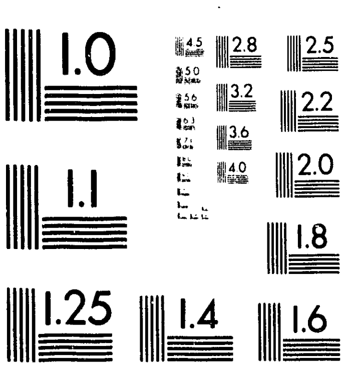
1.0
2.8
2.5
3.2
2.2
3.6
4.0
2.0
1.1
1.8
1.25
1.4
1.6

2/

Refining a Triangulation of a Planar Straight-Line Graph to Eliminate Large Angles

Extended Abstract

Scott A. Mitchell*

May 13, 1993

Abstract

Triangulations without large angles have a number of applications in numerical analysis and computer graphics. In particular, the convergence of a finite element calculation depends on the largest angle of the triangulation. Also, the running time of a finite element calculation is dependent on the triangulation size, so having a triangulation with few Steiner points (added vertices) is also important. Bern, Dobkin and Eppstein [1991] pose as an open problem the existence of an algorithm to triangulate a planar straight-line graph (PSLG) without large angles using a polynomial number of Steiner points.

We solve this problem by showing that any PSLG with v vertices can be triangulated with no angle larger than $7\pi/8$ by adding $O(v^2 \log v)$ Steiner points in $O(v^2 \log^2 v)$ time. We first triangulate the PSLG with an arbitrary constrained triangulation and then *refine* that triangulation by adding additional vertices and edges. Some PSLGs require $\Omega(v^2)$ Steiner points in any triangulation achieving any largest angle bound less than π. Hence the number of Steiner points added by our algorithm is within a $\log v$ factor of worst case optimal.

We note that our refinement algorithm works on arbitrary triangulations: Given any triangulation, we show how to refine it so that no angle is larger than $7\pi/8$. Our construction adds $O(nm+np\log m)$ vertices and runs in time $O((nm+np\log m)\log(m+p))$, where n is the number of edges, m is one plus the number of obtuse angles, and p is one plus the number of holes and interior vertices in the original triangulation. A previously considered problem is refining a constrained triangulation of a simple polygon, where $p = 1$. For this problem we add $O(v^2)$ Steiner points, which is within a constant factor of worst case optimal. The algorithms we present are very practical: For most inputs the number of Steiner points and running time would be considerably smaller than in the worst case.

Submitted to FOCS 93

*Email:samitch@cs.sandia.gov. Applied and Numerical Mathematics Department, Sandia National Laboratories, PO Box 5800, Albuquerque, New Mexico 87185-5800. This work performed at Sandia National Laboratories and was supported by the U.S. Department of Energy under contract DE-AC04-76DP00789, and by the Applied Mathematical Sciences program, U.S. Department of Energy Research.

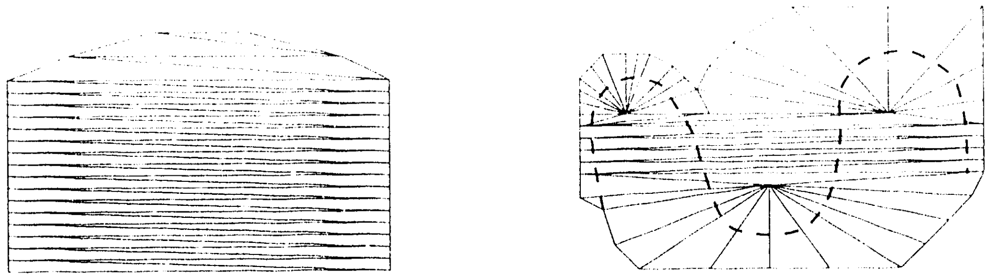

Figure 1: A triangulation refinement may require $\Omega(nm)$ Steiner points to have any constant angle bound (left). In fact, a single Steiner path may require $\Omega(np)$ Steiner points (right).

1 Introduction

1.1 Problem statement and motivation

We are concerned with finding a *conforming* or *Steiner triangulation* of a planar straight-line graph (PSLG). That is, we seek a triangular graph, such that vertices of the input appear as vertices of the output, and edges of the input appear as a union of edges of the output.

Steiner triangulations whose triangles have bounded shape are important for numerical analysis, in particular for a mesh in a finite element method. Babuška and Aziz [1976] shows that the convergence of a finite element method depends on the largest angle of the triangulation. Often one wishes to find a triangulation for a PSLG that is not a polygon. For example, a semiconductor may have two differently doped regions. Hence a description of the semiconductor would include an edge with the interior of the region on both sides. Steiner triangulations without large angles are also of use in functional interpolation and computer graphics (see Barnhill [1983]).

In addition to the shape of the triangles, another important criterion for a triangulation is the number of triangles. For example, in a finite element method calculation the number of triangles directly affects the running time. For many triangulation algorithms the number of triangles produced depends on the input geometry or embedding, and not just the cardinality of the input. Bern, Dobkin and Eppstein [1991] pose as an open problem the existence of an algorithm to triangulate a PSLG without large angles using only a polynomial number of added vertices (called Steiner points). Here "polynomial" is taken to mean polynomial in the input cardinality, independent of the geometry.

1.2 Previous results

For polygonal input there are many results concerning the construction of triangulations without large angles. Bern and Eppstein [1991] shows how to triangulate an arbitrary polygon so that no angle is obtuse by adding $O(n^2)$ Steiner points. Bern, Dobkin and Eppstein [1991] shows how to triangulate various types of input polygons with various angle bounds, using between $O(n \log n)$ and $O(n^{1.85})$ Steiner points. It is unknown if the numbers of Steiner points added by these algorithms are worst case optimal, as the only known lower bound for polygons and point sets is $\Omega(n)$. Bern, Eppstein and Gilbert [1990] shows how to triangulate a point set with no obtuse angles using only a linear number of Steiner points, which is worst case optimal. Eppstein [1992] achieves this and simultaneously approximates the minimum weight Steiner triangulation.

Rupert [1993] shows how to triangulate a PSLG so that no angle is smaller than $\pi/9$, and hence no angle is larger than $7\pi/9$. However, any triangulation that achieves no small angles is doomed to use a non-polynomial number of Steiner points, dependent on the input geometry. There are several previous algorithms that achieve similar results (by dissimilar techniques) for polygonal input. See Bern and Eppstein [1992] for a summary.

Edelsbrunner, Tan, and Waupotitsch [1990] shows how to generate a *constrained triangulation* (one where no Steiner points are allowed) of a PSLG such that the maximum angle is minimized. The technique used is edge-insertion, a global strategy that is a generalization of local edge flip. Mitchell and Park [1993] shows how to generate a *covering triangulation* (one where no Steiner points are allowed on the input edges) of a PSLG such that the maximum angle is approximately minimum, using a polynomial number of Steiner points.

1.3 Overview

We consider the problem of triangulating a PSLG so that no angles are large. We solve this by first triangulating the PSLG with an arbitrary constrained triangulation, and then refining that triangulation. Given any triangulation, we show how to *refine* it by adding additional vertices and edges so that no angle is larger than $7\pi/8$. Our construction adds $O(nm + np\log m)$ vertices and runs in time $O((nm + np\log m)\log(m+p))$. We define p to be one plus the number of holes and interior vertices in the original triangulation. That is, p is the number of two-dimensional connected components of the boundary of the region to be triangulated, plus the number of vertices strictly interior to the region to be triangulated. We define n to be the number of edges, and m to be one plus the number of obtuse angles in the original triangulation. By Euler's formula, in any constrained triangulation of a PSLG with v vertices each of p, n and m is $O(v)$. Hence the final PSLG triangulation has $O(v^2\log v)$ vertices and takes $O(v^2\log^2 v)$ time.

Bern and Eppstein [1991] shows how to refine a constrained triangulation of a simple polygon so that no angle is obtuse using $O(n^4)$ Steiner points. They provide a lower bound example, due to Paterson, that illustrates the key concept in our algorithm. The example shows that a triangulation refinement may require $\Omega(n^2)$ (actually $\Omega(nm)$) Steiner points in order to achieve any angle bound less than π. The example consists of a stack of $\Omega(n)$ long and skinny triangles capped by $m = \Omega(n)$ triangles with obtuse angles directed into the stack as in Figure 1 left. Each obtuse angle in the cap requires a Steiner point on the opposite triangle edge in order to refine the triangulation without large angles. This induced Steiner point in turn induces a Steiner point on the next lower edge, etc. If the figure is made sufficiently wide and short, the Steiner points induced for different obtuse angles are far apart and can not interact with one another. Hence each of the $\Omega(n)$ obtuse angle induces $\Omega(n)$ Steiner points, for a total of $\Omega(n^2)$ Steiner points.

Steiner path. The key concept in our algorithm is the fact that if the final triangulation is to have no large angles, adding a Steiner point on one edge of a triangle may induce the addition of a Steiner point on another edge of the triangle. We call a sequence of induced Steiner points a *Steiner path*. Besides being fairly intuitive, the fact that Steiner paths are sometimes necessary can be proved as a direct result of a lemma about constrained triangulations in Edelsbrunner, Tan, and Waupotitsch [1990] (see Section 2.1).

A variation on Paterson's example provides additional motivation for Steiner paths in Section 2. The stack and the cap can be replaced by triangles all having a vertex in common: We build the example of Figure 1 right by using $p = \Omega(n)$ such constructions separated by a middle stack of size $\Omega(n)$. The Steiner path shown in Figure 1 right is required to intersect each edge of the middle stack $\Omega(p)$ times. Hence a triangulation with a single obtuse angle may require $\Omega(n^2)$ (actually $\Omega(np)$) Steiner points in any refinement that achieves an angle bound less than π.

Algorithm. Our algorithm is as follows. Given a PSLG, we triangulate it with an arbitrary constrained triangulation algorithm, such as the minmax angle triangulation of Edelsbrunner, Tan, and Waupotitsch [1990]. Henceforth we consider that triangulation as our input. For each obtuse angle vertex of the input, we subdivide it into two acute angles by adding the altitude from it to the opposite triangle edge. Hence all triangles are non-obtuse (but also non-conformal), which is important for Section 3. These altitudes introduce a collection Steiner points in the interior of triangle edges. Such points are called *non-conformal* because they make the graph non-triangular. These non-conformal points induce Steiner paths. Any fixed strategy of consecutively chosing the Steiner points on a path is doomed to produce a very long path for some input. As we shall see in Section 2 below, for a given desired angle bound there is some flexibility in picking the next Steiner point on a path. We retain this flexibility, and consecutively determine an ever widening region called a *horn*, such that there is some acceptable Steiner path from the initial vertex to every point in the horn. Only later do we chose exactly which path we take from among all those possible inside the horn. This allows us to bound the length of a particular path by $O(np)$.

Eventually each horn will terminate either by intersecting the boundary of the input or a triangle vertex, by intersecting itself in a special way, or by intersecting another horn. In the first case we create a Steiner path to some Steiner vertex on the input boundary or the triangle vertex. In the second case we resolve the large input angle by creating a Steiner path that ends in a loop (see Figure 3 left). In the third case, we can create a Steiner path to an intersection point of the two horns on a triangle edge (see Figure 3 right).

Because of the third case our algorithm is iterative: We may have to create a Steiner path for the intersection point, which we do in the next iteration. The number of Steiner paths m_i we need to introduce at iteration i decreases geometrically, so there are $O(\log m)$ iterations. The collection of paths may intersect an input edge $O(m_i + p)$ times, which is surprisingly close to our bound of $O(p)$ for single path. Hence at each iteration we add $O(n(m_i+p))$ Steiner points, for a total of $O(nm+np\log m)$. As a practical consideration, the constant in this bound is relatively small. In particular, we derive an upper bound of $3nm + 8np\log_{3/4} m + 2n/3 + m$ Steiner points, and even this is not tight. Furthermore, assuming most inputs do not have long sequences of adjacent triangles with very small angles, for most inputs the refinement algorithm would add considerably fewer points (perhaps only $3m$).

When introducing a path, we just introduce vertices, and not edges between consecutive vertices of the path. We do this because edges for two different paths may cross interior to a triangle. We introduce edges to make the graph conformal only after all paths have been created. We resolve the crossings of Steiner path edges in two ways. If two crossing edges have vertices near a small angle vertex of a triangle, we can swap vertices so that the edges do not cross, and add a diagonal to triangulate the resulting quadrilateral. This strategy does not work near the large angle vertices of a triangle, since paths involving all three of the triangle edges may interact. So instead we introduce one vertex inside the triangle near the edge opposite the small angle, and connect it with an edge to each remaining vertex on the triangle boundary.

The remainder of this paper is organized as follows. Section 2 concerns the development of the Steiner paths, and Section 3 describes how to fix the non-conformal input triangles. In Section 4 we present selected open problems. In the appendix, Section 5, we provide the proofs for many of our lemmas and theorems and bound the running time.

2 Introducing Steiner points

2.1 Steiner path motivation

Consider refining a given triangulation so that no angle is greater than some bound. We the following lemma from Edelsbrunner, Tan, and Waupotitsch [1990], where $\mu(\mathcal{T})$ denotes the maximum angle of a triangulation $\mathcal{T}$.

Lemma 1 *Given a vertex set A, in any constrained triangulation $\mathcal{T}$ containing edge $\overline{WV}$, we have $\mu(\mathcal{T}) \geq \max_{S \in A} \angle WSV$.*

For a Steiner triangulation, the edge opposite a large angle of a triangle must be subdivided in order to reduce the bound of this lemma (adding Steiner vertices elsewhere merely increases A). So suppose we add a vertex S to subdivide an edge. Unless S is on the boundary of the region to be triangulated, there may be a triangle edge $\overline{VW}$ that subtends a large angle at S. Hence to reduce the bound of the lemma we need to subdivide this edge as well, etc., inducing a Steiner path.

To gain some intuition about long Steiner paths, we have a sufficient condition on a triangle T such that $\angle VSW$ is not large: $\angle VSW$ is at least the supplement of the smallest angle of T. Hence Steiner paths only continue through triangles with a small angle.

If $\angle VSW$ is large, we wish to find an acceptable placement of S_1 on $\overline{WV}$ in terms of the angles $\angle VSS_1$ and $\angle WSS_1$ such that the lower bound from Lemma 1 is reasonably small: If we place S_1 so that both of these angles are less than $\alpha \geq \pi/2$, then the lower bound from Lemma 1 is at most α. Requiring $\angle VSS_1 = \angle WSS_1 = \pi/2$ leads to a single choice for S_1. This is too restrictive in that it could lead to long Steiner paths (e.g. infinitely long in Figure 3 left). We chose $\alpha = 3\pi/4$, so that there is a range of acceptable placements for S_1.

In fact, regardless of the choice of α, we need a global strategy for placing the S_i that takes into account the entire induced Steiner path: For any local strategy, there is an input that leads to a very long Steiner path.

We do not fix S_1, but instead only consider it to be in the acceptable range, with the freedom to go back and fix its exact location later. Hence we have more freedom in where to place S_2, or we may even discover a position for S_1 where S_2 is unnecessary due to a triangle vertex of T_1. The longer the path is, the more freedom we have in the placement of the last Steiner vertex (see Figure 2). We are able to take advantage of this freedom to prove that paths are only of length np. Recall that for Section 3, the first point S is always

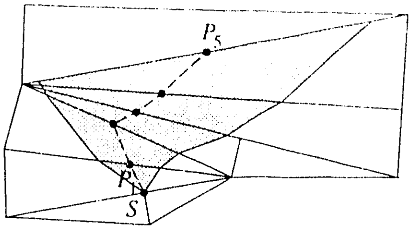

Figure 2: A horn (shaded) and its center path (dashed) terminating on its maw.

on the altitude containing the large angle vertex, and we consider that altitude to be as any other triangle edge. We now formalize.

2.2 Horns

Cone. Consider a Steiner point S on $\overline{WU}$ of $\triangle UVW$. The *cone* at S consists of all points P of $\triangle UVW$ such that $\angle PSW$ and $\angle PSU$ are at most $3\pi/4$. Alternatively, the cone is the intersection of $\triangle UVW$ and the infinite sector at S whose bounding rays make angle $\pi/4$ with $\overline{WU}$. The angle between the bounding rays of the cone is $\pi/2$. See Figure 5.

Maw. The *maw* of a cone is the portion of the cone on the boundary of $\triangle UVW$, excepting S itself. If triangle vertex V is in the maw, then we know that the lower bound from Lemma 1 is at most $3\pi/4$, and there is no need to add any more Steiner points in this path. Otherwise, the maw corresponds to the range of positions for placing S_1 such that the lower bound of Lemma 1 is at most $3\pi/4$. In this case the maw consists of a line segment contained in $\overline{WV}$ or $\overline{UV}$; without loss of generality we assume $\overline{WV}$ throughout this paper.

Each point in the maw defines an acceptable position for S_1, which in turn defines a cone for the next triangle. The union of the maws of these cones defines an acceptable placement for S_2. That is, it defines a range of placements for S_2 such that we can find a placement of S_1 where the cone at S_1 contains S_2, and the cone at S contains S_1. Thus we define the following:

Horn. We iteratively build the *horn* at S as a union of cones. We initialize by setting the horn to be the cone at S. At the next stage of the horn's construction, we add the union of the cones in the next triangle for all points in the maw at the current stage. We define the maw of the horn at a particular stage to be the union of maws for the set of cones just added. See Figure 2.

Center and boundary paths. We define the *center path* of a horn to be the sequence of line segments connecting the midpoints of the maws for consecutive stages, starting at S. See Figure 2. We call the two sequences of segments from the starting vertex of a horn making angle $\pi/4$ with each triangle edge *boundary paths*.

Terminating criteria. We continue the construction of the horn in stages until one of the following occurs (the first three are considered case one in the introduction):

1. The maw contains a triangle vertex. As a heuristic, we also terminate if the maw contains a Steiner path vertex of a previous iteration (see Section 2.5).

2. The maw is on a triangle edge where the next triangle has all angles at least $\pi/8$.

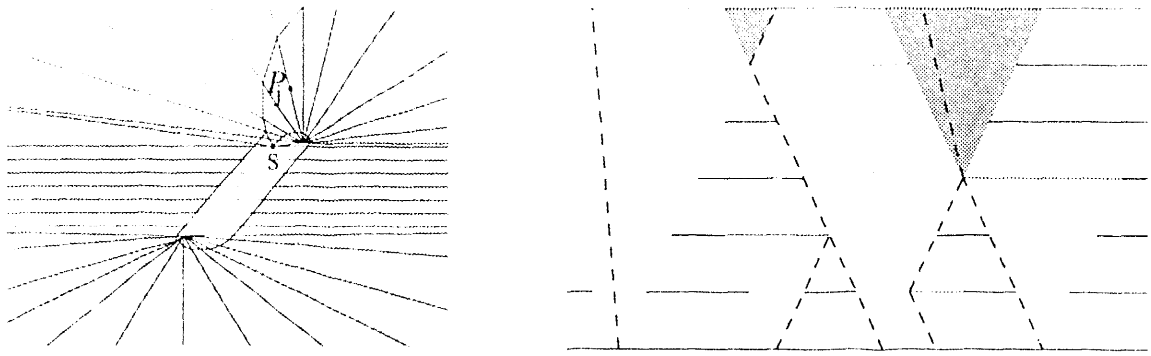

Figure 3: A horn may self intersect (left). The horn shown will self terminate by Item 4 after three more stages, since the horn from center path point P_j will contain P_j. A collection of horns may merge (right). The horns for the next iteration are darkly shaded, and Steiner paths are dashed.

3. The maw is on an edge of the boundary of the region to be triangulated.

4. The maw contains a center path point of a previous stage of the horn, and moreover the horn defined from that center path point contains that center path point (see Figure 3 left).

5. The maw intersects a horn constructed earlier in the current iteration, and either the horns are oriented in the same direction, or in opposite directions and the maw intersects a boundary path or center path of the other horn (see Figure 3 right).

When one of these termination criteria occurs, there is an *acceptable* Steiner path inside the horn. By acceptable, we mean that the lower bound from Lemma 1 is guaranteed to be no more than $3\pi/4$, and the cone from a Steiner path point contains the next point in the path. Given a final Steiner point, it is easy to compute an acceptable path by working from the final Steiner point back to the first point S of the horn. If the horn terminated because of Item 4, then we must form a loop containing the distinguished center path point as well (see Figure 3 left).

In Item 1, the final Steiner point is the vertex contained in the maw. In Item 2 and Item 3 we may pick any point in the maw. In Item 5, we pick the point that is both in the maw and on the boundary path of the other horn. (We may not actually use the corresponding path, depending on whether any later horns in the current iteration terminate on it; see Section 2.5). In Item 4, we pick the center path point P_j that caused the horn to terminate.

2.3 Bounding a single path

In order to bound the number of times horns may intersect a given edge, we need some lemmas about how quickly the maw of a horn grows in relation to its center path length.

Lemma 2 *For the cone at S, the width of the maw is twice the length of the center path.*

In general, the center point of the next cone is not the center point of the next maw, so that the width of the maw is not always twice the center path length. However, we are able to establish a smaller linear bound. See Figure 5 in Section 5.

Lemma 3 *The width of the maw is greater than the length of the center path .*

Lemma 4 *Suppose a horn intersects an edge with maw width M and again at a later stage with maw width M'. Then either $M' > 2M$ or the horn terminates as in Item 4.*

We wish to bound the number of times a center path may cross a given edge. We first show that successive center path points have some ordering along an edge.

Lemma 5 *Suppose a horn intersects an edge E with center path point P_1, and again at a later stage with center path point P_2. If the horn intersect E at a later stage with center point P_3 between P_1 and P_2, then it self terminates as in Item 4.*

Proof. This proof illustrates a general technique we use often: We show that a center path must be long, so that a maw is wide and hence either the center path is far away from some feature or contains that feature.

Consider the horn from P_1 (whether or not P_1 really is the first point of a horn). Let d be the distance from P_1 to P_2. Let M_2 be the maw width at P_2, where from Lemma 3 $M_2 > d$. Let M_3 be the maw width at P_3. From Lemma 4 we have $M_3 > 2M_2$. Hence $M_3 > 2d$, and if P_3 has distance to P_1 less than d, then the maw at P_3 contains P_1 and the horn terminates as in Item 4. ∎

We may use this lemma to analyze the number of intersections a center path may have on an edge if all points of intersection lie on one side of the first center path point. But first we need the following definition.

Hop. The center path between two points P_i and P_{i+1} on an edge E, together with the portion of E between P_i and P_{i+1}, forms the boundary of a compact set in the plane. We call this set a *hop*.

Theorem 1 *Suppose a horn intersects an edge E with center path point P_1. Consider the line L containing E. Then the horn may intersect E at most p times on one side of P_1 before intersecting L on the other side of P_1.*

Proof. By Lemma 5, $P_2, P_3, \ldots$ are consecutive along E, and hence the corresponding hops have disjoint interiors. Each hop contains a vertex of the input in its interior (see Section 5), and hence there can be at most p such hops. ∎

We now consider the case that all of the points of a center path on E to not lie on the same same of the first such point. For this to happen, there needs to be a *reversal.* Intuitively, a reversal is two consecutive hops that travel in opposite directions.

Reversal. A reversal is the center path of a horn from a point P on edge E to a point P_1 on L (or E) to a point P_2 on E, where P_1 and P_2 are on opposite sides of P, and L is the line through E. Also, the center path must not cross E at any point other than P_1 between P and P_2 (otherwise we may find a shorter reversal instead). See Figure 7 in Section 5. Just as for a hop, we say that a reversal contains the input vertices in the region bounded by the center path from P to P_2 and E.

Lemma 6 *The vertices contained by two reversals of a center path are not identical, unless the horn terminates as in Item 4.*

Proof. The proof lies in observation that if hops contains the same vertices and are oriented in the same direction, then they grow closer together as the stages increase. This holds true for two hops for the same Steiner path but different starting stages, and also for two hops for different Steiner paths (used in the next subsection). The rate at which the distance between them decreases can be bounded below in the same way that the rate of increase of the width of a maw can be bounded below. Hence the proof reduces to showing that the "center path" of the region between the two reversals is long, so that the outer reversal must contain the starting center path vertex of the inner reversal. See Section 5. ∎

Theorem 2 *A horn may intersect a given edge at most $O(p)$ times.*

Proof. The hops are partially ordered by containment. Hence there are at most $2p$ unique input vertex sets contained in hops. We enumerate the hops by charging the vertex sets for hops. Using a careful charging scheme, Lemma 6, and Theorem 1, each vertex set gets charged only a constant number of times. See Section 5 for the full proof. ∎

2.4 Bounding the collection of paths

We now consider all of the horns in a given iteration, and how they interact. We seek a bound on the number of times these horns may collectively cross a given edge. We consider two horns with hops oriented in the same direction with respect to E.

Lemma 7 *Suppose a horn H has two consecutive hops M and M_1 (not a reversal), and another horn H' has two consecutive hops M' and M_1', such that M and M' contain the same input vertex set, and similarly for M_1 and M_1'. If neither H nor H' self terminates on M, M', M_1 or M_1' as in Item 4, then M_1 and M_1' intersect as in Item 5.*

Lemma 8 *Suppose a horn H has a reversal R, and another horn H' has a reversal R' with hops containing the same vertex sets as those of R. If neither horn terminates on R or R' as in Item 4, then the two reversals intersect.*

Theorem 3 *The collection of horns may intersect a given edge at most $O(p + m_i)$ times.*

Proof. The proof is very similar to the proof of Theorem 2, and relies on on Theorem 2, Lemma 7 and Lemma 8. See Section 5 for the proof. ∎

2.5 Introducing Steiner paths for the collection

The algorithm for constructing the Steiner paths is iterative. If no horn terminated by Item 5 we could construct the Steiner paths and no more iterations would be required. However, if a horn terminates because of Item 5 then its last Steiner point may induce a path in the next iteration. We have shown above that each iteration will produce only $O(n(m_i + p))$ Steiner path points, where m_i is the number of horns in iteration i. We show below that at each iteration we reduce the number of horns by at least a factor of $3/4$, so the sum of the m_i is $3m$. Furthermore, since the m_i are bounded above by a geometric series, there is only a logarithmic number of iterations in terms of m.

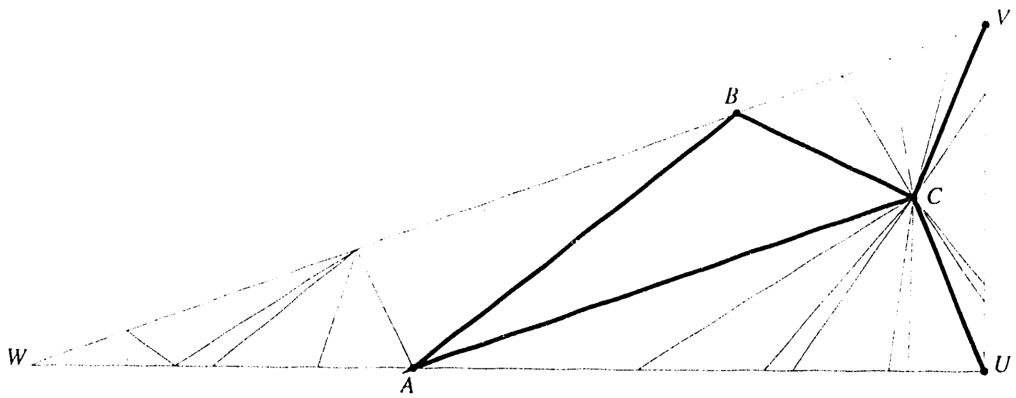

Figure 4: How to triangulate a triangle with Steiner points on its boundary.

Theorem 4 *The number of horns in the next iteration is at most* $3/4$ *times the number of horns in the current iteration. That is,* $m_{i+1} \leq 3m_i/4$.

Proof. The proof of this theorem is a matter of carefully specifying exactly which paths we create when horns intersect. Consider a horn H that intersects a horn H', where H' terminates because of one of Item 1 through Item 4. Then H induces a path in the next iteration, but H' does not. For the horns that terminate on a horn that also terminates because of Item 5, we are careful that at least a constant fraction of these are merged together. See Figure 3 right and Section 5 for the specification and proof. ∎

Theorem 5 *At most* $O(nm + np\log m)$ *Steiner points are added.*

3 Triangulating the non-conformal triangles

We now have a non-conformal triangulation in which some triangles have Steiner points on their edges. We now show how to triangulate these triangles.

3.1 Fixing triangles that have a small angle

We introduce the Steiner path edge $\overline{AB}$ when the angle between the triangle edges containing A and B is less than $\pi/8$. We say such edges are *drawn*. The other Steiner path edges are forever ignored. We have bounds on the angle that Steiner path edges make with triangle edges.

Lemma 9 *The angles at the intersection of a drawn Steiner path edge and a triangle edge are between* $\pi/8$ *and* $7\pi/8$.

Lemma 10 *Steiner path edges may be drawn so that no two cross, while maintaining the angle bounds of Lemma 9.*

We now erase all Steiner path edges $\overline{AB}$ where both $\angle UVA \leq \pi/4$ and $\angle VUB \leq \pi/4$. Since the edges do not cross, there will be an edge $\overline{AB}$ "closest" to $\overline{VU}$ that is not erased:

Any edge with a vertex on $\overline{AU}$ or $\overline{BV}$ will be erased, and any edge with a vertex on $\overline{AW}$ or $\overline{BW}$ will be drawn.

The remaining drawn edges intersect with $\triangle WVU$ to form triangles and trapezoids, together with the region between $\overline{AB}$ and $\overline{UV}$. According to Lemma 9 and Lemma 10 the triangles and trapezoids have largest angle no more than $7\pi/8$. Hence the trapezoids may be triangulated with an arbitrary diagonal and achieve largest angle no more than $7\pi/8$.

It remains to triangulate the region $ABVU$. It may be that $B = V$ or $A = U$. Also there may not be a drawn edge $\overline{AB}$, so that the region degenerates to $\triangle WVU$. The following construction needs no modifications for these cases.

We introduce a vertex C in the interior of the triangle. We place C so that $\angle CVU = \angle CUV = \pi/8$. Since the edges used to subdivide a large input angle in the very first horn iteration are considered as any other triangle edges throughout the remainder of the algorithm, we have that $\triangle VUW$ is non-obtuse (but also non-conformal). Hence $\angle VUW \leq \pi/2$ and $\angle UVW = \pi - \angle VUW - \angle VWU \geq 3\pi/8$. Similarly $\angle VUW \geq 3\pi/8$, so such a C exists inside $\triangle UVW$. We also have that $\angle CUW$ and $\angle CVW$ are at least $\pi/4$.

We triangulate by introducing an edge from C to each of the Steiner and triangle vertices in region $ABVU$. See Figure 4 and Section 5.2.

Lemma 11 *Region $ABVU$ may be triangulated with no angle larger than $7\pi/8$.*

3.2 Fixing triangles that have only large angles

It remains to consider triangles with every angle larger than $\pi/8$. Using a construction almost identical to that of Section 3.1, we may triangulate with no angle larger than $7\pi/8$ (see Figure 4). It is possible to use the fact that the angle at W is large to directly bound the largest angle in the region WAB. For region $ABVU$, the proof of Lemma 11 holds with slightly different angle bounds in various places, but with the same overall bound of $7\pi/8$. See Section 5.3 for the details.

4 Conclusions

There is a tradeoff between the cone angle of the horns and the number of times that a triangle edge is crossed. We state our cone angle to be $\pi/2$. If the cone angle is ϕ, we conjecture that the techniques of Section 3 can be used to obtain triangles with largest angle at most $3\pi/4 + \phi/4$. On a more general note, we have the following problems.

What is the relationship between the largest angle permitted in a triangulation and the number of Steiner points necessary to achieve that bound? How does this depend on the type of input (e.g. convex polygon, simple polygon, polygon with holes, PSLG)?

The cardinality of our PSLG triangulations is within a $\log n$ factor of worst case optimal. Is there an algorithm that is within a constant factor of worst case optimal? A more interesting open problem is the existence of an algorithm that generates triangulations of PSLGs or polygons with cardinality within a factor of optimal for the given input.

References

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

R. Barnhill [1983], Computer aided surface representation and design, *Surfaces in CAGD*, R. Barnhill and W. Boehm, eds., North-Holland, 1-24.

M. Bern, D. Dobkin, and D. Eppstein [1991], Triangulating polygons without large angles, submitted to *International Journal of Computational Geometry & Applications.*

M. Bern and D. Eppstein [1991] Polynomial-Size Nonobtuse Triangulations of Polygons, *7th ACM Symposium on Computational Geometry,* pp. 342-350.

M. Bern and D. Eppstein [1992], Mesh generation and optimal triangulation, *Computing in Euclidean Geometry,* D. Du, and F. Hwang, eds., World Scientific.

M. Bern, D. Eppstein, J. Gilbert [1990] Provably Good Mesh Generation, *Proc. 31st IEEE Symposium on Foundations of Computer Science,* 231-241.

H. Edelsbrunner, T.S. Tan, and R. Waupotitsch [1990], An $O(n^2 \log n)$ Time Algorithm for the MinMax Angle Triangulation, *Proc. 6th ACM Symposium on Computational Geometry,* 44–52.

D. Eppstein [1992], Approximating the Minimum Weight Triangulation, *3rd ACM-SIAM Symp. on Discrete Algorithms,* 48-57.

S.A. Mitchell and J. K. Park [1993], Approximating the MinMax Angle Steiner Triangulation when only Interior Steiner Points are Allowed, manuscript, Org. 1422, Sandia National Laboratories.

J. Rupert [1993], A New and Simple Algorithm for Quality 2-Dimensional Mesh Generation, *Proc. 4th ACM-SIAM Symp. on Discrete Algorithms,* pp. 83-92.

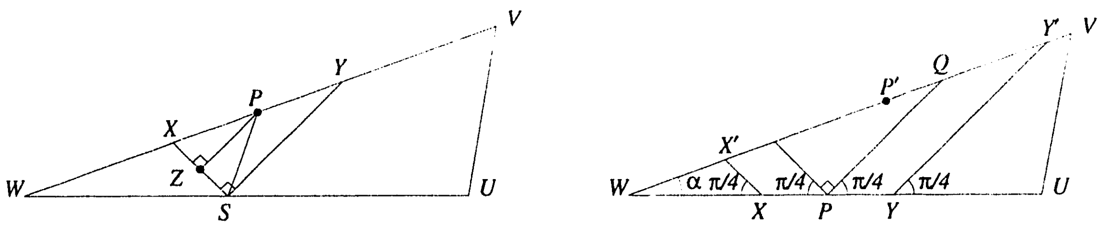

Figure 5: In a cone, the maw is twice the center path (left). In a horn, the maw is more than the center path (right).

5 Appendix

Proof of Lemma 2. The cone forms a triangle $\triangle SXY$ with angle 90° at S. Let Z be the midpoint of $\overline{SY}$, and P_1 the midpoint of XY. Then $\triangle YP_1Z$ is similar to $\triangle YXZ$ with ratio 1/2. Hence $|\overline{ZP_1}| = |\overline{SX}|/2$, and $\triangle SP_1X$ is isosceles. Thus the length of $\overline{SP_1}$ is $|\overline{XY}|/2$. See Figure 5 left. ∎

Proof of Lemma 3. We prove this inductively. For the first stage of the horn the lemma is true by Lemma 2. Consider the center path point P at a given stage of the horn's construction. Suppose that the theorem holds for the triangle edge containing P. We wish to extend the theorem to the next edge. We define $\overline{XY}$ to be the maw at the current stage and $\overline{X'Y'}$ to be the maw at the next stage. Define $\alpha = \angle X'WX$, the angle between the successive triangle edges.

Let P' be the center path point at the next stage. Consider the point Q on $\overline{X'Y'}$ such that $\overline{PQ}$ is parallel to $\overline{YY'}$. If we consider the segment $\overline{XR}$ also parallel to $\overline{YY'}$, we see that Q is the midpoint of $\overline{RY'}$. Hence P' lies closer to W than does Q. Furthermore, it is straightforward but tedious to show that $|\overline{P'W}| > |\overline{PW}|$ by the law of sines. Hence we have that $|\overline{PQ}| > |\overline{PP'}|$. We use $|\overline{PQ}|$ to bound $|\overline{PP'}|$ because it is much easier to analyize. Also, P' approaches Q as X approaches W, so there is nothing to be gained by considering $|\overline{PP'}|$ directly.

From the law of sines $|\overline{PQ}|/\sin\alpha = |\overline{PW}|/\sin(\pi/4-\alpha)$. Since $\sin(a+b) = \sin a\cos b + \sin b\cos a$, this simplifies to $|\overline{PQ}| = |\overline{PW}|\sin\alpha/(\cos\alpha - \sin\alpha)$.

We wish to bound $|\overline{PW}|$ in terms of the difference between the maw widths. Since $\alpha \le \pi/8$ and $\angle X'XW = \pi/4$, we have that $|\overline{X'W}| < |\overline{XW}|$. This implies $|\overline{X'Y'}| > |\overline{XY}| + |\overline{Y'W}| - |\overline{YW}|$. As before, from the law of sines we have $|\overline{Y'W}| = |\overline{YW}|/(\cos\alpha - \sin\alpha)$. Hence $|\overline{X'Y'}| - |\overline{XY}| > |\overline{YW}|(\frac{1}{\cos\alpha-\sin\alpha} - 1)$.

Since $|\overline{PW}| < |\overline{YW}|$, combining the last equations of the preceding two paragraphs we have $|\overline{PQ}| < (|\overline{X'Y'}| - |\overline{XY}|)\frac{\sin\alpha}{1-\cos\alpha+\sin\alpha} < |\overline{X'Y'}| - |\overline{XY}|$. Hence at the current stage the center path increase is less than the maw width increase, and by induction we are done. ∎

Proof of Lemma 4. The perpendicular to the maw in the direction of horn growth must make an angle change of at least π. The maw cannot contain a vertex of a triangle. Hence the center path length between the two intersections is minimized when one vertex of the maw almost contains a triangle vertex V, and the maw "pivots" around V, traveling through

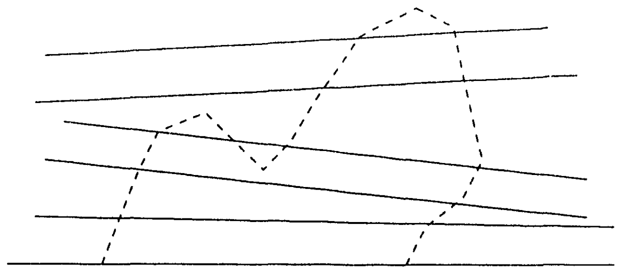

Figure 6: It is impossible for a center path to intersect an edge E twice without enclosing a vertex.

triangles containing V. In this case the center path lies outside a semi-circle of radius M centered at V. Hence the center path has length at least $\pi M/2$. Thus from Lemma 3 we have $M' > M + \pi M/2 > 2M$. ∎

Proof of Theorem 1. [Continued]

The fact that each hop has a vertex of the input in its interior can be proved by contradiction: Suppose that no vertex is interior to a hop. Since no input edge can cross E, each input edge is crossed by the center path an even number of times. Furthermore, since the input edges do not cross, there will be at least one "furthest" edge such that the center path crosses it twice consecutively; see Figure 6. But this is impossible by how the horn is defined, hence a contradiction. ∎

5.1 Ideas related to the proof of Lemma 6

We need some measure of how quickly two horns approach one another when they follow the same sequence of triangle edges. We formalize the distance between two horns as follows:

Inverse horn. An *inverse horn* is the region between two horns in a given sequence of triangles. The inverse horn starts on a triangle edge that the two horns have in common. An inverse horn terminates when it contains a triangle vertex, or when the bounding horns intersect. We assume that the horns have increasing stage as we increase along the sequence of triangles, so that the bounding edges of an inverse horn are comming closer together. Here we are concerned with the case that the two horns merely represent the same horn at different stages. Later, when we bound the length of the collection of horns, we will consider the case that the two horns to be distinct. We define the maw and center path of an inverse horn in the same way as we did for a horn.

An inverse horn may have negative width, in which case it behaves exactly as a horn, and the same lemmas and theorems about horn growth apply to how quickly an inverse horn width may shrink. Otherwise, if an inverse horn has positive width, we shall see that the analogous results hold, but the proofs needed are slightly different.

Lemma 12 *$M_i < M - C_i$, where M is the initial maw width of an inverse horn, M_i is its maw width at stage i, and C_i is the length of the center path to stage i.*

Proof. We prove this inductively in a way very similar to Lemma 3. We define $\overline{XY}$ to be the maw at a stage, and $\overline{X'Y'}$ to be the maw at the next stage, where the triangle

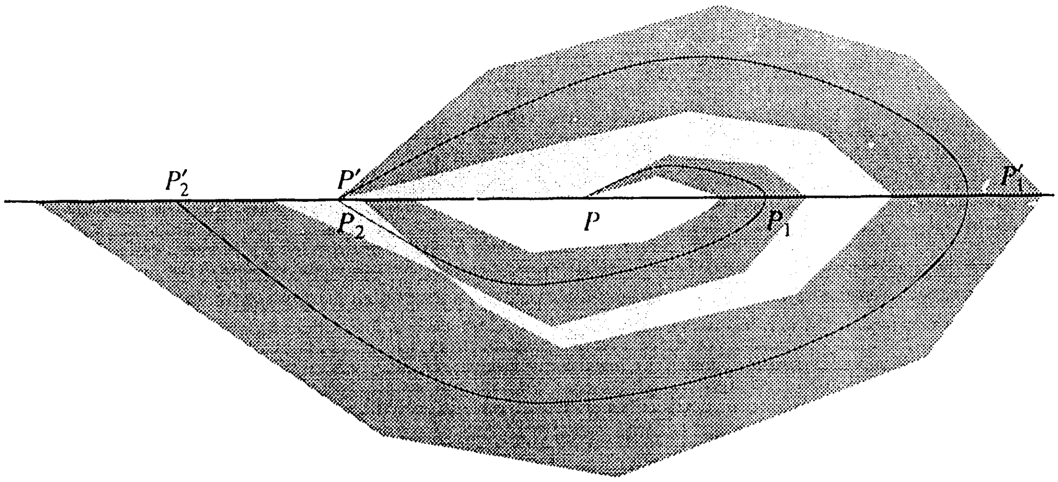

Figure 7: Here we show two consecutive reversals, their horns (darkly shaded), and the inverse horn between them (lightly shaded). The inverse horn terminates with negative width, where the two horns overlap.

edges containing these maws share vertex W. P is the center path point on $\overline{X'Y'}$. We let $\alpha = \angle X'WX$ and Q is defined to be the point on $\overline{WY'}$ such that $\angle QPW = \pi/4$.

Then $|\overline{X'Y'}| = |\overline{Y'W}| - |\overline{X'W}|$. Since $\angle WYY' = \pi/4$ and $\alpha < \pi/8$, we have that $|\overline{Y'W}| < |\overline{YW}|$. Hence $|\overline{X'Y'}| < |\overline{YW}| - |\overline{X'W}| = |\overline{XY}| + |\overline{XW}| - |\overline{X'W}|$.

From the law of sines we have $|\overline{X'W}| = |\overline{XW}| \sin\frac{\pi}{4} / \sin(\frac{\pi}{4} - \alpha) = |\overline{XW}|/(\cos\alpha - \sin\alpha)$. Hence $|\overline{X'Y'}| < |\overline{XY}| + |\overline{WX}|(1 - 1/(\cos\alpha - \sin\alpha))$, or $|\overline{WX}| < (|\overline{XY}| - |\overline{X'Y'}|)(\cos\alpha - \sin\alpha)/(1 - \cos\alpha + \sin\alpha)$.

Define P' to be the center path point of the maw on $\overline{X'Y'}$. It is clear from $|\overline{Y'W}| < |\overline{YW}|$ that $|\overline{PP'}| < |\overline{PQ}|$. Hence it suffices to bound $|\overline{PQ}|$. From the law of sines $|\overline{PQ}| = |\overline{PW}| \sin\alpha/(\cos\alpha - \sin\alpha)$. Hence $|\overline{PQ}| = (|\overline{WX}| + |\overline{XY}|/2)\sin\alpha/(\cos\alpha - \sin\alpha)$. Dropping the $|\overline{XY}|$ term and substituting in our equation for $|\overline{WX}|$ from the previous paragraph, we obtain $|\overline{PQ}| < (|\overline{XY}| - |\overline{X'Y'}|)\sin\alpha/(1 - \cos\alpha + \sin\alpha) < |\overline{XY}| - |\overline{X'Y'}|$.

Hence at each stage, the maw shrinks by at least the center path length, which completes the proof. ∎

We now formally begin the proof of Lemma 6.

Proof of Lemma 6. Consider a reversal R with center path points P, P_1 and P_2 on L, and another reversal R' with center path points P', P_1' and P_2' from later stages of the horn. Let $d = |\overline{PP_1}|$ and $d_1 = |\overline{PP_2}|$. See Figure 7. We consider the horn defined from a center path point P. We suppose R and R' contain the same vertices, and show a contradiction.

Suppose $P_2 = P'$ as in Figure 7. Let M_2 be the width of the maw at P_2. Define the horn from P', and let M_2' be its width at P_2'. If we assume that the horn from P does not self terminate at P_2 then $M_2/2 < d_1$.

Consider the inverse horn between the horns for P and P'. Let X be the center path point of the inverse horn between P_2 and P_2'. Then M^I, the inverse horn width at X, is $2|\overline{XP_2}| - M_2$. The initial inverse horn width is $d_1/2$. From Lemma 12 we have that $M^I < -d_1/2 - 2d - |\overline{XP_2}|$, which we may reduce to $3|\overline{XP_2}| < M_2 - d_1/2 - 2d$.

If the horn from P' does not self terminate at P'_2 then $|\overline{XP'_2}| > M_2/4$, and the above reduces to $d_1/2 + 2d < M_2/4$. Similarly, if the horn from P does not self terminate at P_2 then $M_2/2 < d_1$ and this reduces to $d_1/2 + 2d < d_1/2$, a contradiction. Hence either the horn from P must self terminate at P_2 or the horn from P' must self terminate at P'_2.

We have remaining the cases that $P' \neq P_2$. Suppose P' is on the same side of P as P_2. Either there is a reversal containing R on the path from P_2 to P', in which case we may consider that to be R', or there must be at least one hop from P_2 to P' not containing one of the hops of R. Furthermore, the hop is contained in R'. The hop contains a vertex. Hence R' contains a vertex that R does not.

If P' is on the opposite side of P as P_2, then the path from P' to P'_2 contains a vertex to the right of P_1, and hence not contained in the other reversal. ∎

Proof of Theorem 2. The center path of a horn does not self intersect, except to terminate as in Item 4. Hence two hops are either disjoint, or one completely contains the other. That is, the hops are partially ordered by containment. Hence there are at most $2p$ unique input vertex sets contained in hops. We enumerate the hops by charging the vertex sets for hops, such that each vertex set gets charged only a constant number of times.

If a hop is the smallest hop oriented in a particular direction along L containing a given vertex set, we charge that set for the hop. Each set may be charged once for each hop orientation, or a total of two times in this way. We now consider the case that two hops contain the same vertex set.

Let J be the largest hop contained in a hop I such that I and J are oriented in the same direction along L, and they both contain the same vertex set. By Lemma 5 we have that J precedes I on the center path. If there is no hop following I, we note that I is the last hop of the center path, and there can be only one such hop. We have two cases.

In the first case I and the hop I' following it is not a reversal. By applying Lemma 5, we may show that I' can not contain any other hop, and hence we charge its vertex set for I as well. The set for I' is charged at most once in this way, since there is only one smallest hop containing the vertex set of I', and I' is that hop.

In the second case I and the hop I' following it is a reversal. If J together with the hop J' following it is a reversal, then by Lemma 6, I' contains a vertex set different from J'. Since J was chosen to be the largest hop, there can be no hop between J' and I'. Hence I' is the smallest hop containing its vertex set with its orientation along L, and we charge the vertex set of I' for I. Otherwise, I' contains both J' and the hop preceding J, and these hops are disjoint. Since J was chosen to be the largest hop, together with Lemma 5, I' is the smallest hop containing its vertex set with its orientation along L. Hence we charge the vertex set of I' for I, and this set may be charged at most once in this way.

Thus every hop is either the smallest hop containing its vertex set for a given orientation along L, or is succeeded by such a hop, or is the very last hop of a center path. A hop has only one successor, and hence each of the vertex sets is charged for at most 4 hops. There are at most $2p$ vertex sets. Each hop corresponds to one additional center path point on L. Counting also the last crossing of L, a center path may cross a given edge at most $8p + 1$ times. The constant "8" in this bound is not tight. ∎

Proof of Lemma 7. Let H cross E at points X, Y and Z. Let H' cross E at X', Y' and Z'. Let t be the width of the inverse horn between H and H' on $\overline{YY'}$. It is easy to see that $t < |\overline{XY}|/2$. From Lemma 3 we have that the width of the maw of H at Z is at least $|\overline{XY}| + |\overline{YZ}|$. Since we assume that H does not self terminate, we also have that the width of the maw of H at Z is no more than $2|\overline{YZ}|$. Hence $|\overline{YZ}| > |\overline{XY}|$. From Lemma 12 the inverse horn maw width on $\overline{ZZ'}$ is at most $t - t/2 - |\overline{YZ}| < -3|\overline{XY}|/4 < 0$, and the horns intersect. ∎

Proof of Lemma 8. Label the points that the center paths intersect E as X, Y, Z, X', Y' and Z' respectively. Let t be the inverse horn maw width on $\overline{XX'}$, so $t < |\overline{XX'}|$. Since R' does not self terminate or terminate by containing X, then $|\overline{XY'}| < |\overline{Y'Z'}|$. Hence the inverse horn maw width at on $\overline{ZZ'}$ is at most $t - t/2 - |\overline{Y'Z'}| < 0$, and the horns intersect. ∎

Proof of Theorem 3. The proof is very similar to the proof of Theorem 2, Recall that there are m_i horns. All the hops are disjoint, except that the final hop in each of the m_i horns may intersect with some other hop. For the remaining (not last) hops, two hops are either disjoint, or one completely contains the other. That is, the hops are partially ordered by containment. Hence there are at most $2p$ unique input vertex sets contained in hops. We enumerate the hops by charging the vertex sets for hops, such that each vertex set gets charged only a constant number of times.

If a hop is the smallest hop oriented in a particular direction along L containing a given vertex set, we charge that set for the hop. Each set may be charged once for each hop orientation, or a total of two times in this way. We now consider the case that two hops contain the same vertex set.

Let J be the largest hop contained in a hop I such that I and J are oriented in the same direction along L, and they both contain the same vertex set. Let H be the horn containing J and H' the horn containing I. We rely on the proof of Theorem 2 for the case that $H = H'$. So we assume $H \neq H'$. If the hops following I or J intersect, or I or J is the last hop, we note that this can happen only m_i times in the collection of horns. Otherwise we have two cases.

In the first case I and the hop I' following it is not a reversal. Then from Lemma 7 I' contains a vertex set different from J', and we charge I to the vertex set of J'. Recall J was chosen be be the largest hop contained in I. If J and J' do not form a reversal, then J' is the smallest hop containing its vertex set, except perhaps for a hop that is the first hop of a center path. There are at most m_i first hops, so they may be ignored. Otherwise J and J' form a reversal. If J' is not the smallest hop containing its vertex set, then we note that there is no pair of hops I_1 and I_1' inside J, J' that can also charge the vertex set of J' in this way, and hence the vertex set of J' is charged at most once in this way.

In the second case I and the hop I' following it is a reversal. Since J was chosen to be the largest hop contained in I, we have that I' is the smallest hop containing its vertex set, except perhaps for a hop that is the first hop of a center path. As before, this later case may be ignored. Hence we may charge I to the vertex set of I'.

Each of the $2p$ vertex sets is charged at most twice for each orientation direction, and for at most two hops for each orientation. Hence we have at most $8p$ hops, and at most m_i crossings at the begining of a horn, for a total of $8p + m_i$ crossings. As in the proof of Theorem 6, the "8" is not tight. ∎

Proof of Theorem 4. To prove that the number of Steiner paths at each iteration decreases geometrically, we must specify exactly which acceptable Steiner paths we create. We have a number of cases, depending on why a horn H terminated, and on the horns that intersect it.

Case 1. H terminates because of Item 1, Item 2, Item 3 or Item 4. Then we introduce the acceptable Steiner path described in Section 2.2. This does not generate a horn in the next iteration.

Case 2. H terminates because of Item 5, and the horn H' it terminates on is oriented in the opposite direction along the maws' common triangle edge. (Note that H' must not have terminated because of this case as well, but we may erase any path for H' already created because of Case 1 above.) Then we may construct the Steiner path described in Section 2.2 for each horn. This does not generate a horn in the next iteration.

Case 3. H terminates because of Item 5, but the horn H' it terminates on is oriented in the same direction along the maws' common triangle edge. There is an acceptable Steiner path for both horns to a point of the intersection P as in Section 2.2. However, we may have to construct a horn from P in the next iteration. We have several subcases we must consider in order to exactly bound the number of horns in the next iteration.

Subcase A. $k \geq 2$ horns intersects H' on the same boundary path of H'. This subcase may apply to each boundary path of H'. Among all the horns intersecting H', there is one generating an intersection point P furthest from the starting vertex of H'. We may construct a Steiner path for each horn intersecting H' to a point P' on the boundary path of H'. We may also construct one path for H' from its starting vertex through each P' terminating at P. Thus we had $k \geq 2$ horns in this iteration, and one horn from P in the next.

Subcase B. H is the only horn terminating on a particular boundary path of H', and no Steiner path has been previously created for either H or H'. We introduce the acceptable Steiner path for H and H' to their intersection point P on the boundary path of H. We have two horns in this iteration, and may have one from P in the next.

Subcase C. H is the only horn terminating on a particular boundary path of H', but a Steiner path has been created for either H or H' by the previous cases or subcases. If a Steiner path was introduced on the boundary path of H that intersects H', and either no Steiner path was introduced for H' or one was introduced on the boundary path that intersects H, we complete the Steiner paths as in Subcase B. Otherwise we do nothing. The horn without a Steiner path will generate a horn in the next iteration. A horn H' that terminated because of Case 1 or 2 may have at most two horns terminating on it that meet this subcase. Otherwise, a horn H' may have at most one horn terminating on it that fits this subcase for which we do nothing.

We now count m_{i+1}, the number of horns in the next iteration. Let m_i be the number of horns in the current iteration. Let k be the number of horns terminating as in Case 1 or Case 2, and l be the number of horns terminating as in Subcase A or B. Then we have $l/2$ horns in the next iteration, plus q, the contributions from Subcase C. Hence we have $m_{i+1} \leq l/2 + q$, where $k + l + q = m_i$ and $q \leq 2k + l$. The worst case is achieved when $l = m_i/2$ and $q = l$. Hence $m_{i+1} \leq 3m_i/4$. ∎

Proof of Theorem 5. For each horn in an iteration, we generate at most three Steiner paths in it; one on each boundary path, and one in its interior. Combining this with

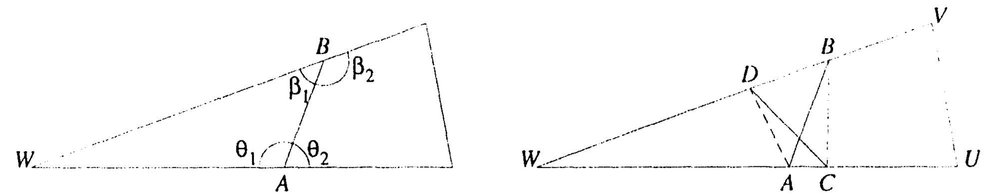

Figure 8: The angles between triangle edges and Steiner path edges are bounded (left). Uncrossing edges only improves angles (right).

Theorem 3, we generate at most $O(m_i + p)$ Steiner points on each of the n edges at each iteration. Since the number of horns must be an integer bigger than one, from Theorem 4 we have at most $\log_{4/3} m$ iterations, and the sum of the np terms is $O(np \log m)$. Similarly, since the m_i are bounded above by a geometric series, the sum of the nm_i terms is $O(nm)$. ∎

5.2 The angle bound for triangles that have a small angle

Proof of Lemma 9. Suppose the path is directed from A to B as in Figure 8 left. Such a Steiner path edge is contained in the cone at A. Hence the angles at A, θ_1 and θ_2 are bounded between $\pi/4$ and $3\pi/4$. The triangle edges meet at vertex W, where $\alpha = \angle AWB \leq \pi/8$. Let angle $\beta_1 = \angle WBA$, then $\beta_1 = \pi - \alpha - \theta_1 \geq \pi - \pi/8 - 3\pi/4 = \pi/8$. Similarly, $\beta_1 \leq \pi - 0 - \pi/4 = 3\pi/4$. This bounds its supplement, β_2, between $\pi/4$ and $7\pi/8$. ∎

Proof of Lemma 10. Consider any pair of edges $\overline{AB}$ and $\overline{CD}$ that cross in triangle $\triangle UVW$ as illustrated in Figure 8 right. We may swap vertices, forming Steiner path edges $\overline{AD}$ and $\overline{CB}$ that do not cross. It is easy to bound the new angles at A by the original angles: $\angle DAV \leq \angle DCU$ and $\angle DAW \leq \angle BAW$. We may similarly bound the new angles at B, C and D. ∎

Proof of Lemma 11. There are four natural subregions of Region $ABVU$ to consider.

Region UVC. Consider any triangle $\triangle CXY$ with X and Y lying on $\overline{UV}$. Now $\angle XCY \leq \angle VCU \leq 7\pi/8$. Assuming U is closer to X than Y, $\angle CYX$ is at most the supplement $\angle CUV$, or $7\pi/8$. Similarly for $\angle CXY$.

Region ABC. This region consists of the single triangle $\triangle ABC$, or is empty if no edge $\overline{AB}$ is drawn. Now $\angle CBA \leq \angle UBA \leq 7\pi/8$ by Lemma 10. Similarly for $\angle CAB$. The worst case for $\angle ACB$ is when $B = V$ and $\angle UVA = \pi/4$. But then $\angle CVA = \angle UVA - \angle UVA = \pi/8$. Hence $\angle ACB$ is at most the supplement of this, or $7\pi/8$.

Regions ACU **and** BCV. Consider the triangle $\triangle CXY$. Then $\angle XCY$ is at most $\angle UCA$, which in turn is at most the supplement of $\angle CUW$, or $3\pi/4$. Suppose U is closer to X than Y. Then $\angle CXU$ is also at most the supplement of $\angle CUW$, or $3\pi/4$.

It remains to bound $\angle CXY$. The worst case is when $Y = A$ and $\angle UVX = \pi/4$. See Figure 9. Since $\angle CXY$ is the supplement of $\angle CXU$, we seek a lower bound on the latter. From the law of sines, we have $\frac{|\overline{CU}|}{\sin \angle CXU} = \frac{|\overline{CX}|}{\sin \angle CUX}$, and $\frac{|\overline{CV}|}{\sin \angle CXV} = \frac{|\overline{CX}|}{\sin \angle CVX}$. Since $\triangle CVU$ is isosceles we have

$$\sin \angle CXU = \sin \angle CXV \frac{\sin \angle CUX}{\sin \angle CVX}.$$

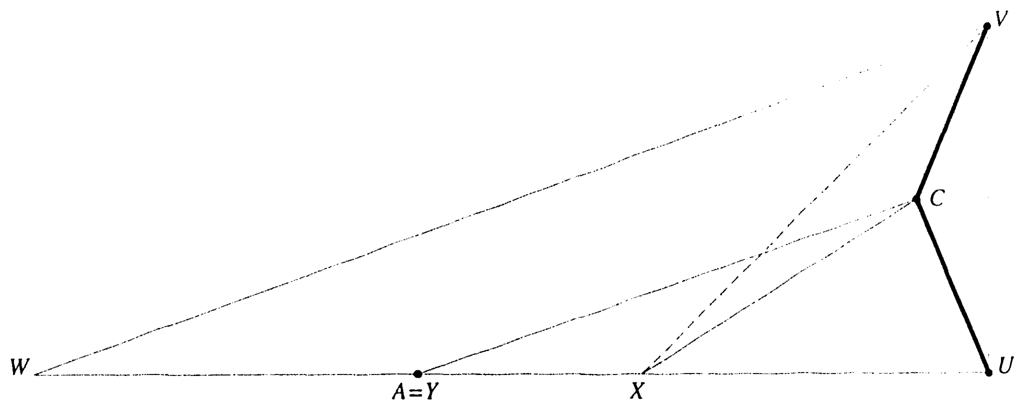

Figure 9: Angle CXY is bounded for any triangle in region ACU.

Note that $\angle CVX = \pi/8$ and $\angle CUX = \angle VUW - \pi/8 \geq \pi/4$. Hence $\angle CUX > \angle CVX$, so that $\angle CXU > \angle CXV$. Now $\angle VXU = \pi - \angle UVX - \angle VUX \geq \pi/4$. Since $\angle CUX + \angle CVX = \angle VXU$, we have $\angle CXU \geq \pi/8$. Hence $\angle CXY \leq 7\pi/8$. ∎

5.3 The angle bound for triangles that have only large angles

Consider fixing a non-conformal triangle that has all angles at least $\pi/8$. One way to triangulate such a triangle is as follows. First, introduce a Steiner point at the intersection of the angle bisectors. Second, introduce an edge between that point and each triangle vertex and Steiner vertex on the boundary of the triangle. Any triangle formed may have largest angle at most the supplement of half of the smallest angle of the triangle, or $15\pi/16$.

However, we can take advantage of the special structure of the Steiner points we introduced in order to triangulate with no angle larger than $7\pi/8$. The method is very similar to that used for a triangle with a small angle, but its description is necessarily different. We do not have the same description of drawn edges. Moreover, the angle that a cone makes with an edge may be close to the angle of the triangle at W, so that the proof of Lemma 9 will not hold. Fortunately, it is possible to use the fact that the angle at W is large to directly bound the largest angle in the region WAB. For region $ABVU$, the same proof as for a triangle with a small angle holds, with slightly different constants in various places, but the same overall bound of $7\pi/8$.

Consider a triangle ΔUVW whose smallest angle occurs at W, but is greater than $\pi/8$. We create a list R of the vertices on $\overline{WV}$, including W and V, sorted from W to V. We create L for $\overline{WU}$ similarly. We draw edges from L to R as follows. We draw an edge from the top of L to the top of R. If W is closer to the top of L than to the top of R, then we pop L, otherwise we pop R. We repeat until $\angle UVA \leq \pi/4$ and $\angle VUB \leq \pi/4$, where A is the top of L and B the top of R.

As in Section 3.1, denote the drawn edge closest to $\overline{UV}$ by $\overline{AB}$. The region WAB is triangulated, and region $ABVU$ is untriangulated.

Region WAB We first analyze the triangles of region WAB. Consider the triangle formed by W and the first edge drawn, ΔWA_1B_1. See Figure 10. Now $\angle WB_1A_1$ and $\angle WA_1B_1$ are bounded by the supplement of $\angle UWV$, or $7\pi/8$. Any triangle of the sequence

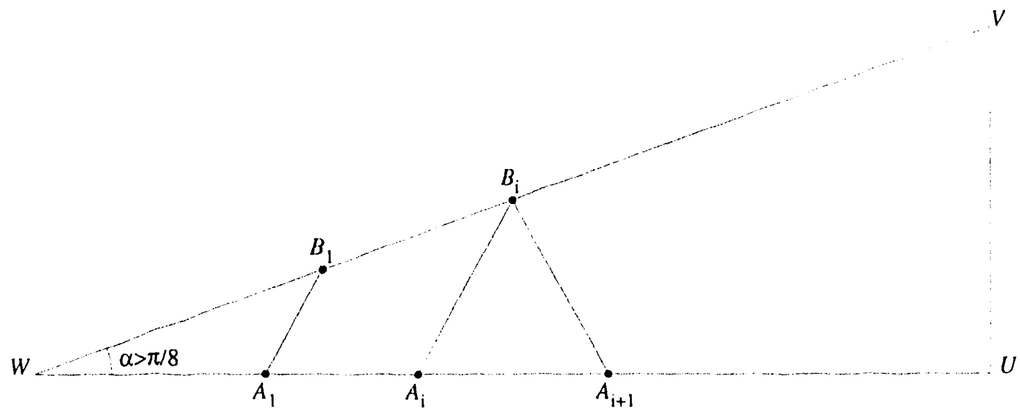

Figure 10: The angles in region ABW are all bounded.

may be analyzed in the same way. Without loss of generality, consider a triangle $\triangle A_iB_iA_{i+1}$, where W is closer to A_i than to B_i. Since $\angle UWV \leq \pi/3$, we have $\angle B_iA_iA_{i+1} \leq 2\pi/3$. Now, $\angle A_iB_iA_{i+1} \leq \angle WB_iU$. But this is at most the supplement of $\angle VWU$, or $7\pi/8$. Similarly, $\angle A_iA_{i+1}B_i \leq \angle WA_{i+1}V \leq 7\pi/8$.

Region $ABVU$. As in Section 3.1, we introduce a vertex C inside $\triangle WUV$ such that $\angle UVW = \angle VUW = \pi/8$. We then introduce an edge from C to each Steiner and triangle vertex in region $ABVU$. Again we have four natural subregions regions to consider.

The analysis differs slightly from Section 3.1 in that $\angle VUW$ and $\angle UVW$ may be smaller. We have $\angle VUW = \pi - \angle UVW - \angle UWV$, and W was chosen so that $\angle VUW \leq \angle VWU$. Hence $\angle VUW \geq \pi/4$, whereas in Section 3.1 the bound was $3\pi/8$. Similarly for $\angle UVW$.

Also our construction is defined differently. We first show that there is no Steiner point S in the region on $\overline{WU}$ such that $\angle UVS > \pi/4$, except perhaps for A itself. This fact was immediate in Section 3.1 by construction, but must be proven here. Suppose there were such a point. Since there is no drawn edge $\overline{SV}$, we must have that $|\overline{WV}| \leq |\overline{WS}|$. Hence $\angle WVS \geq (\pi - \angle UWV)/2$. But then $\angle WVU \geq (\pi - \angle UWV)/2 + \pi/4 \geq 7\pi/12$, a contradiction to the fact that we introduced edges to make all triangles non-obtuse. This also shows that either $\angle VUB \leq \pi/4$ or $\angle UVA \leq \pi/4$.

Region UVC. The analysis of region UVC is identical to that of Section 3.1.

Region ABC. We do not have Lemma 10 to bound $\angle CBA$ and $\angle ABC$. Suppose $\angle CAB > \angle CBA$. Hence W is further from A than from B, and the above shows that $\angle UVA \leq \pi/4$. The worst case is when $\angle UVA = \pi/4$, and the analysis of region ACU below shows that $\angle CAB \leq 7\pi/8$.

The worst case for $\angle ACB$ is when $B = V$ and $\angle UVA = \pi/4$, since either $\angle VUB \leq \pi/4$ or $\angle UVA \leq \pi/4$. As in Section 3.1, then $\angle CVA = \angle UVA - \angle UVA = \pi/8$. Hence $\angle ACB$ is at most the supplement of this, or $7\pi/8$.

Regions ACU **and** BCV. This differs from Section 3.1 in that the upper bound on $\angle VUW$ and $\angle UVW$ is $\pi/4$ instead of $3\pi/8$. Consider any triangle $\triangle CXY$. The same argument as in Section 3.1 shows that $\angle CUW \leq 7\pi/8$, instead of $3\pi/4$.

Consider the argument in Section 3.1 that $\angle CXY \leq 7\pi/8$. Here we have that $\angle CUX = \angle VUW - \pi/8 \geq \pi/8$, instead of $\pi/4$. However, since $\angle CUX$ is still greater than $\angle CVX =$

$\pi/8$, the proof remains correct.

5.4 Combining the cardinality bounds of Section 2 and Section 3

Theorem 6 *Any triangulation can be refined with at most $O(nm + np \log m)$ Steiner points so that no angle is larger than $7\pi/8$. Here n is the number of edges, m one plus the number of obtuse angles and p the number of holes and interior vertices in the original triangulation.*

Proof. From Theorem 5 we have $O(nm + np \log m)$ Steiner points added on triangle edges. We have an additional $O(n)$ vertices in Section 3 by adding one Steiner vertex in the interior of each non-conformal triangle. ∎

Theorem 7 *Any PSLG with v vertices can be triangulated with no angle larger than $7\pi/8$ using at most $O(v^2 \log v)$ Steiner points.*

Proof. By Euler's formula, any constrained triangulation of the input PSLG has at most $3v$ edges and $2v$ triangles. Each triangle can have one obtuse angle. Hence n, m, and p are all bounded by a constant times v. ∎

5.5 Bounding the running time

Theorem 8 *Any triangulation can be refined in time $O((nm + np \log m) \log(m + p))$.*

Proof. For each edge of a triangulation, we maintain a sorted list of the Steiner points it contains. This list is of length $O(m + p \log m)$ after $O(\log m)$ iterations. When we grow a horn to a new edge, it then takes $O(\log(m + p))$ time to determine if a Steiner point or triangle vertex on that edge is contained in the horn maw, and if not the interval in which to place a new Steiner point. Creating the Steiner path when a horn terminates then takes time linear in the length of the path. Triangulating a non-conformal triangle takes time linear in the number of Steiner points on its boundary of the triangle. ∎

DISCLAIMER

DATE
FILMED
10/15/93
END