Approximating the MaxMin-Angle Covering Triangulation

Extended Abstract

Scott A. Mitchell*

Abstract

Given a planar straight-line graph, we seek a *covering* triangulation whose minimum angle is as large as possible. A covering triangulation is a Steiner triangulation with the following restriction: No Steiner vertices may be added on an input edge. We give an explicit upper bound on the largest possible minimum angle in any covering triangulation of a given input. This upper bound depends only on local geometric features of the input. We then show that our covering triangulation has minimum angle at least a constant factor times this upper bound.

This is the first known algorithm for generating a covering triangulation with a provable bound on triangle shape. Covering triangulations can be used to triangulate intersecting regions independently, and so solve several subproblems of mesh generation.

1 Introduction

1.1 Covering triangulations

We propose a new class of triangulations called *covering triangulations*. We define a covering triangulation of an input planar straight line graph as a triangular graph in which the vertex set contains the input vertex set, and the edge set contains the input edge set. If the input is a polygon, then another way to view covering triangulations is that one is allowed Steiner points in the polygon's interior, but not on its boundary. Traditionally, most triangulation algorithms generate either a constrained triangulation or a Steiner triangulation. A *constrained triangulation* has a vertex set that is exactly the vertex set of the input, and an edge set that contains the edge set of the input. A *Steiner triangulation* has a vertex set that contains the vertex set of the input, and every edge of the input is the union of some edges of the triangulation.

From these definitions, we see that a constrained triangulation is a restricted type of covering triangulation, and a covering triangulation is a restricted type of Steiner triangulation.

1.2 Previous results

When seeking a constrained triangulation of a two dimensional input, one typically desires a triangulation that exactly optimizes some measure. The (constrained) Delaunay triangulation optimizes several measures; for example it maximizes the minimum angle (Lawson [1977]). Many algorithms exist for finding a Delaunay triangulation. The algorithms follow several different approaches such as plane sweep, edge flip, and incremental insertion. These algorithms are summarized in Fortune [1992], and Aurenhammer [1991] surveys the history of the Delaunay triangulations. An interesting observation is that like the present work, the Delaunay triangulation can be determined locally. In contrast a constrained triangulation that minimizes the maximum angle or maximizes the minimum height depends on global features of the input geometry (see Edelsbrunner, Tan, and Waupotitsch [1990] and Bern, Edelsbrunner, Eppstein, Mitchell,

* **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 Cornell University and at Sandia National Laboratories. This work 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.

and Tan [1992]). However, Mitchell and Park [1993] has now shown that a covering triangulation whose maximum angle is approximately as small as possible can be determined from the local geometry. Bern [1993] has recently shown that a covering tetrahedralization of a three dimensional polytope can always be constructed, but provides no shape bounds (a constrained tetrahedralization need not exist; see Schönhardt [1928]).

When seeking a Steiner triangulation, one usually settles for a triangulation that is only approximately optimal. This is because of the difficulty in obtaining an exactly optimal solution, and also because in many cases the optimal triangulation has an infinite number of triangles. We have algorithms that approximately maximize the minimum angle in Chew[1989], Bern, Eppstein and Gilbert [1990], and Rupert[1992]. Mitchell and Vavasis[1992] solves this problem in three dimensions. Bern, Dobkin and Eppstein [1991] presents an algorithm that maximizes the minimum height, and also algorithms that minimize the maximum angle. Bern and Eppstein [1992] summarizes much of the computational geometry literature relevant to Steiner and constrained triangulations.

A complete version of this paper appears in Mitchell [1993].

1.3 Application motivation

Triangulation of polyhedral regions is a fundamental geometric problem for numerical analysis. Finite element methods require that complicated input domains be discretized into a *mesh.* A triangulation is a typical choice for a mesh. Triangulations are also desired for other applications as well, such as solid modeling, functional interpolation, and computer graphics.

For numerical stability in the finite element method, it is necessary that the triangles of the mesh have some bound on their shape (Babuška and Aziz [1976]). We seek a triangulation whose minimum angle is as large as possible. This implicitly bounds the largest angle.

The algorithm we present here may find application in a number of mesh generation subproblems. By not adding Steiner points on a polygon's boundary, our algorithm allows us to triangulate intersecting regions independently. One application of this is in mesh refinement. That is, suppose a mesh exists for a given graph, but after running a finite element method we discover that the mesh is too coarse near a heat source. We may then "erase" the mesh in a neighborhood of the heat source. Then a finer mesh could be generated strictly inside this neighborhood, and our algorithm could be used to triangulate the gap between the outer and inner mesh.

Another application is to allow standard mesh generation algorithms to accept degenerate input, such as dangling edges with the interior of the polygon on both sides. Such an edge could be fattened to a hole of finite width, and our algorithm used to triangulate the hole after the exterior has been triangulated with a standard mesh generator. Edges shared by more than one region also arise when triangulating the surface of a three dimensional polytope. Such edges also arise in an octree decomposition of a domain, where one wishes to independently triangulate the two dimensional facets of the intersection of octree boxes with the input polyhedron (Mitchell and Vavasis [1992]).

In these examples, often one desires a covering triangulation with a constant bound on the smallest angle. We show that this is impossible for general input. However, we give an explicit characterization of the minimum angle generated by our algorithm in terms of the input geometry. Hence in the above examples, the boundary of the input may be preprocessed by adding vertices to long edges or "cutting off" sharp angles so that the desired bound is achieved.

2 Bounding the largest angle possible, A

An important feature of our algorithm is that we give explicit bounds on the minimum angle of our triangulation in terms of local geometric features of the input, namely edge lengths and certain derived angles. This is useful in the above examples where there is some flexibility in designing the polygon. This feature is in contrast with some of the constrained triangulation literature: Edge-insertion can be used to find a triangulation which minimizes the maximum angle, but can that angle be determined without generating the triangulation

Figure 1: The maxmin angle constrained triangulation (left) and covering triangulation (right).

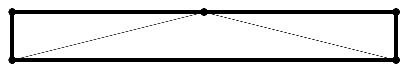

Figure 2: The optimal minimum angle for a covering triangulation may be as bad as linear in the ratio of the length of an edge to its closest interfering point.

itself? Our bounds on the minimum angle is a non-obvious function of the local input geometry. This contrasts with the Steiner triangulation literature: Determining the maximum minimum angle achievable is trivial, as it is an angle between two faces of the input.

Constrained triangulations are inadequate for the applications mentioned above. The four point example of Figure 1 shows that for a polygon the optimal minimum angle when Steiner points are not allowed may be much worse than the optimal minimum angle when Steiner points are allowed on the interior.

A key question is what features of a general planar straight-line graph P determine the optimal minimum angle A of a covering triangulation. For Steiner triangulations, Chew[1989] first noticed that the distance between an edge and an *interfering point*, a point on a closed face disjoint from the edge, determines how small triangles must be in order to have aspect ratio bounded by a constant. For a covering triangulation, because we are not allowed to add Steiner points on ∂P, it may be impossible to make triangles as small as required to have a constant bound on A, even when no edges of the input intersect at a small angle. For example see Figure 2.

Hence A should depend on the ratio of the distance to an interfering point to the length of an edge E. However, A is actually more closely dependent on the angle that the interfering point W makes with E at the closer vertex V of $E = \overline{VU}$. That is, if this angle, $\angle WVU$, is large, it may be possible to triangulate with minimum angle much *more* than the interfering point distance to edge length ratio.

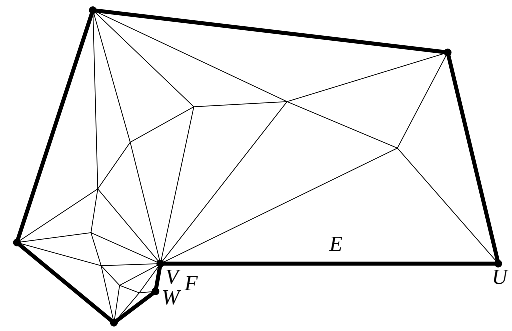

Figure 3: The optimal minimum angle may be much better than linear in the ratio of the length of an edge to its closest interfering point.

For an intuitive reason for why this is so, consider Figure 3. The closest interfering point to edge E is the vertex W of F. Suppose an adversary moves W towards V by shortening F, which changes the interfering point distance proportionately. To compensate we allow ourselves to change the position of the Steiner points. If the adversary halves the length of F, we may change the lengths of the edges containing V according to a geometric series, so that the resulting triangulation has minimum angle about $2^{-1/10}$ times that of the original triangulation. The "10" arises because there are ten similar triangles with vertex V. If F is made sufficiently short, then it will be advantageous to add more similar triangles with vertex V, so that the minimum angle depends logarithmically on $L(F)/L(E)$, where L(E) denotes the length of E.

We may be more precise. Consider an edge $E = \overline{VU}$ of P, and an interfering point W. Suppose V is is the closer vertex of E to W. Let r be the ratio of the distance between W and V to $L(E)$, and assume $r \leq 1$. Let $\omega = \angle UVW$. Then we define a *spiral* as the polar curve $\rho = L(E)s^{\phi}$, where $s = r^{\frac{1}{\omega}}$, the origin is at V, and the $\phi = 0$ axis is aligned with E. The spiral passes through U and W. For example, the spiral for E, V and W passes through the far vertices of the triangles containing V in Figure 3.

We define the corresponding *optimal spiral* to be an approximately best covering triangulation inside a spiral. The optimal spiral consists of a sequence of similar triangles all containing V. See Figure 3. The first and longest edge is E. Each successive edge has length a constant fraction (about e^{-1}) of the previous edge, and the angle at V of any triangle is about $\max(-1/\log s, \pi/2)$. Hence in general we have that an optimal spiral has minimum angle

$$h(E, W) = O(\omega / \max(1, -\log r)).$$

Recall $\omega = \angle UVW$ and $r = L(\overline{VW})/L(\overline{VU})$.

Note that special analysis is needed when r is close to one. The above values are derived from an optimization problem, using the fact that the minimum angle of a triangle is approximately the product of its angle at V and the ratio of the lengths of the edges containing V. The constraints of the optimization problem are merely that the first triangle edge is E, and the last edge is $\overline{VW}$. The solution to the relaxed problem where the last edge is not $\overline{VW}$, but still no triangle contains W, can be shown to be not much larger (recall W is closer to V than U).

In the full paper (Mitchell [1993]) we show that no covering triangulation of P can have minimum angle much larger than that of any optimal spiral. Hence an upper bound on A is A', the minimum of $h(E, W)$ over all edges E and points W of P. Our algorithm constructs a triangulation with minimum angle at least a constant factor times this upper bound A', and hence is within a constant factor of optimal.

3 The algorithm

The algorithm has three main steps. First we triangulate around the edges of P. We handle isolated vertices V of P by a preprocess that adds a short edge containing V to P. We also must bound a PSLG in order to have a well defined region to triangulate, but may use any means to accomplish this. For example, we may add the convex hull edges (and some Steiner points on them to avoid changing the maximum angle possible) or introduce a sufficiently large bounding box. We treat a PSLG edge as two edges, one for each side, unless it is on the boundary of the region to be triangulated. For every edge E of P and each of the vertices U and V of E, we find the spiral with smallest shrinking rate s. For each such spiral we raise s by a power of four, and decrease the initial edge length by a factor of four. This ensures that the spirals for different P vertices are far apart (in terms of A'), but only decreases the minimum angle of the corresponding optimal spiral by a constant factor. We then carefully weld optimal spirals to form a triangulation that completely covers the edges of P: Where two optimal spirals arise from the same input vertex V, we introduce an isosceles triangle where the curves intersect. Where two optimal spirals arise from the same input edge E, we rotate the first edge ($\overline{VB}$ and $\overline{UC}$) of each optimal spiral by the angle of a triangle (at V and U), then add edge $\overline{BC}$ and a diagonal ($\overline{BU}$ or $\overline{CV}$). See Figure 4. The edges we added that do not touch P form the boundary of an untriangulated polygon with holes Q lying inside P.

We can prove that the polygon Q is well shaped in two senses. First, all interior angles of Q are at least 0.39π. Second, somewhat surprisingly, for any edge of Q the ratio between the distance to the closest interfering point of Q to the length of the edge is at least a constant times A'.

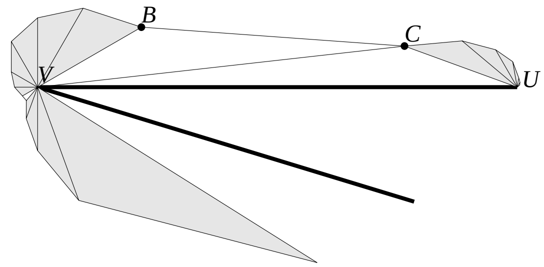

Figure 4: How to weld a spiral (shaded top left) to the spiral arising from the same input vertex (shaded bottom left) and the same input edge (shaded top right).

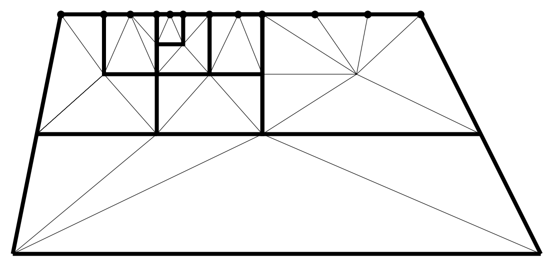

Figure 5: Triangulating between an R edge (top) and a Q edge (bottom). We use the spacing of Steiner points on the R edge to generate a hierarchy of quadrilaterals similar to a balanced quadtree.

Second we *shrink* Q to the polygon R, where R closely resembles a half-sized copy of Q lying inside Q. We take care that R is well shaped in the same sense as Q. We triangulate R using a maxmin angle Steiner triangulation algorithm. We chose to use the two dimensional analog of Mitchell and Vavasis[1992], but Bern, Eppstein and Gilbert[1990] or Rupert[1992] would also be acceptable. This introduces Steiner points on the boundary of R, but there are provable bounds on the number and spacing of Steiner points added along a given edge of R in terms of its length and closest interfering point distance.

Third we *match* R to Q, that is we triangulate the region between Q and R using the edges of Q as a guide. No new vertices on the boundary of either Q or R are introduced in this last step. See Figure 5. That Q is well shaped, combined with the bounds on the Steiner points on ∂R, show that the minimum angle introduced in this step is within a constant factor of A'.

The hardest task is not describing our algorithm but proving its optimality. The proof is a step-by-step argument that we never introduce a triangle with an angle smaller than a constant times A'.

4 Conclusions

If we use the medial axis to determine the interfering point yielding the spiral with smallest shrinking rate for each edge, then the running time of the algorithm is $O(\gamma \log \gamma + t)$. Here γ is the number of edges of the final triangulation, and t is the time taken to triangulate R. Using the two dimensional analog of Mitchell and Vavasis[1992] to triangulate R, t is no worse than $O(\gamma^2 \log \gamma)$. We also note that γ is bounded by the number of edges used to triangulate R, which is not optimal for the covering triangulation problem.

Is there an algorithm for generating a triangulation that approximately maximizes the minimum height? Mitchell and Park [1993] has recently shown how to generate a covering triangulation that approximately minimizes the maximum angle, where like the present work the optimal angle depends on a worst interfering point. However, the maxmin height possible in a covering triangulation appears to depend on a collection of interfering points, and hence its characterization is a more difficult problem.

Is there an algorithm for generating covering triangulations of three dimensional polytopes that maximize the minimum angle between a facet and an edge? Bern [1993] has shown how to generate a covering triangulation of a three dimensional polytope (without bounds on tetrahedron shape). The difficulty in extending the present work to three dimensions is that we must define an optimal spiral at a vertex of possibly high edge degree. The boundary of a three dimensional spiral may be a complicated surface, and not a two parameter curve as in two dimensions.

Acknowledgements

I thank Paul Chew for his insights concerning the development of optimal spirals and for discussing various triangulations with me. I thank Steve Vavasis for reviewing this work and suggesting several simplifications.

References

F. Aurenhammer [1991], Voronoi diagrams — a survey of a fundamental geometric data structure. *ACM Computing Surveys* 23, 345-405.

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

M. Bern [1993], Compatible Tetrahedralizations. *Proc. 9th Symp. on Computational Geometry*, 281-288.

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

M. Bern, H. Edelsbrunner, D. Eppstein, S. Mitchell, and T. S. Tan [1992], Edge-insertion for optimal triangulations, *Proc. Latin American Theoretical Informatics.*

M. Bern, 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, Xerox Palo Alto Research Center. Also, *Proc. 1990 Symposium on Foundations of Computer Science.*

L. P. Chew [1989], Guaranteed-Quality Triangular Meshes, C. S. Cornell, TR 89–983.

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.

S. Fortune [1992], Voronoi Diagrams and Delaunay Triangulations, *Computing in Euclidean Geometry,* D. Du, and F. Hwang, eds., World Scientific.

C.L. Lawson [1977], Software for C^1 surface interpolation, in *J. Rice*, ed., Mathematical Software III, Academic Press 161-194.

S. A. Mitchell [1993], Mesh Generation with Provable Quality Bounds, Cornell CS TR 92-1327.

S. A. Mitchell and J. K. Park [1993], Finding a Covering Triangulation whose Maximum Angle is Provably Small, Sandia National Laboratories, manuscript.

S. A. Mitchell and S. A. Vavasis [1992], Quality Mesh Generation in Three Dimensions, *Proc. 8th Symp. on Computational Geometry*, 212-221.

J. Rupert [1992], A New and Simple Algorithm for Quality 2-Dimensional Mesh Generation, Computer Science Division (EECS), UC Berkeley, Report No. UCB/CSD 92/694.

E. Schönhardt [1928], Uber die Zerlegung von Dreieckspolyedern in Tetraeder. *Math. Ananalen* 98 309-312.