An Optimal Bound for High-quality Conforming Triangulations[1]

Tiow-Seng Tan

Department of Information Systems and Computer Science
National University of Singapore
Lower Kent Ridge Road
Singapore 0511
email: tants@iscs.nus.sg

August 8, 1995

Abstract

This paper shows that for any plane geometric graph $\mathcal{G}$ with n vertices, there exists a triangulation $\mathcal{T}$ that conforms to $\mathcal{G}$, i.e. each edge of $\mathcal{G}$ is the union of some edges of $\mathcal{T}$, where $\mathcal{T}$ has $O(n^2)$ vertices with each angle of its triangles measuring no more than $\frac{11}{15}\pi$. Additionally, $\mathcal{T}$ can be computed in $O(n^2 \log n)$ time.

Keywords. Discrete and computational geometry, two dimensions, plane geometric graph, plane straight-line graph, triangulation, Steiner triangulation, point placement

[1]Research is partially supported by the National University of Singapore under grant RP940641.

1 Introduction

In many engineering applications such as finite element analysis [7], surface interpolation [14], and shape reconstruction [6], two and higher-dimensional domains are frequently decomposed into small and simple elements before numerical computation. One particularly important class of decompositions is the so-called simplicial cell complex, usually referred to as *triangulation*, where a domain is decomposed into simplices—triangles in two and tetrahedra in three dimensions—such that the intersection of two simplices is either empty or a face of both.

For a given domain such as a polygonal region or, more generally, a plane graph with straight edges, there are clearly many ways to decompose or triangulate it into a triangulation, with or without the addition of new vertices to the domain. But accuracy and efficiency of an engineering computation impose optimal criteria such as the triangle shape (with bounds on triangle angles away from 0 and π) and vertex size (with bounds on the number of new vertices), respectively. These criteria of shape and size are somewhat conflicting—good triangles may be the result of adding new vertices. Automatic generation of triangulations has been a subject for research since the 1960s. Nevertheless, many interesting results, both practical and theoretical ones, have been discovered in the recent years too; see, for example, [5, 8, 9, 10, 15, 16, 17, 18, 19, 20, 21] and the survey [4].

This paper considers triangulating a plane geometric graph, i.e. a plane graph with straight edges, using triangles with no large angles. Such resulting triangulations have potential applications in the area of finite element and surface approximation; see, for example, [1, 2, 12]. This paper shows that triangulating a plane geometric graph of n vertices using angles no larger than $\frac{11}{15}\pi$ requires at most $O(n^2)$ new vertices. The previous result by Mitchell [17] achieves angle bound of $\frac{7}{8}\pi$ with $O(n^2 \log n)$ new vertices. This paper thus provides significantly better bounds on triangle shape and vertex size. The proof is constructive with relatively simple steps. The detailed argument about its correctness is however involved, requiring many new insights besides adapting some results from [17]. Note that the quadratic bound on the vertex size is within a constant factor of worst case optimal; see [3, 17] for discussions on the lower bound construction due to M. S. Paterson.

The paper is organized as follows. Section 2 formalizes the problem. Section 3 provides the outline of the method that proves the quadratic bound, and Sections 4 to 9 discuss its details for angle bound $\alpha = \frac{3}{4}\pi$. Then, Section 10 provides details on implementing the proof and extending it to $\alpha = \frac{11}{15}\pi$, and Section 11 concludes the paper. Appendix A documents some technical details required in Section 6.

2 Some Terminologies

We first introduce some notions, then define the problem.

Plane Geometric Graphs. Let S be a set of n *points* or *vertices* in $\Re^2$. An *edge* is a closed line segment connecting two points. Let E be a collection of edges determined by vertices of S. Then $\mathcal{G} = (S, E)$ is a *plane geometric graph* if

(i) no edge contains a vertex other than its endpoints, that is, $ab \cap S = \{a, b\}$ for every edge $ab \in E$, and

(ii) no two edges *cross*, that is, $ab \cap cd \in \{a, b\}$ for every two edges $ab \neq cd$ in E.

One example of a plane geometric graph is a (simple) *polygon* where E forms a single cycle. The cycle is the *boundary* of the polygon. It subdivides $\Re^2$ into a bounded face, its *interior*, and an unbounded face, its *exterior*. A polygon with four edges (or sides) is a *quadrilateral*.

Triangulations. A *triangulation* is a plane geometric graph $\mathcal{T} = (S, E)$ so that E is maximal. By maximality, edges in E bound the convex hull $ch(S)$ of S, i.e. the smallest convex set in $\Re^2$ that contains S, and subdivide its interior into disjoint faces bounded by triangles. With reference to a polygon, we talk about its triangulation as restricted to only its interior.

Conforming Triangulations. A plane geometric graph $\mathcal{G}' = (S', E')$ *conforms* to another such graph $\mathcal{G} = (S, E)$ if $S \subseteq S' \subset ch(S)$ and each edge in E is a union of edges in E'. A triangulation $\mathcal{G}'$ conforms to a plane geometric graph $\mathcal{G}$ is called a *conforming triangulation* of $\mathcal{G}$. Each vertex in $S' - S$ is termed a *Steiner vertex*. The problem studied here is as follows:

> Given a plane geometric graph $\mathcal{G} = (S, E)$, find a conforming triangulation of $\mathcal{G}$ with a small vertex set and with each angle of its triangles measuring at most α.

3 The Outline of Construction

Given a plane geometric graph $\mathcal{G} = (S, E)$ where $|S| = n$, the algorithm of Edelsbrunner, Tan and Waupotitsch [11] can augment it by edges to a triangulation that minimizes its maximum angle over all possible augmentations. We are done if $\mathcal{T}$ has angles each measuring at most the targeted angle bound α. Because of this and the fact that the size of any triangulation of $\mathcal{G}$ is a constant factor of the size of $\mathcal{G}$, we assume in our discussion that the given plane geometric graph

$\mathcal{G} = (S, E)$ is actually a triangulation $\mathcal{T}$ and $\mathcal{T}$ has some bad angles. A *bad angle* is an angle with measure larger than α, whereas a *good angle* is otherwise. For our discussion, the targeted angle bound α is set to $\frac{3}{4}\pi$ unless specified otherwise as in Section 10.

The main difficulty of the problem is as follows. Suppose $\mathcal{T}$ has a bad angle at q in $\triangle pqr$. Then, we can add a Steiner vertex, say t, on pr to subdivide the bad angle at q with the edge tq. This resolution however creates a "large angle" of π at t in the triangle, say prs, sharing the edge pr with $\triangle pqr$. In other words, the trouble of a large angle has propagated to $\triangle prs$ and another subdivision is necessary. Unless ts can subdivide this large angle into two good angles or $\triangle prs$ does not exist as pr is actually an edge bounding $ch(S)$, we need to add one Steiner vertex on either ps or rs and the situation continues to propagate each time to another adjacent triangle. The sequence of Steiner vertices generated in this manner can be viewed as forming a *propagation path* with edges joining two consecutive Steiner vertices. A successful approach to add a bounded number of Steiner vertices has to terminate each propagation path effectively. This is achieved in this paper by fences (Section 4) and dead-ends (Section 5).

The proposed construction has six major steps. It first subdivides each triangle of $\mathcal{T}$ into three quadrilaterals (Section 4) and works with these quadrilaterals throughout the construction. In the second and third steps (Sections 5 and 6), it decides where to add Steiner vertices forming propagation paths. To control the number of Steiner vertices, the construction bounds the number of vertices in each propagation path to linear in size. In the fourth and fifth steps (Sections 7 and 8), the construction manipulates edges of propagation paths to remove all crossings among edges of propagation paths. As a result, some Steiner vertices added by the third step may be removed. In the last step (Section 9), the construction triangulates each quadrilateral one by one using edges of propagation paths lying in the quadrilateral. The output is the union of triangulations of all quadrilaterals.

4 Erecting Fences

For a triangle pqr, let c be the center of its inscribing circle. Each edge joining c with the perpendicular projection of c onto a side of $\triangle pqr$ is a *spoke* of $\triangle pqr$. For example in Figure 4.1, cr', cq' and cp' are the three spokes of $\triangle pqr$. For convenience, we let $\beta = \pi - \alpha$.

Step 1. Erecting Fences.

For each $\triangle pqr \in \mathcal{T}$, add its spokes cr', cp' and cq' to subdivide it into three quadrilaterals $pq'cr', rq'cp'$, and $qr'cp'$ as shown in Figure 4.1. Next, mark $pq'cr'$ as a *fence* if $\angle p \geq \beta$ else as a *non-fence*. Similarly, mark $rq'cp'$ according to $\angle r$ and mark $qr'cp'$ according to $\angle q$.

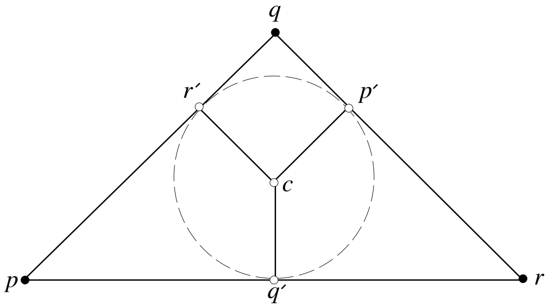

Figure 4.1: Note that $|pr'| = |pq'|$, $|rp'| = |rq'|$, $|qr'| = |qp'|$, and $|cr'| = |cp'| = |cq'|$.

The main purpose of fences is to stop each newly created Steiner vertices (introduced by later steps) lying on a boundary edge of some fence from further generating Steiner vertices. The next lemma shows that each fence with or without Steiner vertices on its boundary edges can always be triangulated with only good triangles. A *good triangle* is a triangle with no bad angle. For convenience, we will treat each edge bounding $ch(S)$ as a (degenerate) fence since fences and edges of $ch(S)$ both terminate the generation of more Steiner vertices.

Lemma 4.1 Let $pq'cr'$ be a fence with $\angle q' = \angle r' = \frac{\pi}{2}$ and $\angle p \geq \beta$. Then, the region bounded by $pq'cr'$, possibly with other vertices on pr' and pq', can always be triangulated with only good triangles.

Proof. When there are no other vertices on pr' and pq', we are done as $pq'cr'$ can be triangulated with two right angle triangles $pq'c$ and $pr'c$. Now, suppose there are other vertices along pr' and pq'. Let the vertices along pr' be sorted as $p = r_0, r_1, \ldots, r_m = r'$ with increasing distance from p. Similarly, let $p = q_0, q_1, \ldots, q_l = q'$ be the corresponding sequence along pq'. We now triangulate $pq'cr'$ to prove the claim. There are two cases: $2\beta \leq \angle p < \pi$ and $\beta \leq \angle p < 2\beta$. For the first case, we simply add edges cq_i, for $i = 0$ to $l-1$, and edges cr_j, for $j = 1$ to $m-1$. Each triangle obtained is of the form cq_iq_{i+1} or cr_jr_{j+1}. For the former, we have $\angle cq_iq_{i+1} < \frac{\pi}{2}$ and $\angle q_icq_{i+1} < \frac{\pi}{2}$ because $\triangle cq_iq'$ has a right angle at q'. And, $\angle cq_{i+1}q_i < \pi - \beta = \alpha$ because in $\triangle cq_0q_{i+1}$ we have $\angle q_0 = \frac{1}{2}\angle r'pq' \geq \beta$. So $\triangle cq_iq_{i+1}$ is good. Similarly, we have good $\triangle cr_jr_{j+1}$.

We next consider the case where $\beta \leq \angle p < 2\beta$. Initially, set $i = 1, j = 1$. Add the edge r_jq_i, and then increment j if r_j is closer to p than q_i else increment i. Repeat the previous statement until $pq'cr'$ is triangulated. Triangle r_mq_lc, the last one obtained, has clearly good $\angle r_m$ and $\angle q_l$, and also good $\angle c$ because $\angle c = \pi - \angle p \leq \pi - \beta = \alpha$. And, $\triangle pr_1q_1$, the first one obtained, is good because $\beta \leq p < 2\beta < \alpha$. The rest of the triangles are either of the form $r_{j-1}r_jq_i$ or $q_{i-1}q_ir_j$. We just consider $\triangle r_{j-1}r_jq_i$; similar argument applies to $\triangle q_{i-1}q_ir_j$. By construction, we have r_{j-1} closer to p than q_i, i.e. $\angle r_{j-1} > \angle q_i$ in $\triangle pr_{j-1}q_i$. Since $\angle p < 2\beta$, we have $\angle q_ir_{j-1}p > \frac{\pi-2\beta}{2}$. Hence, $\triangle r_{j-1}r_jq_i$ has $\angle r_{j-1} = \pi - \angle q_ir_{j-1}p < \frac{\pi}{2} + \beta \leq \alpha$. And, its $\angle r_j$ and $\angle q_i$ are also good because $\angle r_jpq_i \geq \beta$. So, it is a good triangle and the proof is complete. $\square$

5 Setting Traps

The second step, consisting of the planning and the adding sub-steps, is to augment $\mathcal{T}$ with linear number of Steiner vertices, termed dead-ends. We need some terminologies for our discussion.

Paths. A *path* with origin p_0 and last vertex p_k is a sequence of vertices $p_0, p_1, p_2, \ldots, p_k$ where each vertex lies on some edge of $\mathcal{T}$ and each line segment joining consecutive vertices p_i and p_{i+1} does not cross any edges of quadrilaterals. The line segment joining p_i and p_{i+1} is an *edge* of the path or, more specifically, a *directed edge* $\vec{p_i p}_{i+1}$ with *tail* p_i and *head* p_{i+1}. For $\vec{p_i p}_{i+1}$ lying in $pq'cr'$ with $p_i \in pq'$ and $p_{i+1} \in pr'$, we refer to $\angle p_i p_{i+1} p$ and $\angle p_i p_{i+1} r'$ as the two angles at its head, and $\angle p_{i+1} p_i p$ and $\angle p_{i+1} p_i q'$ as the two angles at its tail. Formally, each *propagation path* is a path where the two angles at the tail of each of its directed edges are good.

Path P with vertices $p_0, p_1, p_2, \ldots, p_k$ and path Q with vertices $q_0, q_1, p_2, \ldots, q_k$ are said to be *coherent* if $p_i p_{i+1}$ and $q_i q_{i+1}$, for each $0 \le i \le k-1$, lie in the same quadrilateral. They are *parallel* if $p_i p_{i+1}$ and $q_i q_{i+1}$, for each $0 \le i \le k-1$, lying in the same quadrilateral are parallel. The *length* of a path is equal to its number of directed edges.

A *backward path* is a path where the two angles at the head of each directed edge are α and β, and all good angles of β at the heads are on the same side along the path. Analogously, a *forward path* is a path where the two angles at the tail of each directed edge are α and β, and all good angles of β at the tails are on the same side along the path.

Planning Dead-ends. To plan for sufficient number of dead-ends, Step 2a computes for each point x of each quadrilateral edge whether there exists a *good propagation path*, i.e. a propagation path of length K with origin x and last vertex at an endpoint of a spoke where $K < 6n$ is two times the number of edges of $\mathcal{T}$. From this, we can view each quadrilateral edge, with respect to its quadrilateral, consisting of alternating *good* and (open) *bad segment* where the former includes points that have good propagation paths whereas the latter otherwise. Some of these bad segments will be used by Step 2b to add dead-ends.

Step 2a. Planning Dead-ends.

Consider each non-fence $pq'cr'$; see Figure 5.1. We work on q' and similarly on r'. Trace a backward path starting at q' with its second vertex q''_1 on pr' such that $\angle q' q''_1 r' = \beta$. We first mark $q''_1 r'$ as a (part of some) good segment with respect to $pq'cr'$. At each extension of the backward path from length i at q''_i to length $i+1$ at q''_{i+1} where $q''_i \in p_i q'_i$ and $q''_{i+1} \in p_i r'_i$ of non-fence $p_i q'_i c_i r'_i$ (with $\angle q'_i = \angle r'_i = \frac{\pi}{2}$), we mark $q''_{i+1} t''_{i+1} = q''_{i+1} s''_{i+1} \cap p_i r'_i$ as a (part of some) good segment with respect to $p_i q'_i c_i r'_i$ where s''_{i+1} is the point on the line through $p_i r'_i$ with $\angle t''_i s''_{i+1} q''_{i+1} = \beta$ and t''_i is an endpoint of the good segment $q''_i t''_i$ identified (with respect to $p_{i-1} q'_{i-1} c_{i-1} r'_{i-1}$) by the previous extension. The backward path terminates when one of the following conditions holds: it reaches an edge of a fence, or an edge for the third time, or it will cross a spoke in the next extension.

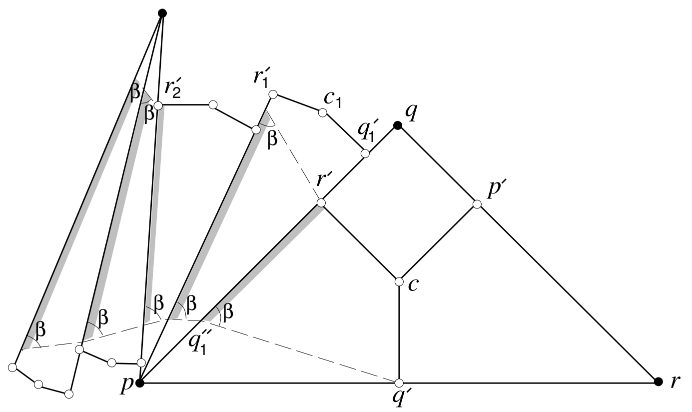

Figure 5.1: The first five good segments identified due to the backward path with origin q' are shown shaded.

A few remarks are in order. First, we could have included in the definition of a good propagation path for one of bounded length to some edge of a fence. This is, however, not efficient for computation. Second, for the same reason, Step 2a identifies some but not all good segments, and simply takes the remaining as bad segments. Third, vertices of backward paths in Step 2a as well as forward paths in Step 2b are for discussion purposes only and are not Steiner vertices added to $\mathcal{T}$; only dead-ends of traps are Steiner vertices added to $\mathcal{T}$ by Step 2b.

Adding Dead-ends. A *trap* consists of a *base* xy, *dead-ends* x' and y', and *boundary paths* P_x from x to x' and P_y from y to y'. Note that dead-ends x', y' may be the same point, and boundary paths P_x and P_y, excluding their last edges, are coherent forward paths starting at x and y so that those good angles of β at tails of their directed edges are inside the region bounded by P_x, P_y, xy and $x'y'$.

Step 2b. Adding Dead-ends.

Consider each non-fence $pq'cr' \subset pqr$; see Figure 5.2(i). We discuss the following for r', and treat q' similarly. Let $p_1r_1'c_1q_1' \not\subset \triangle pqr$ be the non-fence incident to pq (if pq is not a convex hull edge) with $r' \in p_1q_1'$. If r' is an endpoint of a bad segment $x_1r' \subseteq p_1q_1'$ with respect to $p_1r_1'c_1q_1'$, we construct a trap with base x_1r' in the direction into $p_1r_1'c_1q_1'$. Next, let $p_zq_z'c_zr_z'$ be the non-fence containing the last edge of the backward path traced from r' at Step 2a. If the head of this last edge is an endpoint of some bad segment with respect to $p_zq_z'c_zr_z'$, we construct in the direction into $p_zq_z'c_zr_z'$ a trap with base at this bad segment; see Figure 5.2(ii).

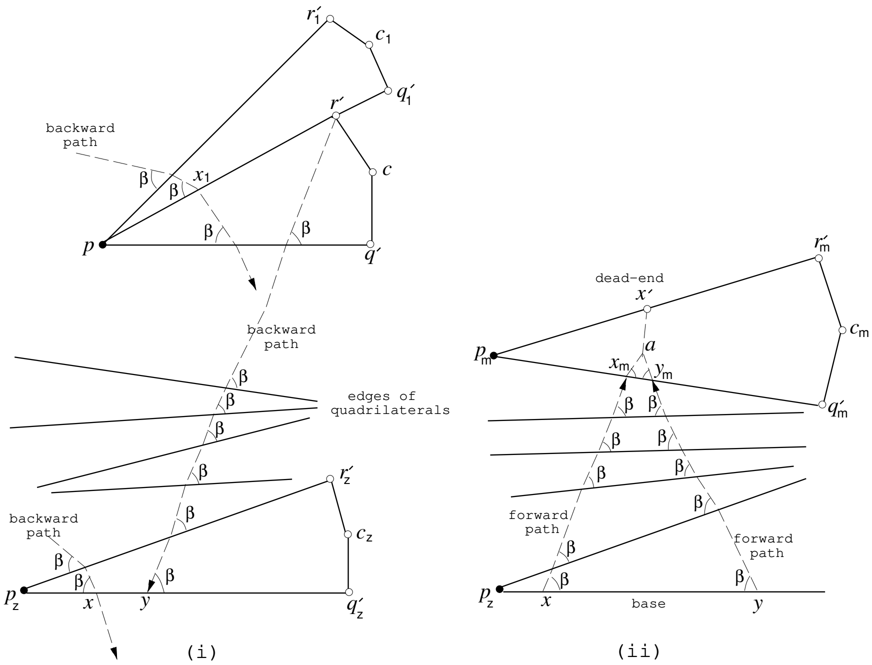

Figure 5.2: (i) shows the backward path traced from r' in Step 2a terminating at y. Segment xy is a bad segment with respect to $p_z q'_z c_z r'_z$, so a trap with base xy in the direction into $p_z q'_z c_z r'_z$ is constructed as shown in (ii).

To construct in the direction into, say $p_z q'_z c_z r'_z$, a trap with base xy on $p_z q'_z$ where $y \in xq'_z$, we trace two coherent forward paths at x and y into $p_z q'_z c_z r'_z$ so that the paths move closer to each other as their lengths increase. Both paths continue until they have just reached, say $\triangle p_m q_m r_m$ (with endpoints of spokes $p'_m \in r_m q_m, r'_m \in p_m q_m, q'_m \in p_m r_m$), at, say $x_m \in p_m q'_m$ (possibly $x_m = x$) and $y_m \in x_m q'_m$ (possibly $y_m = y$), where

1. the two forward paths will intersect at a point, say a, in the interior of $p_m q'_m c_m r'_m$ when further extended for one more step (as in Figure 5.2(ii)), or
2. $p_m q'_m$ is the very first edge visited for the third time by both forward paths.

For case 1, we consider $x' \in p_m q_m$ where ax' is perpendicular to $x_m y_m$. We take x' as the dead-end if ax' does not intersect $c_m r'_m$; otherwise, take r'_m as the dead-end. And, we complete the boundary paths of the trap by joining both forward paths to the dead-end. For case 2, x_m and y_m are the dead-ends and both forward paths are the boundary paths of the trap. The following result is immediate because of case 2.

Lemma 5.1 The length of each boundary path of a trap is at most K. □

In the above, we ignore the fact that (i) x may be the same as p_z so a forward path from p_z is not well defined, (ii) the two coherent forward paths traced may first hit an edge of a fence, or (iii) the two forward paths reach $\triangle p_m q_m r_m$ such that $x_m \in p_m q'_m$ and $y_m \in q'_m r_m$ where we can no longer maintain coherent forward paths in their extensions. In all these, we can actually ignore the construction of a trap. The reason for (i) is explained in Section 6 (in the proof of Lemma 6.3), for (ii) is that each point on xy after all has propagation paths of linear length to an edge of a fence, and for (iii) is as follows. The point x_m is so that $x_m \in r''_m q'_m \subset p_m q'_m$ of non-fence $p_m q'_m c_m r'_m$ where $\angle r'_m r''_m q'_m = \beta$, and similar statement holds for y_m; thus, xy is a good segment not identified by Step 2a. This is so because the region $\mathcal{R}$ enclosed by $P_x, x_m y_m, P_y, yx$ (where P_x and P_y are the two forward paths from x to x_m and y to y_m, respectively) is emptied of endpoints of spokes, except for $q'_m \in x_m y_m$. So, if r''_m is an interior point of $x_m q'_m$, then the backward path from r'_m traced at Step 2a inside $\mathcal{R}$ is parallel to P_x and reaches xy, and Step 2a should have identified some part of xy as a good segment.

Lemma 5.2 The number of traps constructed in Step 2b is at most $14n$.

Proof. Consider $\triangle pqr \in \mathcal{T}$ with endpoints of spokes $p' \in rq, r' \in pq, q' \in pr$ as in Figure 5.1. It is easy to check that at least one of the three quadrilaterals, say $qr'cp'$, is a fence. So, we trace at most one backward path each at r' and p', and at most two at q'. These backward paths result in at most eight traps but at most seven is necessary since the two with bases at q', if any, can always be combined into one (and we pretend we did so in the construction). The claim follows as there are at most $2n$ triangles. ◻

6 Generating Paths

The necessity to propagate vertices in computing triangulation with angles bounded away from π is a consequence of the following simple lemma introduced in [11].

Lemma 6.1 If xy is an edge in a triangulation $\mathcal{A}$ of a point set S, then $\mu(\mathcal{A}) \geq \max_{s \in S} \angle xsy$ where $\mu(\mathcal{A})$ denotes the largest angle in $\mathcal{A}$. ◻

In effect, propagation is to subdivide each long edge into smaller ones so as to remove large angles subtended by the edge.

We adapt a few terminologies from [17]. Let s be a point on pq' of non-fence $pq'cr'$ with $\angle q' = \angle r' = \frac{\pi}{2}$. The *cone* at s consists of all points t of $pq'cr'$ such that both $\angle tsp$ and $\angle tsq'$ are at most α. And, the *maw* of s is the portion of the cone at s on the boundary of $pq'cr'$, except for s. Note that each point on the maw of s is a candidate for extending a propagation path from s. Next, the iterative construction of union of cones at points in maws, starting with s, results in the *horn* of s. Initially at stage 0, the horn is the cone at s. The horn at stage $i+1$ is the union of cones at points in the union of maws, or simply maw, at stage i; see Figure 6.1. It is clear that the two paths bounding the horn of s are actually forward paths starting at s. They are referred as the *boundary paths of the horn.*

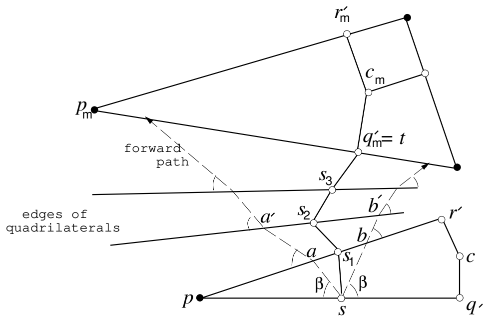

Figure 6.1: The horn at stage 0 is the cone sab, and the maw at this stage is ab. The horn at stage one is the region bounded by the two forward paths from s to a' and s to b', and the maw $a'b'$ at this stage. At stage three, the horn at s contains q'_m, and a propagation path such as s, s_1, s_2, s_3, q'_m can be constructed from q'_m back to s_3 (with q'_m in the maw of s_3), then s_3 to s_2 (with s_3 in the maw of s_2), and so on.

To generate a propagation path with origin s, the construction first computes iteratively the horn of s until it intersects some point t which can be an endpoint of a spoke (on some edge of $\mathcal{T}$), or a dead-end, or any point on an edge of a fence, or a vertex of some existing propagation path, then works straightforwardly from t back to s. Besides the above, t can possibly be some point common to cones of the horn from different stages, which happens when the horn of s *self-intersects*; see Figure 6.2. For all these, we need to place two restrictions on the above choices of t so as to say that a propagation path *terminates properly* and to call it a *properly-terminating propagation path*: P can terminate at a dead-end only when the horn of s has entered the corresponding trap through the base of the trap, and it can terminate at vertex t of some propagation path P' (possibly P itself) only when the last edge of P and some edge of P' share t as head in the same non-fence; see Figure 6.3. These are to avoid having a Steiner vertex (other than an endpoint of a spoke) to serve both as a head of some directed edge and a tail of another directed edge lying in the same non-fence. Also, the construction maintains that propagation paths do not cross spokes. This is possible because of the simple fact that whenever a horn intersects spokes cq' or cr' inside $pq'cr'$ upon entering from pq', it also contains r'.

Step 3. Generating Paths.

Consider endpoints of spokes that are on edges of $\mathcal{T}$ one by one. With each such endpoint as the first vertex, generate a (short) properly-terminating propagation path of linear length. Then, for each dead-end that terminates some propagation paths constructed so far, generate a (short) properly-terminating propagation path of linear length with this dead-end as the origin. At the end, those dead-ends that do not terminate any propagation paths are removed.

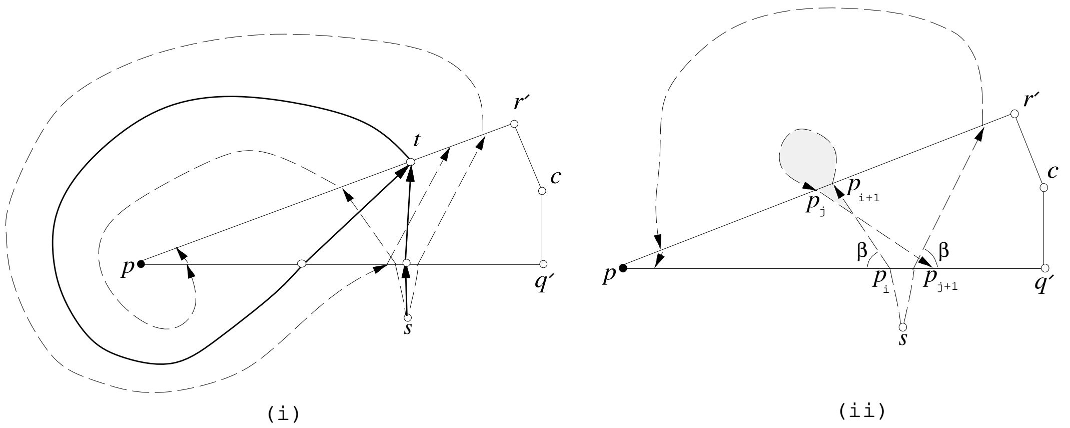

Figure 6.2: Horn self-intersects while passing through $pq'cr'$ in the same direction as in (i) or in the opposite directions as in (ii). The latter is not used to derive propagation paths, whereas the former is used when the intersection contains a vertex such as t whose horn lying inside the horn of s contains itself at some later stage.

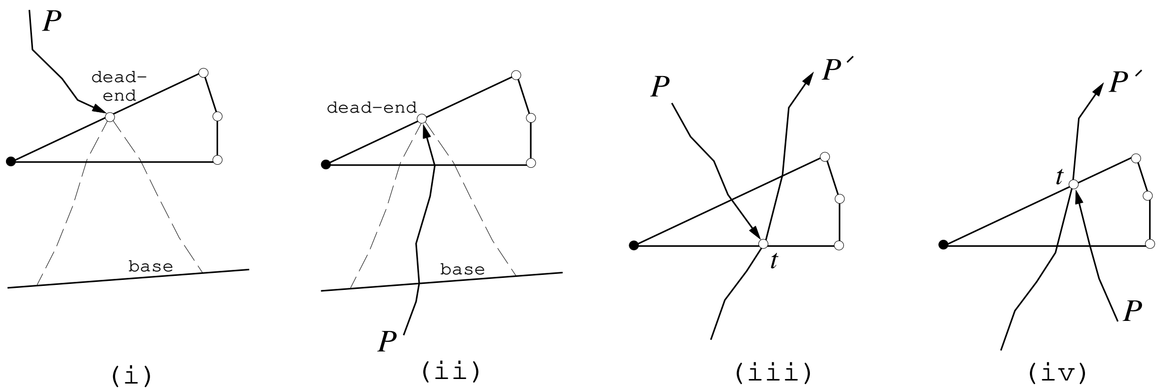

Figure 6.3: Propagation path P terminates properly at a dead-end as in (ii) or at some vertex t of P' as in (iv). The two restrictions prohibits the termination of P as in (i) and (iii).

The next three lemmas imply the quadratic bound on the number of Steiner vertices. Two directed edges $\vec{ab}$ and $\vec{de}$ with endpoints on pq' and pr' of non-fence $pq'cr'$ (with $\angle q' = \angle r' = \frac{\pi}{2}$) are said to have the *same orientation* if a and d (and b and e) are on the same edge of $pq'cr'$, otherwise *opposite orientations*.

Lemma 6.2 Let T be a trap of case 2 (Section 5) with base xy and length of boundary paths L. If the horn of $s \in xy$ entering T self-intersects in opposite directions at a stage less than L, then there exists a propagation path of length L from s that terminates properly at one dead-end of T.

Proof. It suffices to show that the horn of s intersects one boundary path of T and thus contains a dead-end of T at a stage less than L. Suppose non-fence $pq'cr'$ is where the horn first self-intersects in opposite directions. Then, two edges, say $p_i\vec{p}_{i+1}$ and $p_j\vec{p}_{j+1}$, of a boundary path P of the horn cross inside $pq'cr'$. Let p_i precede p_j in P. Suppose $p_i \in pq'$, then $p_{i+1}, p_j \in pr'$ and

$p_{j+1} \in pq'$. Additionally, it is easy to check using incident angles of directed edges that $p_i \in pp_{j+1}$ and $p_j \in pp_{i+1}$ with $\angle p_{i+1}p_ip = \beta$ and $\angle p_{j+1}p_jp = \alpha$, as shown in Figure 6.2(ii), whereas all the other situations (for example, $p_{j+1} \in pp_i$ and $p_{i+1} \in pp_j$ with $\angle p_{i+1}p_ip = \alpha$ and $\angle p_{j+1}p_jp = \beta$) are impossible.

Let P' be the subpath of P from $\vec{p_ip}_{i+1}$ to $\vec{p_jp}_{j+1}$, and let $\mathcal{R}$ be the region bounded by P' (shown shaded in Figure 6.2(ii)). It is clear that one boundary path of T has some part lying in $\mathcal{R}$ and is coherent to P'. Let this part be P'' and let its directed edges be labeled in the same way in increasing order as in P'. Observe that P'' and P' are two forward paths that move closer to each other, as if they are boundary paths of a trap. Moreover, a directed edge of P'' does not cross another of P' of smaller index because of incident angles of directed edges (unless P'' has already crossed a directed edge of P' of the same orientation). Thus, some directed edge of P'' crosses a directed edge of P' of the same orientation since P' self-intersects. ◻

Lemma 6.3 The length of each propagation path constructed by Step 3 is less than $12n$.

Proof. The claim is clearly true when the origin of the propagation path lies on a good segment. So, it suffices to show for each point s lying on bad segment $xy \subseteq p_jq'_j$ with respect to $p_jq'_jc_jr'_j$ (having $\angle q'_j = \angle r'_j = \frac{\pi}{2}$) there exists a properly-terminating propagation path with origin at s, first edge in $p_jq'_jc_jr'_j$, and length less than $12n$.

Let us first review Step 2a on the identification of good segments. Each portion of a good segment identified at each iteration is defined by two endpoints. One is due to the backward path traced from the endpoint of a spoke (such as q' in Figure 5.1). And, the other is either an endpoint of a spoke or obtained by the extension of an endpoint of the good segment identified in the previous iteration. It is not hard to see that some sequence of vertices mentioned in the previous statement form a backward path with the vertex of the former case (such as r' and r'_2 in Figure 5.1) as the origin. To distinguish each such backward path from those explicitly mentioned in Step 2a, we call it an *implicit backward path.*

First we consider at least one of the endpoint of xy, say $x \notin S$ be a vertex of an implicit backward path, say $P_{r'}$, with origin r' and in the direction out of $pq'cr' \subset pqr$. (For example in Figure 5.1, x can be the endpoint on pr'_1 of the dashed line incident to r'.) We trace a forward path P_s with origin s and parallel to $P_{r'}$ until P_s and $P_{r'}$ have vertices separated by an endpoint of a spoke, or until P_s arrives at $qr'cp'$ (which is adjacent to $pq'cr'$). We will discuss the latter with P_s terminating at $u \in qr'$; we can treat the former similarly. We are clearly done if $qr'cp'$ is a fence. If not, we are done too because $u \in p''r' \subset qr'$ where $\angle p'p''r' = \beta$ and P_s plus $\vec{up'}$ is a good propagation path; otherwise $u \in qp''$ which implies that the backward path with origin p' into $qr'cp'$ traced by Step 2a reaches xs, a contradiction to the fact that xy is a bad segment.

Next, we let x and y be vertices of backward paths with origins x_0 and y_0 traced in Step 2a. Then the horn of s at a stage no larger than K encounters one of the three situations in Figure 6.4. Figure 6.4(iii) shows the horn intersecting some good segment, and a propagation path of the

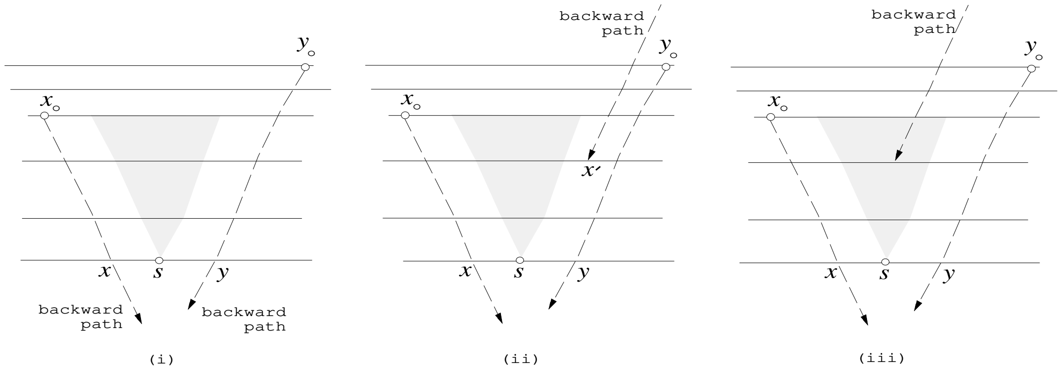

Figure 6.4: Edges of quadrilaterals are drawn as horizontal line segments here. The horn of s (shaded region) encounters either only bad segments as in (i) and (ii), or some good segments as in (iii).

required length can clearly be constructed. Figure 6.4(i) shows the horn entering a trap with x_0 as an endpoint of its base, and Figure 6.4(ii) a trap with x' as an endpoint of its base. Note that the case of x' in (ii) be a vertex of some implicit backward path has already been discussed in the previous paragraph.

Let z be a point common to the horn of s and the base of the trap T that the horn enters. Now if T has one dead-end, we are done. This is because a forward path with origin z and parallel to one boundary path of T meets the other boundary path of T and, thus, the horn of z contains the dead-end of T. In such case, we can derive one properly-terminating propagation path P_z with origin z and length bounded by that of the boundary paths of T. A propagation path from s to z concatenated to P_z gives a properly-terminating propagation path of the required length. On the other hand, if T has two dead-ends (case 2 in Section 5), then an extension of results in [17] (Lemma A.3 in Appendix A) plus Lemma 6.2 imply the existence of a propagation path of the required length. Also, if T was not really constructed because one endpoint of its base coincides with a vertex, say p_1, of $\mathcal{T}$, then the horn of z in this case moves only around edges incident to p_1 and self-intersects at z after one round (by Lemma A.1 in Appendix A), thus a propagation path of the required length that terminates properly at z can be derived. $\square$

Lemma 6.4 Step 3 computes less than $34n$ propagation paths.

Proof. Since a trap has one or two dead-ends, we have in total at most $28n$ dead-ends by Lemma 5.2. So Step 3 generates at most $28n$ propagation paths with origins at dead-ends. The remaining propagation paths are due to the endpoints of spokes of less than $2n$ triangles. $\square$

7 Removing Complications

The next step considers merging of nearby propagation paths. Besides removing redundant portions of paths, i.e. subpaths, this step resolves a complication in triangulating non-fences using edges of propagation paths (in Section 9). Let us label vertices of non-fence $pq'cr'$ (with $\angle q' = \angle r' = \frac{\pi}{2}$) along pq' as $p = q_0, q_1, q_2, \ldots, q_l = q'$ and along pr' as $p = r_0, r_1, r_2, \ldots, r_m = r'$. A *neighboring pair of directed edges* are two directed edges of the same orientation such that the open segment defined by their tails contains no other tail.

Step 4. Merging Paths.

For each non-fence $pq'cr'$, repeat the following as long as there exists a neighboring pair $q_i\vec{r}_{i'}$ and $q_j\vec{r}_{j'}$, $i < j$ and $i' \neq j'$, with $\beta \leq \angle r_{i'}q_jp \leq \alpha$, or, analogously, a neighboring pair $r_i\vec{q}_{i'}$ and $r_j\vec{q}_{j'}$ $i < j$ and $i' \neq j'$, with $\beta \leq \angle q_{i'}r_jp \leq \alpha$. We just describe the former. Let P be some propagation path containing $q_j\vec{r}_{j'}$, and let P' be the shortest subpath of P starting with $q_j\vec{r}_{j'}$ until a vertex z where z can be an endpoint of a spoke, or a vertex on an edge of a fence, or a head common to two or more directed edges, or $z = q_j$ (when $P = P'$). Remove edges and vertices of P', except for its last vertex if shared by other propagation paths. Then if $z \neq q_j$, terminate properly the subpath of P before $q_j\vec{r}_{j'}$ by $q_j\vec{r}_{i'}$.

Lemma 7.1 All propagation paths still terminate properly with merging of paths.

Proof. It suffices to consider propagation paths other than P. Since each vertex, other than an endpoint of a spoke, does not act as both a head and a tail of different directed edges in a non-fence, vertices removed in one iteration are not involved in other propagation paths. Also, when the last vertex z of P' is not removed, it does not result in a bad angle of π because it remains incident to either two directed edges or a directed edge and a spoke at both sides of the edge of $\mathcal{T}$ containing z. So, all propagation paths still terminate properly. ◻

8 Untangling Crossings

For this and the next section, we view each directed edge as an individual with two good angles at its tail, without associating it to a particular propagation path. With reference to non-fence $pq'cr'$ having Steiner vertices along pq' and pr' labeled as in Section 7, the next step is to remove all crossings, in particular, crossings due to directed edges of opposite orientations. It replaces some existing directed edges by new ones while maintaining the following invariants:

I-1. each edge is assigned an orientation such that both angles at its tails are good;

I-2. each Steiner vertex is an endpoint of some directed edge;

I-3. no Steiner vertex is a tail of more than one directed edge;

I-4. directed edges of the same orientation do not cross; and

I-5. for $q_i\vec{r}_{i'}$ and $q_j\vec{r}_{j'}$ with $i < j$ and $i' < j'$, the region defined by their four endpoints can be subdivided into two good triangles using $r_{i'}q_j$. Analogous assertion applies to $r_i\vec{q}_{i'}$ and $r_j\vec{q}_{j'}$.

Initially, I-1, I-2 and I-3 are trivially true. Also, I-4 and I-5 are true; otherwise, it is easy to verify that Step 4 did not complete all possible mergings. In Step 5, the *next available index from* i refers to the smallest index, starting at i and in increasing order, whose corresponding vertex can be used as the head of a new directed edge while maintaining I-4. We say an ordered pair (a, b) is lexicographically smaller than another ordered pair (a', b') if $a < a'$, or $a = a'$ and $b < b'$.

Step 5. Untangling Crossings.

Consider each non-fence $pq'cr'$ one by one. The following is repeated until all crossings are removed. Each iteration involves $q_ir_{i_1}$ crossing $r_jq_{j_1}$, $i < j_1$ and $j < i_1$, such that $(i+j, i_1+j_1)$ is lexicographically the smallest ordered pair.

Case A. $q_i\vec{r}_{i_1}$ crosses $r_j\vec{q}_{j_1}$ (see Figure 8.1).
There are two symmetrical cases depending on which of q_i and r_j has good angles. Let us just discuss the former. For each $q_{j'}\vec{r}_{i_1}$, $i \le j' < j_1$, if $\angle r_jq_{j'}q' \le \alpha$, we move its head from r_{i_1} to r_j (i.e. replace $q_{j'}\vec{r}_{i_1}$ by $q_{j'}\vec{r}_j$). As a result, if r_{i_1} $(\neq r_m)$ is no longer a vertex of any directed edge, we perform the following: if $\angle q_{j_1}r_{i_1}r_m \le \alpha$, add $r_{i_1}\vec{q}_{j_2}$ where j_2 is the next available index from j_1, else add $q_{j_1}\vec{r}_{i_1}$.

Case B. $r_{i_1}\vec{q}_i$ crosses $q_{j_1}\vec{r}_j$ (see Figure 8.2).
We add $q_i\vec{r}_j$ if q_i is not a tail of any directed edge, else add $r_j\vec{q}_i$ if r_j is not a tail of any directed edge. Next, if r_{i_1} is a head of other directed edge, we are done by removing $r_{i_1}\vec{q}_i$. Symmetrically, if q_{j_1} is a head of other directed edge, we are done by removing $q_{j_1}\vec{r}_j$. Otherwise, move the head q_i of all directed edges crossing $q_{j_1}\vec{r}_j$, inclusive of $r_{i_1}\vec{q}_i$, to $q_{j'}$ where $j' = i + 1$.

It is clear that each iteration of either Case A or Case B may introduce new crossings. Still, the process does terminate by removing all crossings. This is because (i) the new crossing introduced by say $q_ar_{a'} \cap r_bq_{b'}$, for $a < b'$ and $b < a'$, is such that $a + b > i + j$, and (ii) there are finitely many

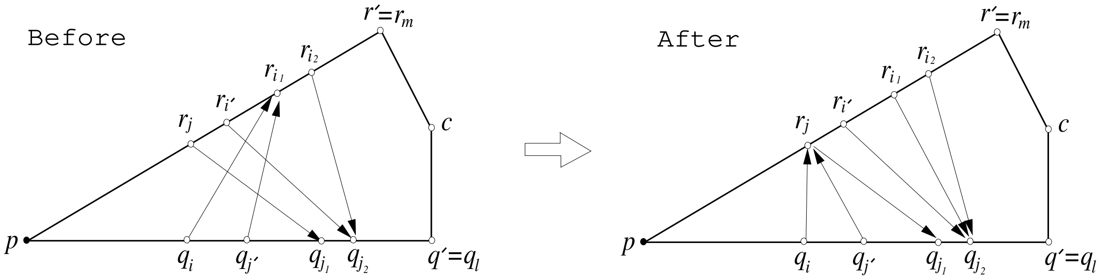

Figure 8.1: Any directed edge e with endpoint on $r_{i'}$ for $j \leq i' < i_1$ crosses $q_i\vec{r}_{i_1}$ by the choice of $(i+j, i_1+j_1)$. Because of this and I-4, the other endpoint q_{j_2} of e is such that $j_2 \geq j_1$, and the orientation of e is opposite to $q_i\vec{r}_{i_1}$. A similar statement applies to directed edges such as $q_{j'}\vec{r}_{i_1}$ for $i \leq j' < j_1$. In the above, $r_{i_1}\vec{q}_{j_2}$ is added to maintain I-2 after those edges with head at r_{i_1} have been moved to r_j.

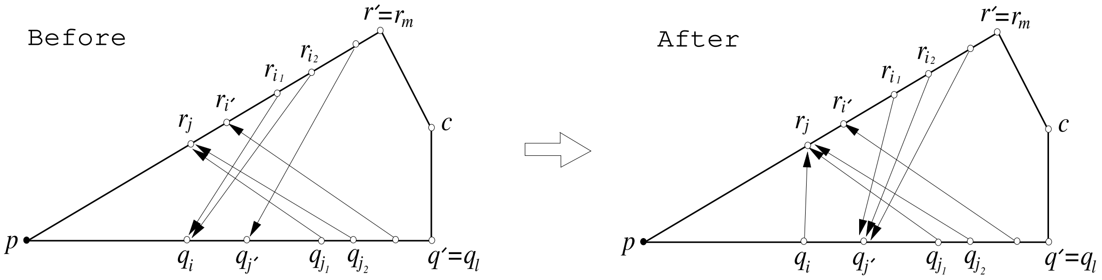

Figure 8.2: Directed edges shown have orientations enforced by the choice of $(i+j, i_1+j_1)$ and I-4. Since neither r_{i_1} nor q_{j_1} is a head of other directed edge, this iteration removed the crossing between $r_{i_1}\vec{q}_i$ and $q_{j_1}\vec{r}_j$ by moving all edges with head at q_i to $q_{j'}$, and adding $q_i\vec{r}_j$.

crossings due to line segments each joining a pair of Steiner vertices. The next lemma is useful in the verification of the invariants after each iteration of Case A and Case B.

Lemma 8.1 For $\triangle pq_ir_{i'}$ with $\angle p < \beta$, at least one of the two exterior angles at q_i and $r_{i'}$ is good. $\square$

Lemma 8.2 Invariants I-1 to I-5 are maintained by each iteration of Case A.

Proof. Refer to Figure 8.1 for the relative positions of vertices mentioned in the following discussion. Let us first show that each stated edge when added satisfies I-1. First, consider q_ir_j. Clearly $\angle q_i$ and $\angle r_j$ inside $\triangle pq_ir_j$ are good as they are parts of two good angles at tails of $q_i\vec{r}_{i_1}$ and $r_j\vec{q}_{j_1}$. As $\angle p < \beta$, Lemma 8.1 implies that one of the two exterior angles of $\triangle pq_ir_j$ at q_i or r_j is good, and we are done for q_ir_j. We can now assume, by symmetry, for the remaining proof that $q_i\vec{r}_j$ is added. We can check similarly that each addition of directed edges $q_{j'}\vec{r}_j$ (obtained from $q_{j'}\vec{r}_{i_1}$), $r_{i_1}\vec{q}_{j_2}$, or $q_{j_1}\vec{r}_{i_1}$ satisfies I-1. This completes the verification for I-1.

The moving of heads of $q_{j'}\vec{r}_{i_1}$ maintains I-4 because $\angle r_j q_{j'} q' \leq \alpha$ for some j' implies that it is also true for smaller j'. (So this moving can be carried out in increasing order of the indices of tails.) By construction, $r_{i_1}\vec{q}_{j_2}$ when added still maintains I-4. Next, notice that if $\angle q_{j_1} r_{i_1} r_m$ is bad, there cannot exist $q_{j'}\vec{r}_{i_2}$ with $i < j' \leq j_1$ and $i_1 < i_2$ because $q_i\vec{r}_{i_1}$ satisfies I-5 before the iteration. So, adding $q_{j_1}\vec{r}_{i_1}$ does not violate I-4 (and I-3), and the invariant is maintained. Invariants I-2 and I-3 are clearly valid.

In the following verification of I-5, we mention only one of the six angles involved since the remaining five can easily be checked to be good. Consider $q_i\vec{r}_j$ with other $q_f\vec{r}_{f'}$ where $j \neq f'$ (and $i \neq f$). If $i > f$, then $\angle q_i r_{f'} r_j = \angle q_i r_{f'} r_{i_1}$, which is good because $q_f\vec{r}_{f'}$ with $q_i\vec{r}_{i_1}$ satisfies I-5. Next, if $i < f < j_1$ and $f' = i_1$, then $\angle r_j q_f q' > \alpha$ implies that $\angle q_f r_j r_{i_1} < \alpha$ (Lemma 8.1). And, if $i < f < j_1$ and $f' > i_1$, $\angle q_f r_j r_{f'} < q_f r_{i_1} r_{f'} \leq \alpha$ by the fact that $q_i\vec{r}_{i_1}$ with $q_f\vec{r}_{f'}$ satisfies I-5. Lastly, if $i < j_1 \leq f$, then $\angle q_f r_j r_{f'} \leq \alpha$ because of $r_j\vec{q}_{j_1}$. Notice that the argument simply applies I-5 on existing edges sharing endpoints with $q_i\vec{r}_j$ to show that $q_i\vec{r}_j$ satisfies I-5. The same approach applies to other newly added edges $q_{j'}\vec{r}_j$, $r_{i_1}\vec{q}_{j_2}$ and $q_{j_1}\vec{r}_{i_1}$. This completes the proof. ◻

Lemma 8.3 Invariants I-1 to I-5 are maintained by each iteration of Case B.

Proof. Refer to Figure 8.2 for the relative positions of vertices mentioned in the following discussion. Let us begin with I-1. First we show that all the eight angles at the endpoints of $q_i r_j$ and $r_{i_1} q_{j_1}$ are good, so directed edges based on them always satisfy I-1. The two exterior angles of $\triangle p q_{j_1} r_{i_1}$ at r_{i_1} and q_{j_1} are good because they are only parts of some good angles at tails of $r_{i_1}\vec{q}_i$ and $q_{j_1}\vec{r}_j$. Next looking at $\triangle q_i r_{i_1} r_j$, we have good $\angle r_j$ followed from the good exterior angle at r_{i_1} because of $r_{i_1}\vec{q}_i$. Similarly, we have good $\angle r_j q_i q_{j_1}$. Now with these four good angles of interest as exterior angles of $\triangle p q_i r_j$ and $\triangle p q_{j_1} r_{i_1}$, we can easily deduce that the other four (interior) angles of interest are also good.

Next, we consider each $r_{i_2}\vec{q}_i$ for $i_2 \geq i_1$ that is moved to $r_{i_2}\vec{q}_{j'}$ where $j' = i + 1$. For $r_{i_1}\vec{q}_i$, we can indeed have $r_{i_1}\vec{q}_{j'}$ as it is within the wedge defined by $r_{i_1}\vec{q}_i$ and $r_{i_1}\vec{q}_{j_1}$. Similarly, each $r_{i_2}\vec{q}_{j'}$ for $i_2 \neq i_1$ satisfies I-1. This completes verifying I-1.

Invariants I-2, I-3 and I-4 are valid by construction. We next verify I-5, as in the previous lemma, by checking one of the six angles involved. Since all the four angles at the endpoints of $q_i r_j$ are good, $q_i\vec{r}_j$ or $r_j\vec{q}_i$ can never be involved to invalidate I-5. Next, consider $r_{i_2}\vec{q}_{j'}$, for $i_2 \geq i_1$, with other $r_g\vec{q}_{g'}$ where $j' \neq g'$. If $g' = i$, we have $\angle r_{i_2} q_i q_{j'} < \angle r_j q_i q_{j'} \leq \alpha$. Now suppose $g \neq i$. Since $r_{i_2}\vec{q}_{j'}$ has the same tail as $r_{i_2}\vec{q}_i$, I-5 is maintained when $i_2 > g$ because of the fact that $r_{i_2}\vec{q}_i$ with $r_g\vec{q}_{g'}$ satisfies I-5. Lastly, for $i_2 < g$, we have $\angle r_g q_{j'} q_{g'} < \angle r_j q_{j'} q_{g'} \leq (\pi - \angle r_j q_{j_1} p) \leq \alpha$. This completes the proof. ◻

9 Triangulating Quadrilaterals

For each non-fence $pq'cr'$, any directed edges in its interior do not cross. Thus, they subdivide $pq'cr'$ into regions of three, four, or five sides. All regions except possibly for the one containing c have *good triangulations*, i.e. triangulations with no bad angles. We must nevertheless handle the region containing c first by merging it with nearby regions. Hence, some directed edges are just guides in the construction and may not be in the final triangulation.

Step 6. Triangulating Quadrilaterals.

Fences and quadrilaterals without subdivision points can easily be triangulated with good triangles; refer to Lemma 4.1. We next consider each non-fence $pq'cr'$ with $\angle q' = \angle r' = \frac{\pi}{2}$. Label its faces from c to p by $R_1, R_2, \ldots, R_{m'}$. Starting with $\mathcal{R} = \emptyset$ and $k = 0$, we then repeat incrementing k and including R_k into $\mathcal{R}$ until we can produce a good triangulation for $\mathcal{R} = \cup_{i=1}^{k} R_i$. Then for the remaining regions R_j, for $j = k+1, k+2 \ldots, m'$, of three or four sides, we triangulate each in a straightforward way.

The following two lemmas show that Step 6 can indeed produce a good triangulation for $pq'cr'$. Let the vertices along pq' be sorted as $p = q_0, q_1, \ldots, q_l = q'$ and along pr' as $p = r_0, r_1, \ldots, r_m = r'$.

Lemma 9.1 Each region R_k for $k = 2, 3, \ldots, m'$ has a good triangulation.

Proof. If R_k is of three sides, then it is pq_1r_1, $q_iq_{i+1}r_j$, or $r_jr_{j+1}q_i$. With the directed edge based on q_1r_1, $\triangle pq_1r_1$ is clearly good. And, $\triangle q_iq_{i+1}r_j$ is good because $\vec{q_ir_j}$ or $\vec{q_{i+1}r_j}$ exists by I-3 and $\angle q_i$ or $\angle q_{i+1}$ is thus between β and α. Similarly, $\triangle r_jr_{j+1}q_i$ is good. Now, if R_k is of four sides, say $q_iq_{i+1}r_{j+1}r_j$. Then it is supported by two directed edges based on q_ir_j and $q_{i+1}r_{j+1}$. If they are of opposite orientations, then $q_iq_{i+1}r_{j+1}r_j$ has two opposite angles each measuring between β and α. So we have a good triangulation for the region by adding the diagonal not incident to these good angles. If they are of the same orientation, a good triangulation is ensured by I-5. We are done as R_k has either three or four sides by I-2. ◻

Lemma 9.2 There exists a $k \leq m'$ such that $\mathcal{R} = \cup_{i=1}^{k} R_i$ has a good triangulation.

Proof. We consider successively larger $\mathcal{R}$ beginning with $\mathcal{R} = R_1$, then $\mathcal{R} = R_1 \cup R_2$ and so on. For each one, we test $\mathcal{R}$ for a good triangulation by joining c with Steiner vertices on the boundary of $\mathcal{R}$ so as to subdivide $\angle q_lcr_m$ further. As we will see, a good $\mathcal{R}$ is found once $\angle q_lcr_m$ can be subdivided into good angles. The following are the details.

Initially, $\mathcal{R} = R_1$ which is (1) $q_l c r_m$, (2) $q_{l-1} q_l c r_m$, (3) $q_l c r_m r_{m-1}$, or (4) $q_{l-1} q_l c r_m r_{m-1}$. For (1), we just go on to enlarge $\mathcal{R}$ since $\angle c$ is bad. For (2) and (3), we add cq_{l-1} and cr_{m-1}, respectively, to subdivide $\angle c$. We are done if we have only good angles around c. For (4), if $|cq_{l-1}| \le |cr_{m-1}|$ we add cq_{l-1}, else add cr_{m-1}, resulting in a region yet to be triangulated. If $\angle c$ of this region is good, we are done by completing the triangulation with the diagonal not incident to c; otherwise, we use the other diagonal to further subdivide $\angle c$ and are done if all angles around c are good.

To enlarge $\mathcal{R}$, we remove the directed edge supporting the current $\mathcal{R}$. In general, we have just removed an edge say $q_{i+1} r_{j+1}$ to obtain a part of $\mathcal{R}$ which has yet to be triangulated. This is $q_i q_{i+1} c r_{j+1}$, $q_{i+1} c r_{j+1} r_j$, or $q_i q_{i+1} c r_{j+1} r_j$, which is handled analogously as in (2) to (4).

It is easy to see that the process terminates, producing only good angles around c. We next verify that only good triangles are produced. Suppose the process terminates with (2) creating $\triangle q_i c r_{j+1}$ (and $\triangle q_i q_{i+1} c$) from $q_i q_{i+1} c r_{j+1}$. Consider $q_i q_l c r_{j+1}$. Since $\angle q_l c r_{j+1} > \alpha$ and $\angle q_l = \frac{\pi}{2}$, both $\angle q_i$ and $\angle r_{j+1}$ are good. So $\triangle q_i c r_{j+1}$ is good. Symmetrically, the triangle created last due to (3) is also good. For (4), the process terminates either similar to (2) and (3) or with a region, say (without loss of generality) $q_i c r_{j+1} r_j$, subdivided into $\triangle q_i c r_{j+1}$ and $\triangle q_i r_{j+1} r_j$. Note that $\triangle q_i c r_{j+1}$ is good as $q_i q_l c r_{j+1}$ has $\angle c > \alpha$ and $\angle q_l = \frac{\pi}{2}$. And, $\triangle q_i r_{j+1} r_j$ is good as $\angle r_j$ is good by our choice of cq_i to triangulate $q_i q_{i+1} c r_{j+1} r_j$, and $\angle q_i$ and $\angle r_{j+1}$ are good by the existence of directed edges based on $q_i r_j$ and $q_{i+1} r_{j+1}$, respectively.

The rest of the triangles resulting from the above are of the form $cq_{i'} q_{i'+1}$ or the symmetrical form $cr_{j'} r_{j'+1}$. For $\triangle c q_{i'} q_{i'+1}$, $\angle q_{i'}$ and $\angle c$ are clearly good as they are acute. And, $\angle q_{i'+1}$ is good from the following: $\angle q_{i'+1} c r_m > \alpha$ because a directed edge incident to $q_{i'+1}$ was removed to enlarge $\mathcal{R}$; $\angle q_{i'+1} c q_l < \pi - \angle q_{i'+1} c r_m < \pi - \alpha$ and $\angle c q_{i'+1} q_{i'} = \angle q_{i'+1} c q_l + \angle q_{i'+1} q_l c < \frac{3}{2}\pi - \alpha = \alpha$. $\square$

Till here, we have established the following theorem:

Theorem 9.3 Triangulating a plane geometric graph $\mathcal{G} = (S, E)$ of $|S| = n$ vertices and $|E| = \mathrm{O}(n)$ edges using angles no larger than $\frac{3}{4}\pi$ requires $\mathrm{O}(n^2)$ Steiner vertices. $\square$

10 Implementing Construction

In this section, we describe an efficient algorithm to implement the above constructive proof. Also, we discuss ways to avoid some redundant Steiner vertices, and extend the construction to a better angle bound of $\alpha = \frac{11}{15}\pi$. Let us assume that each point coordinate can be stored in a constant amount of storage and that basic geometric operations, such as projecting a point onto a line can be carried out in a constant amount of time.

Theorem 10.1 Triangulating a plane geometric graph $\mathcal{G} = (S, E)$ of $|S| = n$ vertices and $|E| = \mathrm{O}(n)$ edges using angles no larger than $\frac{3}{4}\pi$ requires $\mathrm{O}(n^2)$ storage and $\mathrm{O}(n^2 \log n)$ time.

Proof. From $\mathcal{G}$, we can first compute in $\mathrm{O}(n^2 \log n)$ time the triangulation $\mathcal{T}$ that minimizes its maximum angle over all triangulations of $\mathcal{G}$ [11] and store it in a quad-edge data structure of $\mathrm{O}(n)$ storage [13]. If the maximum angle of $\mathcal{T}$ is no larger than $\frac{3}{4}\pi$, then we are done. Otherwise, we proceed to refine $\mathcal{T}$ with the construction given in Sections 4 to 9. To perform those steps efficiently, we maintain for each edge of $\mathcal{T}$ a sorted list of $\mathrm{O}(n)$ vertices of propagation paths on the edge. And, we link up vertices with pointers that act as directed edges of propagation paths. Also, we keep some general information about each quadrilateral edge (such as its type as a fence or non-fence, its good and bad segments, etc.) and about each vertex (such as its type as a dead-end or an endpoint of a spoke, its outgoing pointer, etc.). All in all, these structures requires $\mathrm{O}(n^2)$ storage.

With the above, Step 1 to Step 3 are straightforward. The splitting of an edge of $\mathcal{T}$ into edges of quadrilaterals, and the subsequent walking from one quadrilateral to an adjacent one can each be done in constant time with the quad-edge data structure. The checking of vertices within a maw, and the insertion of a vertex to an edge of $\mathcal{T}$ (or a quadrilateral) can be done in logarithmic time. Therefore, Step 1 to Step 3 are bounded by $\mathrm{O}(n^2 \log n)$ time. Step 4 is simple by first separating directed edges of opposite orientations into two sets and then working on each set in say increasing order of neighboring pairs. This step thus takes $\mathrm{O}(n^2)$ time.

Next, Step 5 is slightly involved. Let us discuss it with notation introduced in Section 8 and with reference to non-fence $pq'cr'$ of $\mathrm{O}(n)$ Steiner vertices. Note first that the sum $i+j+i_1+j_1$ (from the lexicographically smallest ordered pair) is increasing in successive iterations because of I-4. Thus, the bound of $\mathrm{O}(n)$ on the sum also bounds the number of iterations. Note second that the next available index along pq' (or pr') in Case A is also increasing in successive iterations. The reason is that each vertex v that we passed in a search for the next available index remains unavailable because of blocking by other edges, or v becomes unblocked by the removal of directed edges but is now out of range of the searching. Thus, a careful implementation of Step 5 needs only $\mathrm{O}(n)$ time in locating vertices along pq' (or pr') by marching across pq' and pr' once in some coherent way. With these notes, we can implement Step 5 to run in $\mathrm{O}(n)$ time by maintaining for each Steiner vertex a list of directed edges with heads at the vertex. Each list stores elements in increasing order of indices of tails, and is a queue that allows a removal of an element or an appending of another queue in constant time. So, crossings of directed edges inside all quadrilaterals can be removed in $\mathrm{O}(n^2)$ time.

Lastly, Step 6 is straightforward and runs in time linear to the number of vertices in the resulting triangulation. $\square$

Reducing Steiner Vertices. The above construction generates considerably many Steiner vertices, though it achieves the worst case optimal bound within a constant factor. We see in the following, a modified construction that avoids many redundant Steiner vertices. Let us start with Step 1. Instead of subdividing each $\triangle pqr$, with largest angle at q, into three quadrilaterals, we subdivide it into two or four triangles as follows. Also note that a fence now has three edges (with only one is a spoke).

Case a. $\angle q > \alpha$.
Subdivide $\triangle pqr$ into $\triangle pqq'$ and $\triangle rqq'$ where $q' \in pr$ and $\angle pqq' = \angle rqq'$. We treat both triangles as non-fences and qq' as a spoke.

Case b. $2\beta \leq \angle q \leq \alpha$.
Same as Case a. But, mark $\triangle pqq'$ as a fence if $\angle p \geq \beta$ else as a non-fence, and similarly mark $\triangle rqq'$ according to $\angle r$.

Case c. $\angle q < 2\beta$.
Subdivide $\triangle pqr$ into $\triangle pr'q'$, $\triangle qp'r'$, $\triangle rq'p'$ and $\triangle p'q'r'$ where p', q', r' are the three perpendicular projections of the center of inscribing circle of $\triangle pqr$ onto qr, rp, pq, respectively. Also, treat edges $p'q', q'r'$ and $r'p'$ as spokes, and mark each of the triangles (excluding $\triangle p'q'r'$) as a fence if its angle opposite its spoke is no less than β, else as a non-fence.

With this subdivision, we then plan and set up traps in Step 2a and Step 2b. Then, Step 3 generates propagation paths selectively for dead-ends and endpoints of spokes—so endpoints of spokes are just like dead-ends that may or may not be Steiner vertices in the final construction. A dead-end as before originates a propagation path if it terminates some propagation paths, whereas an endpoint of a spoke originates a propagation path in the following situations. A q' of Case a always originates a propagation path. A q' of Case b and p', q', r' of Case c each originates a propagation path if it terminates some other propagation paths, or if the two smaller triangles in $\triangle pqr$ sharing this endpoint have Steiner vertices on their boundaries. (Other refined rules in selecting vertices for propagation are omitted in this paper for simplicity in the presentation.)

Step 4 as before removes unnecessary subpaths. Note that each spoke in a triangle can be considered as a directed edge in either orientation since all the four angles defined by its two endpoints with edges of the triangle are good. Spokes are not involved in merging. Next, Step 5 remains the same. Finally, Step 6 works in a similar way as before: for a non-fence, we just use the straightforward method in Lemma 9.1 (treating each spoke as a directed edge of either orientation); otherwise, we have a fence that can be solved with Lemma 4.1. It is easy to check that $\triangle p'q'r'$ of Case c is good; each of its angle is just a half of the sum of two angles of $\triangle pqr$. However, there is one minor problem—q' in Case b or p', q', r' in Case c which are used in Step 6 may not be Steiner vertices in the construction as mentioned. Nevertheless, a simple shifting of edges can remedy the problem in all cases; see Figure 10.1.

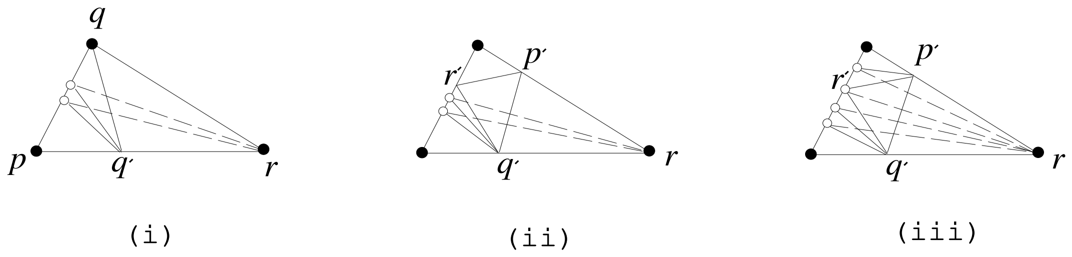

Figure 10.1: Vertex q' of Case b (i), p', q', r' of Case c (ii), and p', q' of Case c (iii) are not Steiner vertices in the construction. Edges incident to these vertices are shifted to r as shown in dashed line segments. All new angles that resulted from the shifting can be checked to be good using the simple extension of Lemma 8.1 to $\angle p < 2\beta$.

Reducing Angle Bound. The angle bound of $\alpha = \frac{3}{4}\pi$ can slightly be improved to $\alpha = \frac{11}{15}\pi$ (i.e. 132^o), but with a larger constant in the quadratic bound on vertex size. In this case, $\beta = \frac{4}{15}\pi$ (i.e. 48^o). All results developed starting from Section 4 are valid except for Lemma 4.1, Lemma 8.1 (with $\angle p < 2\beta$) and Lemma 9.2. We need not worry about Lemma 9.2 since it is no longer relevant to the above modified construction. As for the other two lemmas, problem arises when $\pi - 2\beta < \angle p < 2\beta$. One way to resolve this is to perform the following before Step 1: for each $\triangle pqr$ of $\mathcal{T}$ with largest angle between $\pi - 2\beta$ and 2β, we subdivide it by a new vertex s inside pqr into triangles with angles at s equal to $\frac{2}{3}\pi$. The existence of such a subdivision can be verified as follows. Let $\angle q \geq \angle p \geq \angle r$. We can identify an interior point s' on qr such that $\angle ps'r = \frac{2}{3}\pi$ since $\angle p + \angle q > \frac{2}{3}\pi$. Hence, we can draw a circle through p, s', r whose arc inside $\triangle pqr$ represents the loci of points forming angles of $\frac{2}{3}\pi$ with pr. Note that $\angle qtr$ is π when $t = s'$, and decreases as t moves along the arc towards the intersection point s'' of the arc with pq. Since $\angle qtr < \frac{2}{3}\pi$ when $t = s''$ as $\angle q > \frac{\pi}{3}$, some point s along the arc is such that $\angle qsr = \frac{2}{3}\pi$.

Corollary 10.2 Triangulating a plane geometric graph $\mathcal{G} = (S, E)$ of $|S| = n$ vertices and $|E| = \mathrm{O}(n)$ edges using angles no larger than $\frac{11}{15}\pi$ requires $\mathrm{O}(n^2)$ storage and $\mathrm{O}(n^2 \log n)$ time. $\square$

11 Concluding Remarks

This paper shows that there exists for any plane geometric graph a conforming triangulation with a quadratic bound on its vertex set and $\frac{3}{4}\pi$ bound on its angles. It is possible to extend the result to a slightly better bound on angles, but with a larger constant factor in the quadratic bound on the vertex set. The paper mentions such an improvement for angle bound of $\frac{11}{15}\pi$. On the whole, it improves the result of Mitchell [17] by reducing a logarithmic factor on the bound on the vertex set and by reducing at least $\frac{\pi}{8}$ on the bound on angles. The new bound on the vertex set is asymptotically optimal and can be computed in slightly more than quadratic time. The computation is simple and practical, without introducing many unnecessary vertices. On the other hand, it remains open whether the bound on angles can be reduced further.

The main idea of the paper is on the control of the lengths of propagation paths using fences and traps. There is a similar idea in Edelsbrunner and Tan [10] where the corresponding notion is termed walls. The current paper is, nevertheless, much more complex and has a number of new strategies to address issues on the number of fences, traps and the crossings of propagation paths. It remains interesting to see whether some of these ideas and strategies can be applied to other triangulation problems.

Acknowledgement

I would like to thank Scott Mitchell for providing details of his paper [17]. And, I am grateful to one anonymous referee for many suggestions on improving the presentation of the paper.

References

[1] I. Babuška and A. K. Aziz. On the angle condition in the finite element method. *SIAM J. Numer. Anal.* **13** (1976), 214–226.

[2] R. E. Barnhill and F. F. Little. Three and four-dimensional surfaces. *Rocky Mountain J. Math.* **14** (1984), 77–102.

[3] M. Bern and D. Eppstein. Polynomial-size nonobtuse triangulation of polygons. *Intl. J. Comput. Geom. Appl.* **2(3)** (1992), 241–255.

[4] M. Bern and D. Eppstein. Mesh Generation and Optimal Triangulation. In *Computing in Euclidean Geometry*, 2nd Edition, D. Z. Du and F. K. Hwang, eds., World Scientific, Singapore, 1995, 47–123.

[5] M. Bern, S. A. Mitchell and J. Ruppert. Linear-size non-obtuse triangulation of polygons. *In* "Proc. 10th Ann. Sympos. Comput. Geom., 1994", 221–230.

[6] J. D. Boissonnat. Shape reconstruction from planar cross sections. *Computer Vision, Graphics, and Image Process.* **44** (1988), 1–29.

[7] J. Cavendish. Automatic triangulation of arbitrary planar domains for the finite element method. *Intl. J. Numer. Methods Engrg.* **8** (1974), 679–696.

[8] T. K. Dey. Decompositions of Polyhedra in Three Dimensions. Techn. Rep. CSD-TR-91-056, Ph.D. Thesis, Comput. Sci. Dept., Purdue Univ., IN, 1991.

[9] H. Edelsbrunner and N. Shah. Incremental topological flipping works for regular triangulations. *In* "Proc. 8th Ann. Sympos. Comput. Geom., 1992", 43–52.

[10] H. Edelsbrunner and T. S. Tan. An upper bound for conforming Delaunay triangulations. *Discrete Comput. Geom.* **10** (1993), 197–213.

[11] H. Edelsbrunner, T. S. Tan, and R. Waupotitsch. An $O(n^2 \log n)$ time algorithm for the minmax angle triangulation. *SIAM J. Sci. Statist. Comput.* **13** (1992), 994–1008.

[12] J. A. Gregory. Error bounds for linear interpolation on triangles. *The Mathematics of Finite Element and Applications II*, J. R. Whiteman, ed., Academic Press, NY, 1975, 163–170.

[13] L. J. Guibas and J. Stolfi. Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams. *ACM Trans. Graphics* **4** (1985), 74–123.

[14] C. L. Lawson. Software for C^1 surface interpolation. In *Math. Software III*, J. R. Rice, ed., Academic Press, 1977, 161–194.

[15] E. A. Melissaratos and D. L. Souvaine. Coping with inconsistencies: a new approach to produce quality triangulations of polygonal domains with holes. *In* "Proc. 8th Ann. Sympos. Comput. Geom., 1992", 202–211.

[16] S. A. Mitchell. Mesh Generation with Provable Quality Bounds. Techn. Rep. CS-TR-92-1327, Ph.D. Thesis, Comput. Sci. Dept., Cornell Univ., NY, 1993.

[17] S. A. Mitchell. Refining a triangulation of a planar straight-line graph to eliminate large angles. *In* "Proc. 34th Ann. IEEE Sympos. Found. Comput. Sci., 1993", 583–591.

[18] S. A. Mitchell. Finding a covering triangulation whose maximum angle is provably small. *In* "Proc. 17th Ann. Computer Science Conference, ACSC-17, 1994", *Australian Comp. Science Comm.* **16** (1994), 55–64.

[19] L. Palios. Decomposition Problems in Computational Geometry. Techn. Rep. CS-TR-368-92, Ph.D. Thesis, Comput. Sci. Dept., Princeton Univ., NJ, 1992.

[20] J. Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. *In* "Proc. 4th ACM-SIAM Sympos. Discrete Algorithms, 1993", 83–92.

[21] T. S. Tan. Optimal Two-Dimensional Triangulations. Techn. Rep. UIUCDCS-R-92-1783, Ph.D. Thesis, Comput. Sci. Dept., Univ. Illinois, Urbana, IL, 1993.

A Analyzing Step 3

We prove again some important results of Mitchell [17]; our new proofs are necessary because those results are applied here to a more general setting. For a horn of s, its *center path* is the sequence of line segments where each connects the midpoints of the maws of two consecutive stages, starting at s. Each midpoint mentioned is also termed a *center path point*. The length of a line segment is the distance of its two endpoints, and the *length of a center path* is the sum of the lengths of its line segments. And, the *maw width* of a maw is the distance between its two endpoints. For our following discussion, horns refer to those that do not intersect endpoints of spokes; otherwise our analysis of Step 3 is trivial.

Lemma A.1 For a point s on an edge of $\mathcal{T}$, the maw width of its horn at some stage is greater than the length of the center path up till that stage.

Proof. It suffices to show that the increase in the maw width from one stage to the next is no less than the increase in the length of the center path. Referring to Figure A.1, we show $|a'b'| - |ab| \geq |ee'|$. Note that the case of stage 0 is so that $a = b = e$. We have $|ab| + |de| < |a'b''|$ from simple

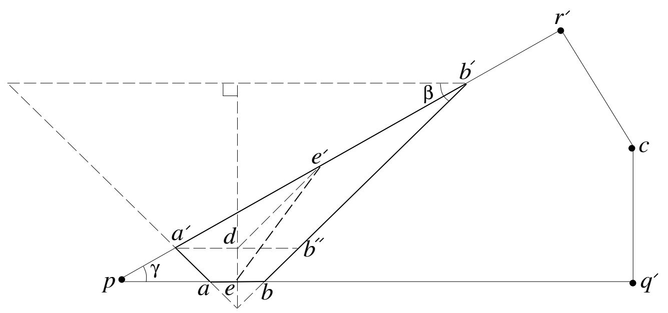

Figure A.1: The part of the horn of s inside $pq'cr'$ is $abb'a'$. That is, the horn of s arrives at $ab \subset pq'$ and then extends to $a'b' \subset pr'$. The midpoint of ab and $a'b'$ are e and e', respectively. Point $b'' \in bb'$ is so that $a'b''$ is parallel to ab, and $d \in a'b''$ is so that de is perpendicular to ab. Notice that $e'd$ is parallel to $b'b''$.

observation, and $|de| + |de'| > |ee'|$ from the triangle inequality. So, we just need to show that $|a'b'| - |a'b''| \geq |de'|$. The law of sines for $\triangle a'de'$ gives $|de'| = |a'e'| \sin\gamma / \sin\alpha = \frac{1}{2}|a'b'| \sin\gamma/\sin\alpha$, and for $\triangle a'b''b'$ gives $|a'b''| = |a'b'| \sin(\beta - \gamma)/\sin\alpha$. So,

$$\begin{aligned}
\frac{|a'b'| - |a'b''|}{|de'|} &= \frac{2\sin\alpha}{\sin\gamma}\left(1 - \frac{\sin(\beta-\gamma)}{\sin\alpha}\right) \\
&= \frac{2}{\sin\gamma}(\sin\beta - \sin\beta\cos\gamma + \cos\beta\sin\gamma) \\
&> \frac{2}{\sin\gamma}(\cos\beta\sin\gamma) \\
&= 2\cos\beta \;>\; 1.
\end{aligned}$$

□

Let T be a trap having two distinct dead-ends (case 2 in Section 5) with base xy. For a point $z \in xy$, it is clear that the horn of z into T moves closer in successive stages to both boundary paths of T. This notion of closeness can be formalized from the concept of inverse horn [17]. For our purposes, we define an *inverse horn* as the region between two coherent forward paths (in a given sequence of non-fences) where good angles of β at tails of their directed edges are inside the region. The notions of maw, maw width, center path, and boundary paths for horn are extended analogously to inverse horn. As mentioned, the maw width of an inverse horn decreases (rather than increases) from one stage to the next. When the maw width is negative, the two boundary paths of the inverse horn have crossed.

Lemma A.2 Let us consider the inverse horn defined by coherent forward paths originated at w and z. Initially, $\ell = |wz|$ is the maw width. Then, the maw width ℓ' at some later stage is so that $\ell - \ell' > \frac{4}{5}\ell_c$ where ℓ_c is the length of the center path up till that stage.

Proof. Analogous to the previous lemma, it suffices to show that from one stage to the next, the decrease in maw width is at least $\frac{4}{5}$ the increase in the center path length. Referring to Figure A.2, we show in the following $|ab| - |a'b'| > \frac{4}{5}|ee'|$. The law of sines for $\triangle a'de'$ gives $|de'| = \frac{1}{2}|a'b'|\sin\gamma/\sin\beta$, and for $\triangle a'b''b'$ gives $|a'b''| = |a'b'|\sin(\pi-\beta-\gamma)/\sin\beta = |a'b'|\sin(\beta+\gamma)/\sin\beta$. So,

$$\frac{|a'b''| - |a'b'|}{|de'|} = \frac{2}{\sin\gamma}\left(\sin(\beta+\gamma) - \sin\beta\right),$$

and, for $0 < \gamma < \beta = \frac{\pi}{4}$, $\frac{|a'b''|-|a'b'|}{|de'|}$ is decreasing in γ. This means

$$\frac{|a'b''| - |a'b'|}{|de'|} > \frac{4}{5},$$

so the claim is proved because $|ab| > |a'b''| + |de|$ and $|de| + |de'| > |ee'|$. $\square$

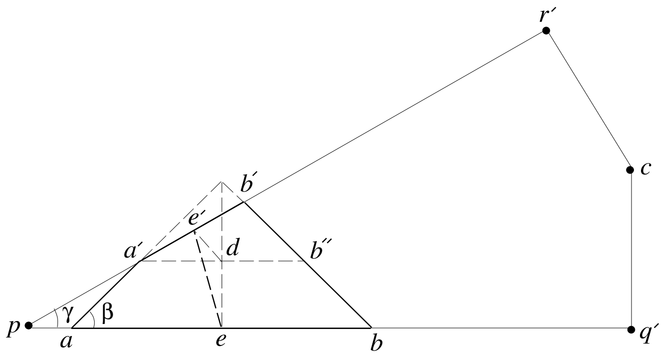

Figure A.2: Region $abb'a'$ is the part of the inverse horn inside $pq'cr'$. The inverse horn arrives at $ab \subset pq'$ and then extends to $a'b' \subset pr'$. The midpoint of ab and $a'b'$ are e and e', respectively. Point $b'' \in bb'$ is so that $a'b''$ is parallel to ab, and $d \in a'b''$ is so that de is perpendicular to ab. Notice that $e'd$ is parallel to $b'b''$.

With the above lemmas, we next complete the argument of Lemma 6.3 using the next result, which implies the existence of a properly-terminating propagation path with origin z and length $O(n)$.

Lemma A.3 Let z be a point on the base of a trap T having two distinct dead-ends (case 2 in Section 5). Then, the horn of z inside T either self-intersects or intersects at least one boundary path of T at a stage no larger than the length of the boundary paths of T.

Proof. Let pq be the edge intersected three times by each boundary path of T (so pq contains the dead-ends of T). For the proof, it suffices to consider the horn $\mathcal{H}_z$ of z_1 inside T where z_1 is some point on the first intersection of the horn of z and pq. Assuming $\mathcal{H}_z$ does not self-intersect inside T, we prove that $\mathcal{H}_z$ intersects a boundary path of T. Each boundary path of T with one boundary path of $\mathcal{H}_z$ form an inverse horn. Refer to Figure A.3; the inverse horn $\mathcal{I}$ of interest has initial maw z_1w_1 where w_1 (on a boundary path of T) lies on z_1v and v is some point on the second intersection of $\mathcal{H}_z$ with the line through pq. (In Figure A.3, v is z_2.)

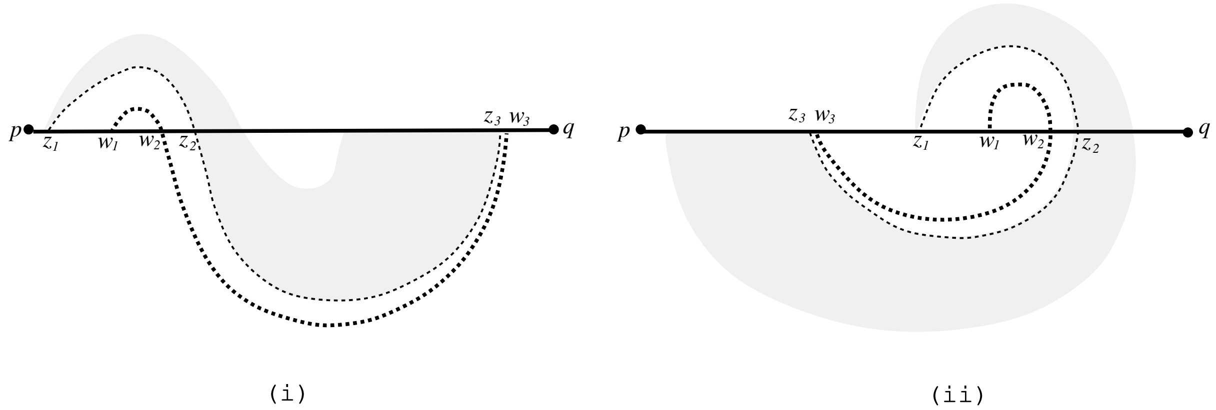

Figure A.3: The edge pq intersects one boundary path of $\mathcal{H}_z$ at z_1, z_2, z_3 and one boundary path of T at w_1, w_2, w_3. The figure shows two representative situations of the inverse horn $\mathcal{I}$: z_2 and z_3 are on the same side of z_1 in (i), but on different sides in (ii). Note that w_2 may be a point in z_1w_1 rather than w_1z_2, and w_3 of (ii) may lie in z_1w_1 rather than z_3z_1.

For Figure A.3(i), the maw widths of $\mathcal{H}_z$ at z_2 and at z_3 are both greater than $|z_1z_2|$ by Lemma A.1. So, the center path length of $\mathcal{I}$ from midpoint of w_2z_2 to midpoint of w_3z_3 is at least $2|z_1z_2|$. Since $\mathcal{H}_z$ does not self-intersect, then Lemma A.2 implies that the maw width $\mathcal{I}$ at w_3z_3 is at most $|w_2z_2| - \frac{4}{5}(2|z_1z_2|) < 0$, as required. Besides Figure A.3(i), there are three other situations to consider for z_2 and z_3 on the same side of z_1. First, $\mathcal{H}_z$ may leave pq at z_1 and return to pq at z_2 from below (rather than from above as shown in the figure). In doing so, the center path of $\mathcal{H}_z$ intersects some point v on the line through pq but not on pq. If $z_2 \in pq \cap z_1v$, then the maw of $\mathcal{H}_z$ at z_2 contains z_1 (Lemma A.1) because the center path of $\mathcal{H}_z$ has traveled a distance larger than $|z_1v|$; if $z_2 \notin pq \cap z_1v$, we can use the argument in the next paragraph. Second, $\mathcal{H}_z$ may leave pq at z_2 and then return to pq at z_3 from above, intersecting the line through pq at points not on pq. In this case, the above arguments also apply. Third, if $z_3 \in z_1z_2$, then the maw of $\mathcal{H}_z$ again contains z_1.

For Figure A.3(ii) with z_2 and z_3 on different sides of z_1, the center path point of $\mathcal{I}$ at z_1w_1 must be to the right of that at z_3w_3 since $z_3 \notin z_1z_2$ and the maw width of $\mathcal{I}$ at z_3 is smaller than that at z_1. Thus, the center path of $\mathcal{I}$ traveled a distance of at least $\ell_1 = |z_1w_1|$ to the midpoint of z_3w_3, Lemma A.2 implies that its maw width $\ell_3 = |z_3w_3| < \ell_1 - \frac{4}{5}\ell_1 = \frac{1}{5}\ell_1$. With this bound on ℓ_3, a better bound on the distance traveled is $\ell_1 + (\frac{\ell_1}{2} - \frac{\ell_3}{2}) > \frac{7}{5}\ell_1$, which in turn implies that ℓ_3 is negative (Lemma A.2). Notice that the same argument can also handle the situation similar to Figure A.3(ii) with $q \in w_1w_2$, i.e. z_2 and w_2 are not on pq. $\square$

Note that Lemma A.3 can be proved even if Lemma A.2 is weakened to having $\ell - \ell' > \frac{67}{100}\ell_c$. This just requires a few more rounds of calculation of the bound on the distance traveled by the inverse horn for Figure A.3(ii). This and because Lemma A.1 remains valid as long as $\beta > \frac{\pi}{3}$, we can pick $\alpha = \frac{11}{15}\pi$ (or $\beta = \frac{4}{15}\pi$) to derive $\ell - \ell' > \frac{67}{100}\ell_c$ for Lemma A.2 and, as a result, an improved angle bound.