/4}$, $2^{1/2}$ and $2^{1/4}$ as in Fig. 1. The endpoints of the longest and shortest segments form a Delaunay triangle with smallest angle $\arctan 2^{-3/4}$. The circumcenter of this triangle encroaches upon the longer segment causing the midpoint of the longest segment to be inserted. Now the adjacent segments have lengths 1, $2^{3/4}$, $2^{1/2}$ and $2^{1/4}$ and the ratio of the shortest and longest segment is still $2^{3/4}$. Again this gives a poor quality triangle and the midpoint of the longest segment is inserted. This cycle repeats indefinitely. The result is an upper bound on α^*.

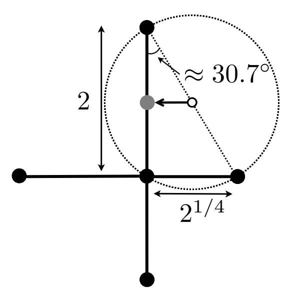

Fig. 1.

Theorem 2. $\alpha^* \leq \arctan 2^{-3/4} \approx 30.7°$.

Experiments

Next we consider two experiments which investigate the claim that Ruppert's algorithm terminates in practice for $\alpha \geq 33°$. First, Ruppert's algorithm is run 500 times on inputs containing $10,000$ randomly generated vertices. The success rate for several different values of α is given in Table 2(a): the algorithm always succeeded for $\alpha = 33.3°$ and always failed for $\alpha = 33.5°$. Fig. 2 contains a histogram of the maximum ratio of the local feature size to output mesh size (which is closely related to the constant in Theorem 1) of each run for several alpha values. The second experiment is inspired by the example in Fig. 1. Input are created consisting of four adjacent segments of random length and 100 additional random vertices. After testing 500 example inputs, Table 2(b) contains the success rate of the algorithm.

Table 2.

(a)

Minimum Angle	Success Rate
33.3	100.0%
33.4	98.6%
33.5	0.0%

(b)

Minimum Angle	Success Rate
31.4	100.0%
32.0	99.6%
32.6	98.6%

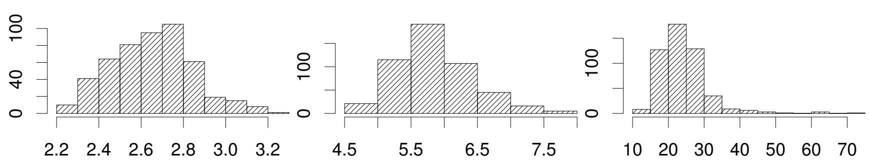

(a) $\alpha = 20.7°$ (b) $\alpha = 30.0°$ (c) $\alpha = 33.0°$

Fig. 2. Histogram of the ratio of output mesh size to local feature size for 500 runs of Ruppert's algorithm with 10,000 randomly placed input vertices.

Theorem

By carefully analyzing the example in Fig. 1, we strengthen Theorem 1.

Theorem 3. $\alpha^* \gtrapprox 22.2°$.

Proof. If $\mathbf{q}$ is a midpoint, $\mathbf{p}_{k+1}(\mathbf{q})$ is either a midpoint or input vertex, and $\{\mathbf{p}_i(\mathbf{q})\}_{i=1}^{k}$ are all circumcenters, say that $\mathbf{q}$ results from a *length-k circumcenter sequence*. The circumcenter sequence will be said to begin at $\mathbf{p}_{k+1}(\mathbf{q})$ and end at $\mathbf{q}$; see Fig. 3(a). Ruppert's algorithm is analyzed by considering circumcenter sequences of different lengths. Let $\mathbf{q}$ be a midpoint which is inserted at the end of a length-k circumcenter sequence, and let s_1 and s_2 be the input segments containing $\mathbf{p}_{k+1}(\mathbf{q})$ and $\mathbf{q}$, respectively; if $\mathbf{p}_{k+1}(\mathbf{q})$ is an input point, treat it as a disjoint segment from s_2 in the following analysis.
Length-0 Circumcenter Sequences. The non-acute input requirement is designed to ensure that for a length-0 circumcenter sequence $s_1 \cap s_2 = \emptyset$. Then

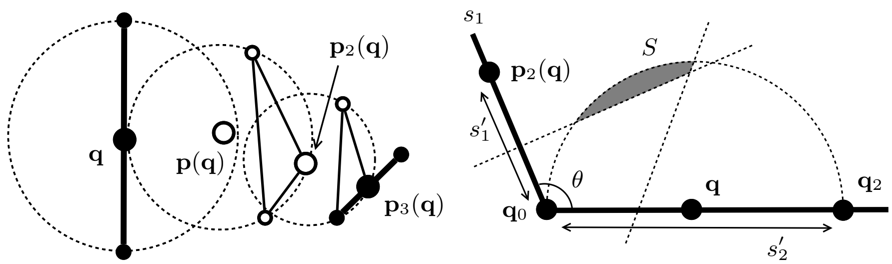

(a) A length-2 circumcenter sequence. (b) $\mathbf{p}(\mathbf{q})$ must lie in the shaded region.

Fig. 3. Diagrams for the proof of Theorem 3

$$\mathrm{lfs}(\mathbf{q}) \leq |\mathbf{q} - \mathbf{p}(\mathbf{q})| + r_{\mathbf{q}} \leq 2r_{\mathbf{q}}. \tag{1}$$

Length-2+ Circumcenter Sequences. We consider an inductive estimate associated with a length-2+ circumcenter sequence.

$$\begin{aligned}\mathrm{lfs}(\mathbf{q}) &\leq |\mathbf{q} - \mathbf{p}(\mathbf{q})| + \mathrm{lfs}(\mathbf{p}(\mathbf{q})) \leq r_{\mathbf{q}} + |\mathbf{p}(\mathbf{q}) - \mathbf{p}_2(\mathbf{q})| + \mathrm{lfs}(\mathbf{p}_2(\mathbf{q})) \\ &\leq r_{\mathbf{q}} + r_{\mathbf{p}(\mathbf{q})} + r_{\mathbf{p}_2(\mathbf{q})} + \mathrm{lfs}(\mathbf{p}_3(\mathbf{q})) \leq \left(\bar{C} + C_\alpha \sqrt{2}(2\sin\alpha)^2\right) r_{\mathbf{q}}. \end{aligned} \tag{2}$$

Length-1 Circumcenter Sequences. If $s_1 \cap s_2 = \emptyset$, then

$$\mathrm{lfs}(\mathbf{q}) \leq |\mathbf{q} - \mathbf{p}_2(\mathbf{q})| \leq |\mathbf{q} - \mathbf{p}(\mathbf{q})| + |\mathbf{p}(\mathbf{q}) - \mathbf{p}_2(\mathbf{q})| \leq 2r_{\mathbf{p}(\mathbf{q})} \leq 2\sqrt{2}r_{\mathbf{q}}. \tag{3}$$

Miller, Pav, and Walkington studied circumcenter sequences beginning and ending with vertices on the same input segment [3]. They showed that any such circumcenter sequence has length at least three. Thus $s_1 \neq s_2$.

Next suppose s_1 and s_2 are adjacent input segments. Let $\mathbf{q}_0$ denote the input vertex shared by s_1 and s_2. Let $s_2' \subset s_2$ be the segment which is split when inserting $\mathbf{q}$. First we claim that $q_0 \in s_2'$; i.e., s_2' is at the end of s_2. If $q_0 \notin s_2'$, then $\mathrm{dist}(q_0, s_2') \geq |s_2'|$ since all midpoints are inserted by splitting segments in half. Moreover, the non-acute input assumption ensures that q_0 is the nearest point on s_1 to the diametral ball of s_2', $B(s_2')$: $\mathrm{dist}(B(s_2'), s_1) \geq |s_2'|$. However $\mathbf{p}(\mathbf{q})$ lies in the diametral ball of s_2', so $r_{\mathbf{p}(\mathbf{q})} \leq \frac{1}{\sqrt{2}}|s_2'|$. Since $\mathbf{p}_2(\mathbf{q}) \in s_1$ and $\mathbf{p}(\mathbf{q}) \in B(s_2')$, $\frac{1}{\sqrt{2}} \geq r_{\mathbf{p}(\mathbf{q})} = |\mathbf{p}(\mathbf{q}) - \mathbf{p}_2(\mathbf{q})| \geq \mathrm{dist}(B(s_2'), s_1) \geq |s_2'|$. This contradiction means $q_0 \in s_2'$.

So $q_0 \in s_2'$ as in Fig. 3(b). Let s_1' be the subsegment of s_1 containing $\mathbf{q}_0$. First suppose that $\mathbf{p}_2(\mathbf{q}) \notin s_1'$. Let $\mathbf{q}^*$ be the endpoint of s_1' opposite $\mathbf{q}_0$. Then,

$$\begin{aligned}\mathrm{lfs}(\mathbf{q}) &\leq |\mathbf{q} - \mathbf{q}^*| + \mathrm{lfs}(\mathbf{q}^*) \leq r_{\mathbf{q}} + r_{\mathbf{q}^*} + C_\alpha r_{\mathbf{q}^*} \leq r_{\mathbf{q}} + (1 + C_\alpha)\frac{1}{2} r_{\mathbf{p}_2(\mathbf{q})} \\ &= r_{\mathbf{q}} + (1 + C_\alpha)\frac{1}{2}|\mathbf{p}_2(\mathbf{q}) - \mathbf{p}(\mathbf{q})| \leq \left(1 + \frac{1}{\sqrt{2}} + \frac{C_\alpha}{\sqrt{2}}\right) r_{\mathbf{q}}. \end{aligned} \tag{4}$$

Next suppose that $\mathbf{p}_2(\mathbf{q}) \in s_1'$ as in Fig. 3(b). Since $\mathbf{p}(\mathbf{q})$ is the circumcenter of a triangle with vertex $\mathbf{p}_2(\mathbf{q}) \in s_1$ and $\mathbf{p}(\mathbf{q})$ encroaches s_2',

$$|\mathbf{p}(\mathbf{q}) - \mathbf{p}_2(\mathbf{q})| \leq \min\left(|\mathbf{p}(\mathbf{q}) - \mathbf{q}_0|, |\mathbf{p}(\mathbf{q}) - \mathbf{q}_2|\right) \text{ and } |\mathbf{p}(\mathbf{q}) - \mathbf{q}| \leq \frac{1}{\sqrt{2}}|s_2'|.$$

The set of points S satisfying these inequalities is shaded in Fig. 3(b). Let

$$f(\theta, |s_1'|, |s_2'|) = \max_{\mathbf{p}(\mathbf{q})^* \in S} |\mathbf{p}_2(\mathbf{q}) - \mathbf{p}(\mathbf{q})|, \text{ and } d^* = \min_{\frac{f(\theta,|s_1'|,|s_2'|)}{|s_1'|} < \frac{1}{2\sin\alpha}} \frac{r_{\mathbf{q}}}{r_{\mathbf{p}_2(\mathbf{q})}},$$

where we emphasize that f only depends on the lengths and angle between s_1 and s_1. If $\frac{f(\theta,|s_1'|,|s_2'|)}{|s_1'|} < \frac{1}{2\sin\alpha}$, then no suitable $\mathbf{p}(\mathbf{q})$ exists as the circumcenter of an appropriate poor-quality triangle and with parent $\mathbf{p}_2(\mathbf{q})$.

If $d^* \geq 2^{-1/4}$ at most three consecutive length-1 circumcenter sequences can occur: there are at most four adjacent segments at a given input vertex (by the non-acute input requirement) and the fourth insertion radius is not small enough to re-split a segment. Considering the worst case (three consecutive length-1 circumcenter sequences preceded by two circumcenters) and applying the inductive hypothesis,

$$\operatorname{lfs}(\mathbf{q}) \leq \bar{C} r_{\mathbf{q}} + C_\alpha r_{\mathbf{p}_9(\mathbf{q})} \leq \left[\bar{C} + C_\alpha (d^*)^{-3}\sqrt{2}\,(2\sin\alpha)^2\right] r_{\mathbf{q}} \qquad (5)$$

where $\bar{C}$ is another constant independent of α.

Based on the definition, d^* can be computed for a specified small angle threshold α. (This computation has been omitted to satisfy page restrictions.) We find that for $\alpha \lessapprox 22.2°$, $d^* \geq 2^{-1/4}$ and $(d^*)^{-3}\sqrt{2}\,(2\sin\alpha)^2 < 1$, allowing C_α to be selected large enough such that $\bar{C} + C_\alpha(d^*)^{-3}\sqrt{2}\,(2\sin\alpha)^2 < C_\alpha$. This constant is also sufficient in the other cases (1), (2), (3) and (4) which completes the proof. □

References

1. L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. In *Proc. 9th Symp. Comput. Geom.*, pages 274–280, 1993.
2. H. Erten and A. Üngör. Quality triangulations with locally optimal Steiner points. *SIAM J. Sci. Comput.*, 31:2103–2130, 2009.
3. G. Miller, S. Pav, and N. Walkington. When and why Ruppert's algorithm works. In *Proc. 12th Int. Meshing Roundtable*, pages 91–102, 2003.
4. S. Pav. *Delaunay Refinement Algorithms.* PhD thesis, Carnegie Mellon, 2003.
5. J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms*, 18(3):548–585, 1995.
6. J. Shewchuk. Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator. In M. Lin and D. Manocha, editors, *Applied Computational Geometry Towards Geometric Engineering*, pages 203–222. Springer, 1996.
7. J. Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Comput. Geom.*, 22(1–3):86–95, 2002.