

Triangulations Losing Bundles and Weight

Cem Boyacı Hale Erten Alper Üngör
 Dept. of Computer & Information Science & Engineering
 University of Florida
 {cboyaci,herten,ungor}@cise.ufl.edu

Abstract

We introduce bundle-free triangulations, that are free of large collection of triangles overlapping a circle empty of vertices. We prove that bundle-free Steiner triangulations can be used as an approximate solution for the minimum weight Steiner triangulation problem. We present new algorithms, implementations and experimental study for computing minimum weight Steiner triangulations.

1 Introduction

Consider the following well-known problem, known as the Minimum Weight Triangulation (MWT) problem:

Problem 1 [MWT] *Given a set of points in the plane, compute the triangulation that minimizes the summation of edge lengths.*

The summation of the edge lengths in a triangulation is referred as the *weight* of the triangulation. The MWT problem has been studied by many and has recently been proven to be NP-hard [MR06]. When points additional to the input points are allowed, we get another interesting triangulation problem, called the Minimum Weight Steiner Triangulation (MWST) [Epp94]. Additional points are called the *Steiner points*.

Problem 2 [MWST] *Given a set of points in the plane, compute the Steiner triangulation that minimizes the summation of edge lengths.*

An approximation algorithm for computing MWST is proposed by Eppstein [Epp94]. This algorithm uses quadtree refinement. Here, we propose alternative solutions for the MWST problem. Our solutions rely on a new geometric structure, called bundle-free triangulations. We show that quality triangulations are a special type of bundle-free triangulations. As a result, the Delaunay refinement algorithms can be used to solve MWST. Alternative heuristics for computing bundle-free triangulations with smaller weight are also studied.

There are various approximation algorithms [LK98, PH87, RS06] as well as heuristics and implementations

[AAH99, BS98, BDE02, DKC00] for computing minimum weight triangulations. Minimum weight Steiner triangulations, however, are much less studied. For instance, the complexity of the MWST problem remains as an open problem. We are aware of two approximation algorithms for computing MWST [Epp94, CL02]. [Epp94] presents an approximation algorithm for computing MWST of point sets and convex polygons. [CL02] extends this work to arbitrary polygons (with holes). We are not aware of a proper implementation of neither of these methods. Existing quadtree based quality triangulation software seem to be a reasonable alternative for this purpose.

2 Approximation Algorithms for MWST

Our algorithms and contribution leverage on the quadtree refinement result of Eppstein [Epp94], which we review in the next section. In addition, we rely on the following definition. Given a planar point set P , the *local feature size* at any point x in the plane, denoted $\text{lfs}_P(x)$, is the radius of the smallest disk centered at x that intersects two points in P [Rup93].

2.1 Quadtree MWST

A *quadtree* is a recursive partition of a region of the plane into axis aligned squares. One square, the root, covers the entire region. A square can be divided into four child squares with horizontal and vertical line segments through its center. The collection of squares then forms a tree with smaller squares at lower levels of the tree. A quadtree is said to be *balanced* if the size of any two adjacent leaf squares differ at most by a factor of two. Below, we summarize two results given by Eppstein [BEG94, Epp94] establishing a connection between the quadtree refinement, minimum weight Steiner triangulations, and the local feature size function.

Theorem 1 *Quadtree refinement gives an approximate solution for the MWST problem. Moreover, the size (side length) of each leaf quadtree cell is an approximation of the local feature size function in the cell.*

2.2 Bundle-free triangulations and their weight

Definition 1 A set of k triangles is called a **bundle**, if there exists a disk empty of all the input points and intersecting all k triangles.




Figure 1: Two sets of 9-bundle triangles is shown. This triangulation contains other 9 (and more)-bundle triangles.

A triangulation is said to be k -bundle-free if it contains no set of triangles that forms a k -bundle. In general, there may not be bundle free triangulation of a point set for small bundle constants, even if the Steiner points are allowed. However, for any input and sufficiently large bundle constant, it is possible to compute bundle free triangulations. This can be achieved by quality triangulation algorithms as described in Section 2.3. Achieving a quality constraint implies a bundle free triangulation. The converse is not true.

Lemma 2 Let T be a Steiner triangulation of a point set P and P' be the vertex set of T . If T is k -bundle-free and $lfs_{P'} = \Omega(lfs_P)$ then the weight of T is at most a constant factor of the weight of the minimum weight Steiner triangulation of P .

Proof. We overlay the bundle free triangulation T with a balanced quadtree subdivision of the domain that approximates the local feature size. For each quadtree leaf cell, we give a bound on the summation of the edge length of the bundle free triangulation. Let C be a quadtree cell with side length l and E is the set of edges of T that intersects C . We classify E into two subsets. Let E_1 be the set of edges that are incident to a vertex inside C . Let $E_2 = E - E_1$.

We first give a bound on the size of E_1 . Let V be the set of vertices inside C . Since the quadtree approximates the local feature size, $|V| \leq c_1$ where c_1 is a positive constant. We count the number of triangles incident to vertices in C , by locating two relatively small disks that are tangent to each other around each vertex in C . See Figure 2 (left). Each disk is empty of points by construction, and can not intersect with k or more triangles. Hence, the number of triangles incident to a vertex inside C is constant.

We now give a bound on the size of E_2 . We present a construction that carefully places a set of disks empty

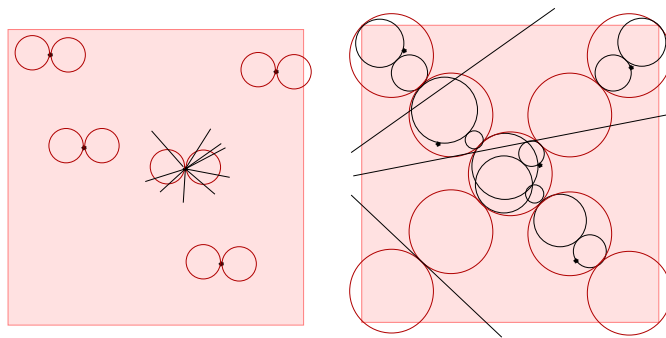


Figure 2: Counting the edges vertices incident to a vertex inside the quadtree cell. Counting the edges crossing the quadtree cell.

of other points in two stages. See Figure 2 (left). First, we pack a set of spheres along the diagonals of the cell C . We choose the radii of these disks a constant factor smaller than the side length of C such that each disk contains at most a single vertex. Next, we replace each disk that contains a vertex with two other disks excluding the point and still maintaining tangency along the diagonals of the cell. The disk at the center, if contains a vertex, is replaced with four disks in order to maintain four tangency points. Let D be the set of the resulting set of disks. By construction, $|D|$ is bounded by a constant. Our construction assures that an edge that does not have an endpoint inside C has to either intersect a disk in D or be tangent to two disks in D . There could be at most a constant number of edges tangent to two touching disks in D . For each disk, there could be at most $k-2$ edges intersecting it. Otherwise, the triangulation would not be bundle free. Hence, $|E_2|$ is bounded by a constant.

Since, $|E|$ is bounded by a constant and length of the portion of each edge inside C is at most $\sqrt{2}l$, the weight of the bundle free triangulation is within a constant factor of the weight of the balanced quadtree. ■

2.3 Delaunay Refinement for MWST

Delaunay refinement algorithms are originally proposed to compute quality triangulations (which have lower bound on the smallest angle). It is shown that the local feature size function with respect to their output is within a constant factor of the the local feature size function with respect to their input [Rup93]. The following lemma together with the lemma in the previous section suggests that Delaunay refinement algorithms can be used to give an approximate solution for the MWST problem.

Lemma 3 Let T be a triangulation with a lower bound α on its smallest angle. Then, T is bundle free for some constant k_1 .

Proof. Let c be an empty circle inside the domain of T . By a simple packing argument (considering arcs of angle 2α) there could be at most π/α triangles that has two or more edges intersecting c . This, in turn, implies that there could be at most $\pi/\alpha + 2$ triangles that has only one edge intersecting c . So, there are at most $2\pi/\alpha + 2$ triangles intersecting the circle. Hence, T is k_1 bundle free for $k_1 = 2\pi/\alpha + 3$. ■

Corollary 1 *Delaunay refinement approximates MWST.*

3 Experiments

Implementations. We are not aware of any implementation of the quadtree based MWST algorithms [Epp94, CL02]. However, the quadtree algorithms for MWST and for quality triangulations rely on the same principles. Hence, for comparison, we use a quadtree refinement software (Tripoint by Scott Mitchell) originally developed to produce quality triangulations.

Delaunay refinement implementations are parameterized by the quality constraint. We have shown already that Delaunay refinement gives an approximate solution for the MWST problem, for any constant quality guarantee (as long as the output point set respects the local feature size). In practice, it is worth to find which constraint values lead to smaller weight triangulations. Plots of the weight of Steiner triangulations with respect to varying angle constraint (Figure 3) reveal that the minimum weight is achieved generally when the angle constraint is in the range of $5^\circ - 10^\circ$.




Figure 3: Weight of the quality triangulation of the Ellipse and Florida data sets as the constraint angle changes from 0° to 20° . The weights are slightly smaller when the offcenters are used as the Steiner points instead of circumcenters within the Triangle software of Shewchuk.

We also implemented a bundle removal heuristic as a relaxation of the Delaunay refinement algorithm, allowing skinny triangles (with small angles) as long as they do not form a bundle. As a second heuristic edge-flip operations are utilized to improve the weight greedily.

Data Sets. We tested these algorithms on different type of planar data sets (point sets, convex polygons, polygons with holes) of various distributions (Random, Kuzmin, etc.). The Tripoint software has a limitation on the input type (NA entries in Table 1).

Results. Table 1 presents a summary of our experimental results. The boundary of the input must belong to every triangulation of the input. Hence, we report the weight of the boundary separately and exclude it in the other entries. Figure 4 shows the output of the mentioned algorithms on two different data sets.

It is rather clear from our experiments that the Delaunay based software are by far superior to the quadtree refinement implementation. For most data sets, we observed significant difference, between the weight of the initial Delaunay triangulation and the refined triangulation. The bundle removal heuristic is only slightly (5% or so) better performing than the Delaunay refinement.

4 Discussion

Our definition and approximation result on bundle free triangulations does not rely on the Delaunay property, while our current implementation does. It would be interesting to employ other triangulation algorithms (e.g., LMT-heuristic or other MWT heuristics) as part of our bundle removal strategy. In this direction, our preliminary experiments suggest some room for improvement when edge flips are integrated into our bundle removal heuristic.

Approximation constants for our results rely on the the bundle-free constant ($k_1 = 27$ for $\alpha = \pi/6$ which is realizable by Delaunay refinement), and how well the quadtree refinement approximates the local feature size and the MWST. As a result our overall approximation constant is too large to be relevant in practice. Finding a tight approximation bound is left as an open problem.

Both the quadtree refinement and the Delaunay refinement result in points sets that are well-spaced. With this observation in mind, we list the following interesting open problems.

Problem 3 *For a given well-spaced point set, is there a polynomial-time algorithm for computing the minimum weight triangulation?*

Problem 4 *Is the MWST of a well-spaced point set has finite size?*

Problem 5 *Given a point set P , is there an algorithm to detect whether the minimum weight triangulation and the minimum weight Steiner triangulation of P is the same?*

References

- [AAH99] O. Aichholzer, F. Aurenhammer, and R. Hainz. New results on MWT subgraphs. *Information Processing Letters*, 69(5):215–219, 1999.
- [BDE02] P. Bose, L. Devroye, and W.S. Evans. Diamonds are not a minimum weight triangulation’s best friend. *Int. J. Comput. Geometry Appl.*, 12(6):445–454, 2002.

Data Set			Weight					
name	# points	# edges	Boundary	DT	QTR	DR	BRH	BRH+EFH
Africa	54	54	30.7	100	454.9	96.4	98.1	88.5
Boeing	30	30	26.6	100	256.3	61.6	53.9	50.1
California	291	3603	1.2	100	305.1	88.9	83.4	76.6
Ellipse	1000	0	1	100	42.9	11.8	10.7	9.9
Florida	304	304	17	100	319.5	90.8	85.7	78.5
Kuzmin2K	2000	0	82.6	100	NA	91.6	64.0	63.2
New York	282	282	18.3	100	328.6	89.5	84.9	77.3
Plate	65	65	14.2	100	NA	58.6	54.6	52.7
PuertoRico	77	77	22.2	100	345.5	93.3	86.5	80.4
Random	1004	4	3.5	100	NA	97.3	95.9	93.8
Superior	522	522	13.9	100	271.8	88.0	78.8	74.2
Texas	641	641	11	100	284.4	78.1	73.0	69.0
Turkey	216	216	24.4	100	391.5	92.0	92.6	85.8

Table 1: For each data set, we present the weight of the output of the of the Delaunay triangulation (DT), the quadtree refinement (QTR), the Delaunay refinement (DR), the bundle removal heuristic (BRH), and the bundle removal with edge flip heuristic (BR+EFH) implementations are shown from left to right. The weight of the boundary of the domain is reported separately in the fourth column and excluded in these calculations. Also, the weights are normalized such that the Delaunay triangulation excluding the boundary has weight 100.




Figure 4: Output of the Delaunay triangulation, the quadtree refinement, the Delaunay refinement, the bundle removal heuristic, and the bundle removal with edge flip heuristic implementations are shown from left to right.

- [BEG94] M. Bern, D. Eppstein, and J.R. Gilbert. Provably good mesh generation. *J. Comp. System Sciences*, 48:384–409, 1994.
- [BS98] R. Beirouti and J. Snoeyink. Implementations of the LMT heuristic for minimum weight triangulation. In *Proc. Symp. on Comp. Geometry*, pages 96–105, 1998.
- [CL02] S.-W. Cheng and K.-H. Lee. Quadtree, ray shooting and approximate minimum weight Steiner triangulation. *Comput. Geom.*, 23(2):99–116, 2002.
- [DKC00] Y. Dai, N. Katoh, and S.-W. Cheng. LMT-skeleton heuristics for several new classes of optimal triangulations. *Comput. Geom.*, 17(1-2):51–68, 2000.
- [Epp94] D. Eppstein. Approximating the minimum weight triangulation. *Disc. & Comp. Geometry*, 11:163–191, 1994.
- [LK98] C. Levcopoulos and D. Krznaric. Quasi-greedy triangulations approximating the minimum weight triangulation. *J. Algorithms*, 27(2):303–338, 1998.
- [MR06] W. Mulzer and G. Rote. Minimum weight triangulation is NP-hard. In *Proc. Symp. on Computational Geometry*, pages 1–10, 2006.
- [PH87] D.A. Plaisted and J. Hong. A heuristic triangulation algorithm. *J. Algorithms*, 8(3):405–437, 1987.
- [RS06] J. Remy and A. Steger. A quasi-polynomial time approximation scheme for minimum weight triangulation. In *Proc. of the 38th ACM Symp. on Theory of Computing*, pages 316–325, 2006.
- [Rup93] J. Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. In *Proc. ACM-SIAM Symp. on Disc. Algorithms*, pages 83–92, 1993.