New Results on Minimum Weight Triangulations and the LMT-Skeleton *
(Extended Abstract)

Reinhard Hainz, Oswin Aichholzer, Franz Aurenhammer
Institute for Theoretical Computer Science
Graz University of Technology
Klosterwiesgasse 32/2, A-8010 Graz, Austria
e-mail: {rhainz|oaich|auren}@igi.tu-graz.ac.at

1 Introduction

A minimum weight triangulation (MWT) of a planar point set S minimizes the sum of the edge lengths among all triangulations of S. Computing a minimum weight triangulation is an important and interesting problem, whose complexity status is unknown. It is one of the open problems listed at the end of Garey and Johnson's book about NP-completeness [7].

In this work we present new methods to compute a subgraph of every MWT, based on [6]. This leads to an algorithm that computes the modified LMT-Skeleton (see Section 3 for an definition) in near linear time and space for uniformly distributed points. Moreover we give an improvement of the simple LMT-Skeleton which produces in general more edges and is in many cases equal to the intersection of all local optimal triangulations. Finally we show that for convex (and *near-convex*) polygons each edge of the β-Skeleton (cf Section 2) with $\beta > \sqrt{\frac{4}{3}}$ belongs to every MWT, too.

The best known methods for computing the MWT are based on these skeletons [2]. In fact, Jack Snoeyink reported that his group has computed the MWT for up to 60,000 uniformly distributed points within half an hour [10]. If all points of S are connected by a skeleton then the remaining edges of the MWT can be obtained from the polygonal regions by dynamic programming in polynomial time. If the skeleton is not connected then we have to insert additional edges in a depth-first-search like manner. Thus the running time of these algorithms depend on the number of components of the skeleton produced in the first stage. Up to now there is no skeleton of the MWT known which is guaranteed to be connected for every planar point set, but for moderately sized point sets the number of components is small and therefore the above mentioned algorithms work well [10].

The results presented in this extended abstract are based on the work of Hainz [8]. For a more complete description and discussion see also there.

*Research partially supported by the Spezialforschungsbereich F 003, *Optimierung und Kontrolle*.

2 The β-Skeleton

The β-Skeleton ($\beta \geq 1$) of S is a subgraph of the Delaunay triangulation. It consists of all edges xy with the property that the two disks with radius $\beta d(x,y)/2$ passing through both x and y, are empty of any further points from S. Mark Keil [9] proved that the β-Skeleton for $\beta \geq \sqrt{2}$ is contained in every MWT. This result was improved in [4] to $\beta \geq 1.17682$. A simple four point example shows that this does not hold for $\beta \leq \sqrt{\frac{4}{3}} \approx 1.1547$.

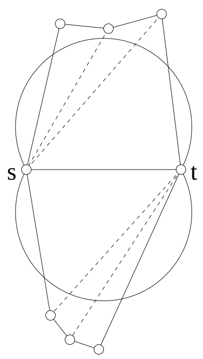

Figure 1: Example for a diagonal st satisfying the condition for pseudo star-shaped polygons.

Let us call a simple polygon P *pseudo star-shaped* if for every diagonal st of P all points of P to the left of st are visible from s and all points to the right of st are visible from t (or vice versa). See Figure 1 for an example. Note that any convex polygon is clearly pseudo star-shaped. We are now able to show the following result.

Theorem 1 *Let P be a simple pseudo star-shaped polygon. Then the β-Skeleton of P with $\beta > \sqrt{\frac{4}{3}}$ belongs to any MWT of P.*

Proof omitted in this abstract.

3 The LMT-Skeleton

A triangulation T is called *local optimal* if every quadrilateral of T is either non-convex or contains its shorter diagonal. Clearly every MWT is local optimal. The

LMT-Skeleton of a planar point set S is a subset of the intersection of all local optimal triangulations of S and thus a subset of every MWT of S. It was first defined algorithmically in [6]: The LMT-Algorithm first eliminates all edges which are not in at least one local optimal quadrilateral of S (these edges are called *lmt-exclusive*). An edge now becomes a LMT-Skeleton-edge if it is not intersected by any remaining edge. The algorithm repeats these two steps for all edges until no further edge can be either removed or made an LMT-Skeleton-edge. Note that the considered quadrilaterals have to consist of non removed edges.

It is surprising that this simple algorithm produces a skeleton which is most likely connected even for very large point sets (uniformly random distributed). Experiments show that this holds for sets of up to 15000 points. But note that for special point sets there exists a 19 point counterexample [2,3].

The running time of a straight forward implementation of the above algorithm is $O(n^6)$ and was given in [6]. Dickerson et al improved this to $O(kn^2)$, where k is the number of iterations. In experiments k tends to behave like $O(\log n)$ for uniformly distributed point sets. More sophisticated near linear time algorithms are presented in this work and are also reported by [10].

An *exclusion region* for an edge st $(s, t \in S)$ is a domain $(\subseteq \Re^2)$ that proves st to be not an edge of any MWT if it contains two points x, y $(x, y \in S\backslash\{s, t\})$ such that xy intersects st. One key fact is that the diamond exclusion region given by Das et.al. [5] is an exclusion region for the MWT and the expected number of remaining edges is only linear (about $20n$). Using this method to reduce the edge set before applying the LMT-algorithm may give different results compared to the simple LMT-Skeleton, but of course still a subset of the MWT. To emphasize this we call the edges of the LMT-Algorithm with reduced input the *modified LMT-Skeleton*.

Using the following observations, we are able to speed up the modified algorithm:

1. In [1] the authors introduced a linear time and space algorithm that computes a superset of the edges that remain after applying the diamond exclusion region. Experiments show that this set is about nine times larger then the set of all edges that come through the diamond test.

2. Any "inner" (not on or near the convex hull) point $a \in S$ has an expected constant degree. It can be shown that it is sufficient to check edges emanating from a only against intersections with edges having both endpoints as neighbours of a.

3. The LMT-Algorithm has *local* properties, see text below.

The first two observations show that the modified LMT-algorithm would run in $O(kn)$ time, if all points of S have similar properties as the inner points. Unfortunately this may not hold for points and edges near the boundary of S. But experiments show that this does not slow down the algorithm and that, for uniformly distributed points, we can expect a near linear running time, cf. Fig. 2.

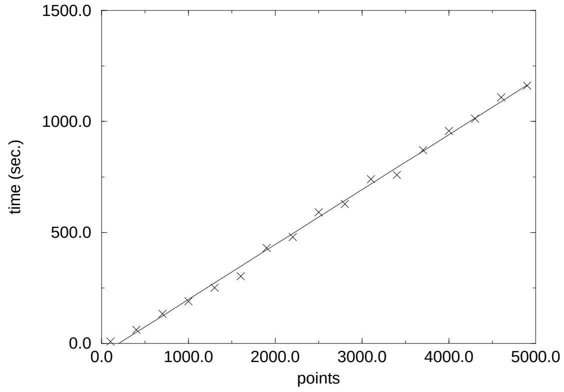

Figure 2: Running time of the modified LMT-Algorithm (PC 486) and their linear regression.

Let us now briefly consider the so called *LMT-Levels* which play an important role on the above observations. An edge is called of LMT-Level 1 if it is not a diagonal in any local optimal quadrilateral. An edge is called of LMT-Level i $(i > 1)$, if it is not a diagonal in any local optimal quadrilateral entirely consisting of edges which are not of LMT-Level l, $1 \leq l < i$. All remaining edges (i.e. edges which can not be removed by the LMT-algorithm) are called of LMT-Level 0.

With the definition of LMT-Levels it is easy to prove that there exists a sorting of the edges such that one iteration of the LMT-Algorithm is sufficient. On the other hand this idea also provides lower bounds for special point sets (e.g. $k = \Omega(n)$, see [8]).

A very useful observation from the algorithmic point of view is the following local property of the LMT-Algorithm.

Lemma 1 *Let L be the (modified) LMT-Skeleton of a planar point set S. Every edge $e \in L$ of LMT-Level i $(i \geq 1)$ has an common endpoint with at least one edge $e' \in L$ of LMT-Level $(i-1)$.*

Proof omitted in this abstract.

Using all above observations we developed an algorithm that runs in expected near linear time and linear space.

4 The improved LMT-Skeleton

The LMT-Skeleton is part of the intersection of all local optimal triangulations (for short: *lot*). But Figure 3 shows that in general not all edges of the lot are contained in the LMT-Skeleton and thus answers an open question in [6].

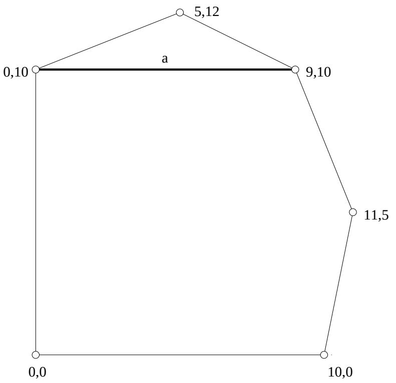

Figure 3: The LMT-Skeleton does not coincide with the intersection of all local optimal triangulations (*lot*): *a* belongs to the *lot* but not to the LMT-Skeleton.

We therefore consider the following extension to the simple LMT-Algorithm. If an edge e is a diagonal of a local optimal quadrilateral V then we also require that every edge f on the boundary of V satisfy one of the following conditions:

1. f is on the convex hull of S **or**

2. f is a diagonal of a local optimal quadrilateral that shares one triangle with V.

It is not hard to see that any quadrilateral must satisfy this conditions if it is part of at least one local optimal triangulation. Utilizing the methods of the simple LMT-Algorithm we implemented this new approach, and call the result of the algorithm the improved LMT-Skeleton.

With this modification we get the following nice property for simple polygons.

Theorem 2 *For simple polygons the improved LMT-Skeleton coincides with the intersection of all local optimal triangulations.*

Proof omitted in this abstract.

In many cases the LMT-Skeleton of a planar point set S is connected, which implies that the improved LMT-Skeleton for S is equal to *lot*. Note, that for uniformly distributed point sets the simple LMT-Skeleton is a good approximation to *lot*, i.e. in general there are only few edges added by the improved LMT-Skeleton.

Figure 4 shows a typical example for uniformly distributed points. There are $O(n)$ holes[1] and the largest hole has only a constant number of vertices.

References

[1] O. Aichholzer, R. L. S. Drysdale, G. Rote, *A Simple Linear Time Greedy Triangulation Algorithm for Uniformly Distributed Points*, IIG - Report-Series 408 (February 1995).

[1]A hole is a non-triangulated simple polygon with at least 4 vertices.

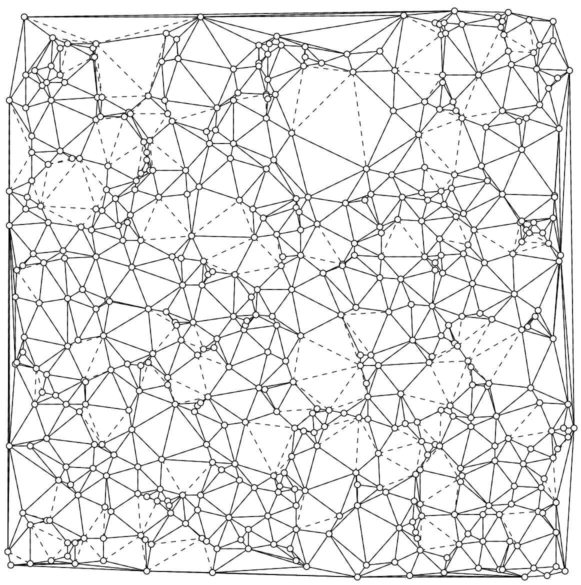

Figure 4: The improved LMT-Skeleton (solid) of 500 random points (uniformly distributed) and the remaining edges of a MWT (dashed). The improved LMT-Skeleton contains 58 holes with at most 14 edges.

[2] P. Belleville, M. Keil, M. McAllister, J. Snoeyink, *On Computing Edges That Are in All Minimum-Weight Triangulations*, Proc. 12th Annu. ACM Sympos. Comput. Geom. (1996), V7-V8.

[3] P. Bose, L. Devroye, W. Evans, *Diamonds are not a minimum weight triangulation's best friend*, Proc. 8th Canad. Conf. Comput. Geom. (1996), 68-73.

[4] Siu-Weng Cheng, Yin-Feng Xu, *Approaching the largest β-skeleton within a Minimum Weight Triangulation*, 12th Annual Symposium on Computational Geometry, Philadelphia, 1996.

[5] G. Das, D. Joseph, *Which Triangulations Approximate the Complete Graph?*, Optimal Algorithms, Int. Symp. Proc., LNCS 401, Springer-Verlag Berlin 1989, pp 168-192.

[6] Matthew T. Dickerson, Mark H. Montague, *A (Usually?) Connected Subgraph of the Minimum Weight Triangulation.*, Proc. 12th Annu. ACM Sympos. Comput. Geom. (1996), 204-213.

[7] M. Garey, D. Johnson, *Computers and Intractability*, A Guide to the Theory of NP-completeness, Freeman, 1979.

[8] Reinhard Hainz, *Minimum Weight Triangulierungen*, Diplomarbeit at the Institute for Theoretical Computer Science, Graz University of Technology, 1997.

[9] J. Mark Keil, *Computing a subgraph of the minimum weight triangulation*, Computational Geometry 4 (1994) 13-26.

[10] J. Snoeyink, Private Communication.