Karl-Franzens-Universität Graz & Technische Universität Graz

SPEZIALFORSCHUNGSBEREICH F 003

OPTIMIERUNG
und
KONTROLLE

Projektbereich
DISKRETE OPTIMIERUNG

O. Aichholzer F. Aurenhammer R. Hainz

New Results on MWT Subgraphs

Bericht Nr. 140 – September 1998

New Results on MWT Subgraphs

Oswin Aichholzer
Franz Aurenhammer
Reinhard Hainz

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

Abstract

Let P be a simple polygon in the plane and let $MWT(P)$ be a minimum-weight triangulation of P. We prove that the β-skeleton of P is a subset of $MWT(P)$ for all values $\beta > \sqrt{\frac{4}{3}}$ provided P is convex or near-convex. This settles the question of tightness of this bound for a special case and gives evidence for its validity in the general point set case.

We further disprove the conjecture that the so-called LMT-skeleton coincides with the intersection of all locally minimal triangulations, $LMT(P)$, even for convex polygons P. We introduce an improved LMT-skeleton algorithm which, for simple polygons P, exactly computes $LMT(P)$, and thus a larger subgraph of $MWT(P)$. The algorithm achieves the same in the general point set case provided the connectedness of the improved LMT-skeleton, which is given in allmost all practical instances.

Keywords: Computational geometry, minimum-weight triangulation, β-skeleton, LMT-skeleton

1 Introduction

A *triangulation* of a set S of n points in the Euclidean plane is a maximal set of non-crossing line segments (called *edges*) which have both endpoints in S. A triangulation of S that minimizes the sum of edge lengths is called a *minimum-weight triangulation*, $MWT(S)$, of S. Despite of the simplicity of this concept, its structural and computational properties are not well understood. For example, it is not known whether there exists an algorithm computing an $MWT(S)$ in time polynomial in n. For a catalogue of properties of optimal triangulations, and minimum-weight triangulations in particular, the reader may consult the recent survey paper by Aurenhammer and Xu [1].

Many efforts have been put into the study of subgraphs of $MWT(S)$. Gilbert [8] pointed out that the shortest edge defined by S always belongs to $MWT(S)$. Another simple observation is that *unavoidable edges*, which are edges not being crossed by any other edge defined by S, have to appear in any triangulation of S and thus are in $MWT(S)$. For example, all edges of the convex hull of S are unavoidable. The number of unavoidable edges does not exceed $2n-2$, see Xu [11], but usually is very small as most of these edges occur on the convex hull.

Only in recent years, several less trivial subgraphs of $MWT(S)$ have been identified. One of them arises from a class of empty neighborhood graphs defined by Kirkpatrick and Radke [10], and is called the *β-skeleton*, $\beta(S)$, of S. This graph is defined locally, and β is a parameter controlling the size of the neighborhood of an edge, to be empty of points in S for that edge to be included in $\beta(S)$. Interestingly, $\beta(S)$ is a subgraph of every $MWT(S)$ provided β is large enough. The original bound $\beta \geq \sqrt{2}$ in Keil [9] has been improved later in Cheng and Xu [5] to $\beta > 1.1768$. The largest value for which a (simple, four-point) counterexample is available is $\sqrt{\frac{4}{3}}$. To close this gap is an open problem. We show in this note (Section 2) that this can be achieved for the β-skeleton of convex polygons, and a certain class of star-shaped polygons. This strengthens the conjecture that the lower bound $\sqrt{\frac{4}{3}}$ is tight for arbitrary point sets S.

An essentially distinct subgraph of every $MWT(S)$ can be defined in a global way, via intersection of triangulations. Call a triangulation T of S *locally minimal* if every 4-sided polygon drawn by T and not containing points from S is optimally triangulated. That is, every convex quadrilateral contains the shorter one of its two diagonals. Let $LMT(S)$ denote the intersection of all locally minimal triangulations for S. Then $LMT(S)$ is a subgraph of every $MWT(S)$, as these triangulations of course are locally minimal, too. Whereas it is not known how to compute $LMT(S)$ in polynomial time, a large subgraph of $LMT(S)$, the so-called *LMT-skeleton* of S, can be computed by a simple and cute method, proposed in Belleville et al. [3] and in Dickerson and Montague [7].

The fact that the LMT-skeleton of S tends to be a connected graph even for large point sets S comes as a surprise, and for the first time allows for a rapid construction of $MWT(S)$ for practical purposes. Several variants of the LMT-skeleton have been considered recently, see [3, 7, 4, 2], but the question whether these skeletons coincide with $LMT(S)$ has remained open. In this note (Section 3) we give a counterexample. We further propose a new variant, the so-called *improved LMT-skeleton*, and show that this structure is identical to $LMT(S)$ when restricted to simple polygons. As a consequence, $LMT(S)$ for arbitrary point sets S coincides with the improved LMT-skeleton (and thus can be constructed in polynomial time) provided the connectedness of this structure, which is given in almost all practical instances. In this sense, the improved LMT-skeleton exploits the global subgraph approach to its utmost generality.

The following notation is used throughout. S denotes a set of n points in the Euclidean plane. For two points p and q in S, let pq be the (straight line) edge connecting them. When appropriate, an edge will also be considered just as a pair of points. The *length* of an edge pq is the Euclidean distance between p and q, denoted by $|pq|$. The *weight* of a set of edges is the sum of their lengths. Two edges are said to *cross* when they intersect in their interiors. When talking about triangulations or skeletons for some simple polygon P in the Euclidean plane, we will consider the restriction of these structures to the closure of P. That is, we only consider diagonals and boundary edges of P as possible triangulation or skeleton edges.

2 The β-skeleton

Let p and q be two distinct points in S and let $\beta > 1$. Following Kirkpatrick and Radke [10], the edge pq is included in the β-skeleton $\beta(S)$ of S if the two circles of diameter $\beta \cdot |pq|$ and passing through both p and q do not enclose any point in S. We will need the following observation Keil [9] used for relating $\beta(S)$ to $MWT(S)$.

Lemma 1 *Let pq be an edge of the β-skeleton $\beta(S)$ for $\beta > \sqrt{\frac{4}{3}}$, and let x and y be two points in S such that the line segment xy intersects the edge pq. Then $|xy| > \max\{|px|, |py|, |qx|, |qy|, |pq|\}$.*

In fact, only the weaker version of Lemma 1 for $\beta > \sqrt{2}$ is proved in [9] but, as mentioned there, the stronger version above still holds. The interested reader may check this by replacing triangle angles of $\frac{\pi}{4}$ by $\frac{\pi}{3}$ in the original proof.

As mentioned in the introduction, it is not known whether the β-skeleton is a subset of a minimum-weight triangulation for values of β close to $\sqrt{\frac{4}{3}}$. However, for the special case of convex polygons, the following can be shown.

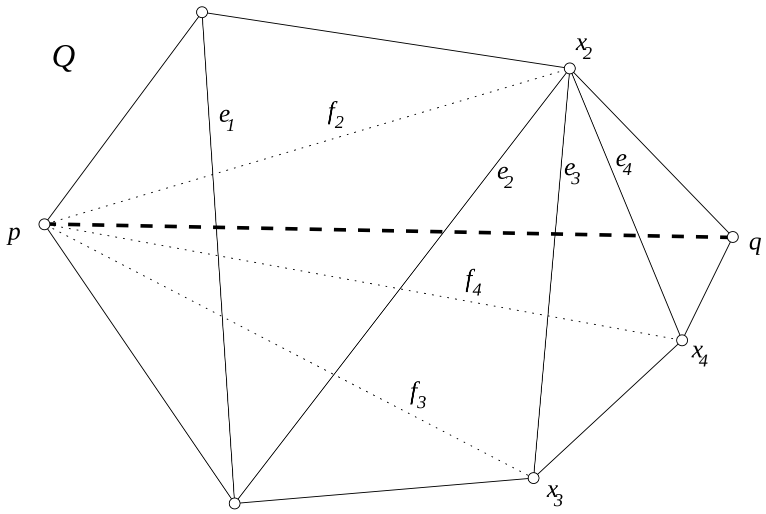

Figure 1: Proof of Theorem 1: pq is not in $MWT(P)$ and thus not in $\beta(P)$.

Theorem 1 *Let P be a convex polygon, let $MWT(P)$ be an arbitrary minimum-weight triangulation of P, and let $\beta(P)$ be the β-skeleton of P for some $\beta > \sqrt{\frac{4}{3}}$. Then $\beta(P)$ is a subset of $MWT(P)$.*

Proof. We prove the assertion by contradiction. Assume there is an edge pq of $\beta(P)$ which is not in $MWT(P)$. Then edge pq has to intersect some triangles of $MWT(P)$ properly. Let Q be their union, and let $MWT(Q)$ be the restriction of $MWT(P)$ to the convex subpolygon Q. Each of the $k \geq 1$ diagonals of $MWT(Q)$ is crossed by pq. Let $e_1, \ldots, e_k$ denote their total ordering with respect to increasing distance from point p. We now construct a new triangulation of Q which contains pq and has a weight smaller than $MWT(Q)$, giving a contradiction.

Consult Figure 1. First consider e_1, the edge closest to p. We have $|e_1| > |pq|$ by Lemma 1. So if $k = 1$, replacing e_1 by pq already yields the desired triangulation.

Else, for $i = 2, \ldots, k$, consider edge e_i, which has one endpoint x_i that is no endpoint of e_{i-1}. Let f_i denote the edge px_i. Again, $|e_i| > |f_i|$ by Lemma 1. So the weight of $\{e_1, \ldots, e_k\}$ exceeds the weight of $\{pq, f_2, \ldots, f_k\}$. It remains to be observed that $pq, f_2, \ldots, f_k$ indeed induce a triangulation of Q, as these edges are non-crossing (they all emanate from point p), and their number is k. □

Theorem 1 applies to a slightly larger class of polygons, as only visibility from p rather than convexity is required in the proof above. So the theorem is true for all polygons that are star-shaped as seen from the endpoints of their β-skeleton edges, for at least one endpoint per edge.

3 The improved LMT-skeleton

The LMT-skeleton of a finite point set S, introduced in Belleville et al. [3] and in Dickerson and Montague [7] (see also [6]), is based on the concept of locally minimal triangulations (see Section 1). Its definition is procedural and can be stated as follows.

Consider some edge set $E \subset S \times S$. An edge $e \in E$ is called *redundant in E* if e is no edge of the convex hull of S and there is no point-empty quadrilateral formed by E that has e as its shortest diagonal. Note that a redundant edge cannot appear in any locally minimal triangulation which is subset of E. Edge e is called *unavoidable in E* if no other edge in E crosses e. Unavoidable edges have to appear in every triangulation which is subset of E. The LMT-skeleton algorithm puts $E = S \times S$ and proceeds in several rounds. Each round identifies all edges redundant in the current set E, and then eliminates them from the set. When no more edge in the reduced set E can be classified as redundant, the algorithm includes all edges that are unavoidable in E into the LMT-skeleton, and then stops. It is clear that the produced LMT-skeleton is a subset of $LMT(S)$, the intersection of all locally minimal triangulations of S. The number of rounds (but not the LMT-skeleton) depends on the ordering in which the edges are examined.

Let us first exhibit an example where the LMT-algorithm fails to produce all edges of $LMT(S)$. In the convex and 7-sided polygon in Figure 2 only one edge (the longest diagonal, p_3p_6) is classified as redundant. Thus none of the remaining diagonals is unavoidable, that is, the LMT-skeleton algorithm leaves the polygon's interior empty. On the other hand, the polygon allows only for a single locally minimal triangulation, which therefore coincides with $LMT(S)$ (and with $MWT(S)$).

The reason why the algorithm does not produce $LMT(S)$ in general is buried in the definition of a redundant edge. An edge e may not be classified as redundant because of being shortest in some quadrilateral Q, but for some edge f of Q, no quadrilateral witnessing f's non-redundancy need to share a triangle with Q. In this case, e also cannot appear in any

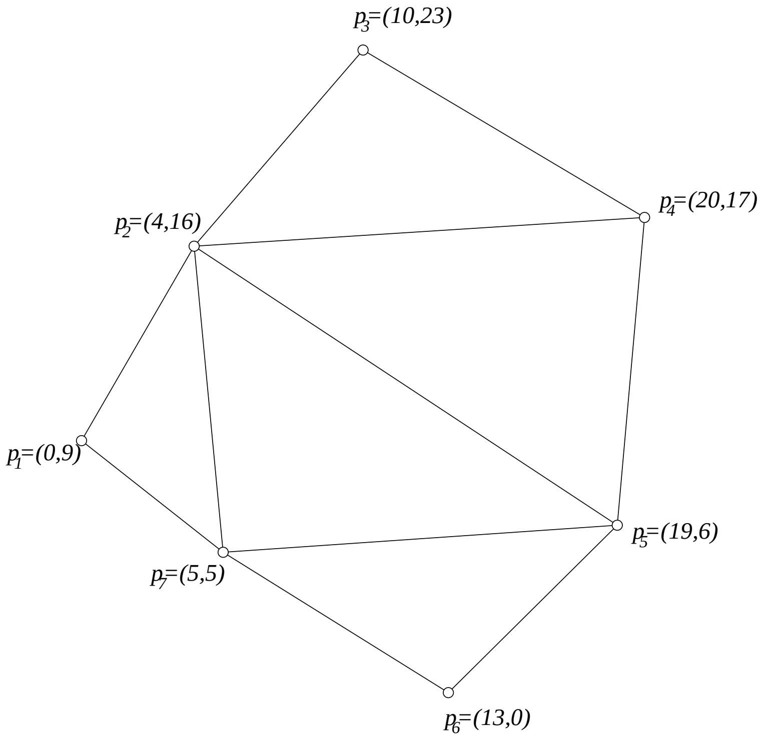

Figure 2: Convex 7-gon with a single locally minimal triangulation but empty LMT-skeleton.

locally minimal triangulation. We therefore strengthen the notation of redundancy. To be non-redundant in some set E of edges, e must either be a convex hull edge or be shortest in some point-empty quadrilateral Q from E, and for each edge f of Q that is no convex hull edge, there must exist some point-empty quadrilateral from E that has f as shortest diagonal *and* that shares a triangle with Q. The set of edges produced by the resulting modified algorithm is a superset of the original LMT-skeleton but still is a subset of $LMT(S)$. We call this set the *improved LMT-skeleton* of S, or $skel^+(S)$ for short.

We proceed to prove the following result on the improved LMT-skeleton for simple polygons. To adapt the concepts above — defined for arbitrary point sets S — to the polygon case it suffices to replace 'convex hull edge' by 'polygon boundary edge' in the definition of redundancy, and to initialize the set E so as to contain all boundary edges and diagonals of the polygon.

Theorem 2 *For any simple polygon P, $skel^+(P)$ coincides with $LMT(P)$.*

Proof. Consider some edge $e \in LMT(P)$. We assume $e \notin skel^+(P)$ and show that this leads to a contradiction.

Let E be the subset of edges that remains after repeatedly eliminating redundant edges (diagonals of P) with the improved LMT-algorithm. By definition, $skel^+(P)$ then consists of all edges that are unavoidable in E. Hence, by our assumption of $e \notin skel^+(P)$, edge e is not unavoidable in E. That is, there is another edge $f \in E$ that crosses e. As being contained in E, edge f cannot be redundant. Also, as crossing e, edge f is no boundary edge of P. So there must exist some quadrilateral Q formed by E that has f as its shortest diagonal. We include f into an initially empty set T of edges and note that — by construction with the improved LMT-algorithm — for each edge g of Q that is no boundary edge of P, there must be some quadrilateral Q' which has g as its shortest diagonal and which shares a triangle with Q. We add each such edge g to the set T, too, and repeat this process for the edges of Q' and so on, until no more edges can be added to T.

In fact, constructing T triangulates the polygon P. Each edge of T is a diagonal of some quadrilateral and of P as well, so quadrilaterals sharing a triangle are encountered in a tree-like fashion. Triangulation T, on the one hand, is locally minimal and, on the other, does not include edge e. But this contradicts the assumption of $e \in LMT(P)$. □

The proof above may fail for the case of general point sets S. We do not know whether the construction of T is guaranteed to lead to a triangulation of S in this case. Edges crossing each other might be included, though we did not succeed to give an example. Still, the following observation can be made for general point sets.

Corollary 1 *Let S be a finite and planar point set, and assume that $skel^+(S)$ forms a connected graph. Then $skel^+(S) = LMT(S)$.*

Proof. If $skel^+(S)$ forms a connected graph then it subdivides the convex hull of S into simple polygons. Each polygon boundary edge belongs to $LMT(S)$ because of $skel^+(S) \subset LMT(S)$, and in the interior of each polygon we have the identity of $skel^+(S)$ and $LMT(S)$ by Theorem 2. □

4 Concluding remarks

We have solved two open questions on β-skeletons and LMT-skeletons for restricted cases, contributing evidence to the conjecture that Theorem 1 and Theorem 2 might be true for the general point set case. In particular, the improved LMT-algorithm will indeed construct the intersection of all locally minimal triangulations $LMT(S)$, except for very specially constructed point sets S, which is not the case for previous LMT-algorithms. Its superiority carries over to the practically more relevant situation where most edges are removed from the start set $E = S \times S$ by pre-exclusion tests before an LMT-algorithm is run.

In order to define subsets of an $MWT(S)$ richer than $LMT(S)$, let us consider the following generalization of local minimality. For fixed k, call a triangulation of S *k-minimal* if it is minimum-weight within each of its k-sided and point-empty polygons. Let $LMT_k(S)$ denote the intersection of all k-minimal triangulations of S. Clearly, $LMT_i(S) \subset LMT_j(S)$ for $3 \leq i < j \leq n = |S|$. $LMT_3(S)$ is the set of unavoidable edges and $LMT_4(S) = LMT(S)$. Interestingly, there exist point sets S (in convex position) where $MWT(S)$ is unique but $LMT_{n-1}(S)$ is empty. We raise the question of constructing $LMT_k(S)$ efficiently for general k. At present, we do not even know of a polynomial-time algorithm for computing 5-minimal triangulations. Popular strategies like edge flipping or greedy edge insertion, which are well known to produce 4-minimal (i.e. locally minimal) triangulations, are easily shown to fail.

References

[1] F.Aurenhammer, Y.-F.Xu, *Optimal Triangulations*, Encyclopedia of Optimization, Kluwer Acad. Publ., to appear.

[2] R.Beirouti, J.Snoeyink, *Implementations of the LMT heuristic for minimum weight triangulation*, Proc. 14th Ann. ACM Symp. on Computational Geometry, 1998, 96-105.

[3] P.Belleville, M.Keil, M.McAllister, J.Snoeyink,$IMP(P)$ *On computing edges that are in all minimum-weight triangulations*, Proc. 12th Ann. ACM Symp. on Computational Geometry, 1996, V7-V8.

[4] S.-W.Cheng, N.Katoh, M.Sugai, *A study of the LMT-skeleton*, Proc. Int. Symp. on Algorithms and Computation (ISAAC), Lecture Notes in Computer Science 1178, Springer Verlag, 1996, 256-265.

[5] S.-W.Cheng, Y.-F.Xu, *Approaching the largest β-skeleton within a minimum-weight triangulation*, Proc. 12th Ann. ACM Symp. on Computational Geometry, 1996, 196-203.

[6] M.T.Dickerson, J.M.Keil, M.H.Montague, *A large subgraph of the minimum weight triangulation*, Discrete & Computational Geometry 18 (1997), 289-304.

[7] M.T.Dickerson, M.H.Montague, *A (usually?) connected subgraph of the minimum weight triangulation*, Proc. 12th Ann. ACM Symp. on Computational Geometry, 1996, 204-213.

[8] P.D.Gilbert, *New results in planar triangulation*, M.S. thesis, Coordinated Science Laboratory, University of Illinois, Urbana, 1979.

[9] M.Keil, *Computing a subgraph of the minimum weight triangulation*, Computational Geometry: Theory and Applications 4 (1994), 13-26.

[10] D.G.Kirkpatrick, J.D.Radke, *A framework for computational morphology*, G.T.Toussaint (ed.), Computational Geometry, Elsevier, Amsterdam, 1985, 217-248.

[11] Y.-F.Xu, *On stable line segments in all triangulations*, Appl.Math.-JCU 11B, 1996, 235-238.