



Constant-Level Greedy Triangulations Approximate the MWT Well

OSWIN AICHHOLZER
FRANZ AURENHAMMER

oaich@igi.tu-graz.ac.at
auren@igi.tu-graz.ac.at

Institute for Theoretical Computer Science, Graz University of Technology, Klosterwiesgasse 32/2, A-8010 Graz, Austria

GÜNTER ROTE

rote@ftug.dnet.tu-graz.ac.at

Institut für Mathematik, Technische Universität Graz, Steyrergasse 30, A-8010 Graz, Austria

YIN-FENG XU

yfxu@xjtu.edu.cn

School of Management, Xi'an Jiaotong University, Xi'an Shaanxi, 710049, PRC

Received April 23, 1996; Revised January 23, 1997

Abstract. The well-known greedy triangulation $GT(S)$ of a finite point set S is obtained by inserting compatible edges in increasing length order, where an edge is compatible if it does not cross previously inserted ones. Exploiting the concept of so-called light edges, we introduce a definition of $GT(S)$ that does not rely on the length ordering of the edges. Rather, it provides a decomposition of $GT(S)$ into levels, and the number of levels allows us to bound the total edge length of $GT(S)$. In particular, we show $|GT(S)| \leq 3 \cdot 2^{k+1} |MWT(S)|$, where k is the number of levels and $MWT(S)$ is the minimum-weight triangulation of S .

Keywords: greedy triangulation, minimum-weight triangulation

1. Introduction

A *triangulation* of a given set S of n points in the plane is a maximal set of non-crossing line segments (called *edges*) whose both endpoints belong to S . Besides the Delaunay triangulation and the minimum-weight triangulation, the greedy triangulation (GT) is among the three most prominent ones. It is obtained by inserting compatible edges in increasing length order, where an edge is compatible if it does not cross previously inserted ones. Various algorithms for computing the GT are known, and the GT has been used in several applications. See, e.g., (Dickerson et al., 1994) for a short history.

One use of the greedy triangulation is a length approximation to the minimum-weight triangulation (MWT). For a given point set S , the MWT minimizes the total edge length for all possible triangulations of S . Unfortunately, there are no known efficient algorithms for computing a MWT for general point sets. Therefore, efficiently computable approximations to the MWT are of importance.

Although the GT tends to be short in practical applications, and is provably short for uniformly distributed point sets (Levcopoulos and Lingas, 1992) and for point sets in

convex position (Levcopoulos and Lingas, 1987), its worst-case length behaviour is fairly bad. The GT can be a factor of $\Omega(\sqrt{n})$ longer than the MWT; see (Levcopoulos, 1987). Only very recently, a matching upper bound has been proved (Levcopoulos and Krznic, 1996).

In this note, we prove an upper bound of the form $|GT(S)| \leq c_k \cdot |MWT(S)|$, where c_k is a constant depending on the shape of S but not on its size. In Section 2, we introduce a decomposition of $GT(S)$ into k disjoint sets of edges, called *levels*, where the parameter $k \geq 1$ results from the shape of S . The level decomposition provides an alternative way of viewing, or defining, $GT(S)$ without having at hands the sorted list of edges spanned by S . It is based on the concept of *light* edges, introduced in (Aichholzer et al., 1996), and allows us to gain more insight into the structure of $GT(S)$. In particular, we show $c_k \leq 3 \cdot 2^{k+1}$ for $k \geq 2$ in Section 3. This generalizes the result $c_1 = 1$ in (Aichholzer et al., 1996) and implies that a GT with constantly many levels is a constant approximation to the MWT. Section 4 studies the number of levels in a GT and offers a short discussion of the presented topic.

As we have learned recently, level decompositions of GTs have been considered independently in (Jansson, 1995), for the sake of an efficient parallel computation of greedy triangulations. An improved length bound of the form $|GT(S)| \leq O(k) \cdot |MWT(S)|$ appeared to be implicit in the independent work (Levcopoulos and Krznic, 1996). The constant hidden in the O-notation is very high, though.

2. A level decomposition of GT

The usual procedural definition of $GT(S)$, as given in Section 1, resorts to the length properties of the edges spanned by S as well as to their crossing properties.¹ In particular, an edge which is not crossed by any shorter edge will surely belong to $GT(S)$. Let us call an edge *light* in this case. Below is a catalog of basic properties of light edges.

Lemma 1. *Let L denote the set of all light edges defined by S .*

- (a) *L is a non-crossing set of edges.*
- (b) *L contains all edges bounding the convex hull of S .*
- (c) *L is a subset of $GT(S)$.*
- (d) *In general, L is not a subset of $MWT(S)$.*

The following result is less obvious and is proved in (Aichholzer et al., 1996). Let $|A|$ be the *weight* of a given set A of edges, that is, the sum of the lengths of all the edges in A .

Lemma 2. $|L| \leq |MWT(S)|$.

In conjunction with Lemma 1(c), Lemma 2 immediately implies: If L happens to form a triangulation of S , then $|L| = |GT(S)| = |MWT(S)|$. In any case, we learn that at least a subset of the edges in $GT(S)$ can be bounded in length by the weight of $MWT(S)$.

In general, the following fact prevents particular edges in $GT(S)$ from being light. An edge e , though showing up in $GT(S)$, may still be crossed by some edge f shorter than e ,

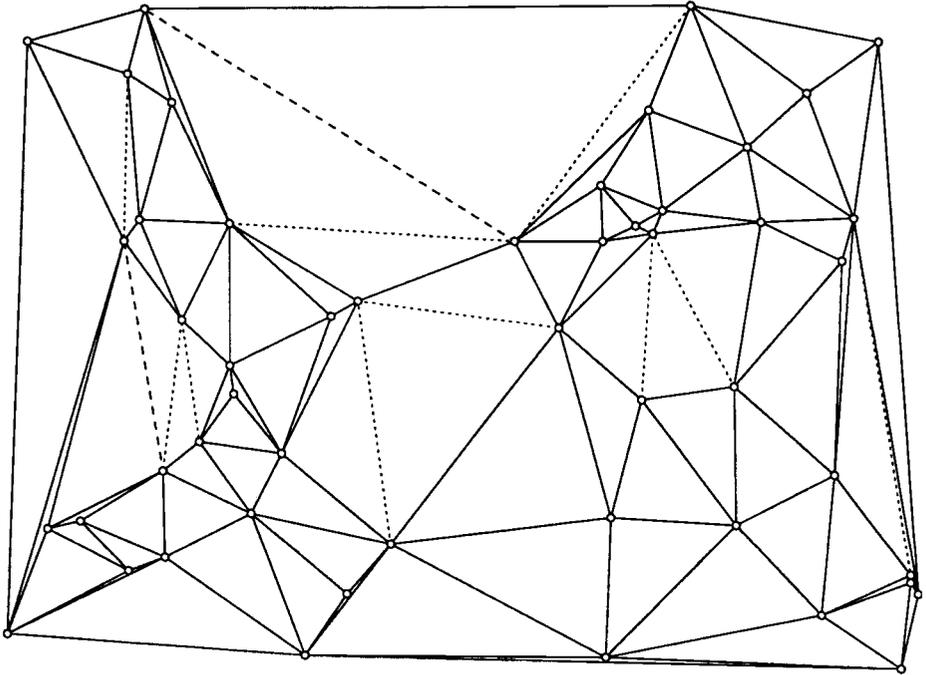


Figure 1. GT with three levels shown solid, dotted, and dashed.

as f might be non-compatible, that is, f is crossed by shorter edges which have already been inserted into $GT(S)$. This fact suggests to generalize the definition of light edges.

The edges in L are called *light of level 1*. Let E be the total set of edges defined by S , and let C_1 collect all edges of E that are crossed by some edge in L . Notice that each edge in L , and therefore no edge in C_1 , appears in $GT(S)$. Define $E_2 = E \setminus (L \cup C_1)$. An edge $e \in E_2$ is called *light of level 2* if e is not crossed by a shorter edge in E_2 . Let L_2 be the set of all edges which are light of level 2, and let C_2 collect all edges of E_2 that are crossed by some edge in L_2 . Again, each edge in L_2 , and therefore no edge in C_2 , appears in $GT(S)$. By setting $E_3 = E_2 \setminus (L_2 \cup C_2)$ we now can define, in the obvious way, the set L_3 of edges which are light of level 3 (see figure 1). Repeating this process until $E_{k+1} = \emptyset$ yields a hierarchy of levels L_1, L_2, \dots, L_k with $L_1 = L$.

It is evident that levels are pairwise disjoint, and that no edge of level i can cross an edge of level j , for $1 \leq i, j \leq k$. More specifically, we have:

Lemma 3. $GT(S) = L_1 \cup L_2 \cup \dots \cup L_k$.

Lemma 3 gives an alternative, though still procedural, definition of $GT(S)$ which is of interest in its own right. It provides a more structured view of $GT(S)$ than does the original definition and does not require the sorted order of the edges spanned by S . Clearly, a level decomposition of $GT(S)$ can be computed in polynomial time. We leave open the question

whether the level decomposition proves useful in the design of new and efficient greedy triangulation algorithms.²

From Lemma 2 we learn that a GT with a single level is in fact a MWT. This suggests the conjecture that, if $GT(S)$ has only a few levels, it should not decline in weight too much from $MWT(S)$. We affirmatively settle this conjecture in the following section.

3. Bounding the weight of a k -level GT

The goal of this section is to establish an upper bound on the weight of $GT(S)$ that depends on the number of levels of $GT(S)$.

Theorem 1. *Let S be a finite set of points in the plane, and let k be the number of levels attained by $GT(S)$. Then $|GT(S)| \leq c_k \cdot |MWT(S)|$, where $c_1 = 1$ and $c_k = 3 \cdot 2^{k+1}$ for $k \geq 2$.*

The remainder of this section contains a proof of Theorem 1. The special case $k = 1$ follows from Lemma 2. For $k \geq 2$, the proof is mainly based on an appropriate weighting scheme for the points in S . It proceeds in three stages: (1) each point is associated with an initial weight, such that the sum of these weights can be related to the weight of $MWT(S)$ (Lemma 4); (2) point weights are updated stepwise, where each step corresponds to a level of $GT(S)$, and the increase of weight per point is controlled (Lemma 5); and (3) the sum of the final point weights is used to bound the weight of $GT(S)$ (Lemma 7).

The weight of each point $p \in S$ is obtained by assigning to p a certain star of incident edges. A star can also contain edges which are not in $GT(S)$. The initial star of p , $\sigma_0(p)$, consists of three edges such that the angles between consecutive edges are less than π , and the sum of the edge lengths is minimum. (An exception are points lying on the convex hull of S . Their initial star consists of the two respective convex hull edges instead.) Here, and during the subsequent extensions and modifications of stars, the weight of a point p is defined to be the length of the longest edge in its star. Let $w_0(p)$ be the initial weight of p .

Lemma 4. $\sum_{p \in S} w_0(p) \leq 2 \cdot |MWT(S)|$.

Proof: For any point $p \in S$, any triangulation of S has to contain edges incident to p whose total length is at least $w_0(p)$. In particular, this is true for $MWT(S)$. The factor 2 is obtained because each edge of $MWT(S)$ is counted twice in this way, once for each endpoint. \square

Let $GT(S)$ consist of k levels. For each point $p \in S$, its star is now updated during k steps. Let $\sigma_i(p)$ denote the star of p after step i , for $i = 1, \dots, k$. The following three invariants for $\sigma_i(p)$ are maintained.

- (1) $\sigma_i(p)$ contains—among possible other edges—all incident edges in $GT(S)$ which are of level at most i .
- (2) All angles between consecutive edges in $\sigma_i(p)$ are less than π .
- (3) No edge of $\sigma_i(p)$ is crossed by any edge in $GT(S)$ of level $\leq i$.

To maintain invariant (1), all incident level- i edges are added to $\sigma_{i-1}(p)$. Adding such an edge is called an *expansion* of a star. Clearly, an expansion does not destroy invariant (2).

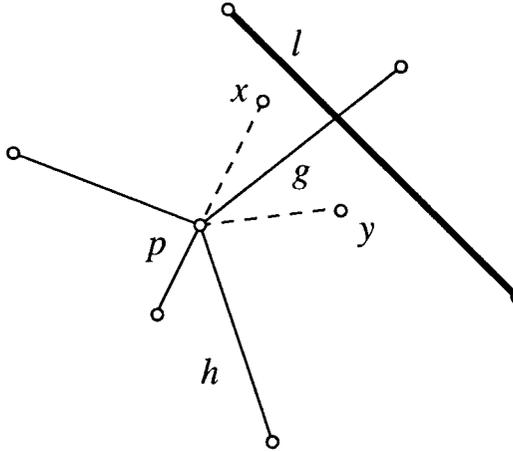


Figure 2. Modification of a star.

After having expanded the stars for all points in S , invariant (3) may be violated. Assume this is the case for point p . For each edge g in p 's star that is crossed by edges of level $\leq i$ we do the following; see figure 2. Let l be the edge of level $\leq i$ that crosses g closest to p . Consider the (topologically closed) triangle spanned by l and p , and let x and y be the points in this triangle that are hit first when g is rotated about p . (Note that x or y can be an endpoint of l .) We remove g , and add the edges px and py if they have not been part of the star yet. This action is called a *modification* of a star.

Indeed, after the modification, no edge of level $\leq i$ crosses px or py . By construction, any such edge e would have to cross g , too. As g is no edge of $GT(S)$, g was already present in $\sigma_{i-1}(p)$. Hence, by invariant (3), e cannot be of level $\leq i - 1$. It also cannot be of level i , as this would contradict the definition of l , x , and y . Notice finally that invariant (2) is maintained during a modification.

Recall that the weight of a point p after step i , $w_i(p)$, was defined to be the length of the longest edge in $\sigma_i(p)$.

Lemma 5. $w_i(p) \leq 2 \cdot w_{i-1}(p)$, for $i = 1, \dots, k$.

Proof: We argue that an expansion, as well as a modification, can produce edges of length at most twice the length of the longest edge in $\sigma_{i-1}(p)$.

We first consider an expansion of $\sigma_{i-1}(p)$. Let l be an edge of level i that is added, and let g and h be the edges adjacent to l in $\sigma_{i-1}(p)$; see figure 3. We have $|g|, |h| \leq w_{i-1}(p)$. Assume $|l| > w_{i-1}(p)$ as there is nothing to prove, otherwise. By invariant (2), l then crosses the edge qr , where q and r are the second endpoints of g and h , respectively. Consider the convex hull of all points of S that lie in the triangle pqr , including q and r but excluding p . Let e be the edge of the convex hull which is crossed first by l when coming from p . (Note that $e = qr$ is possible.) By construction, e cannot be crossed by any edge of level $\leq i - 1$ as such an edge would have to cross g or h , too, contradicting invariant (3).

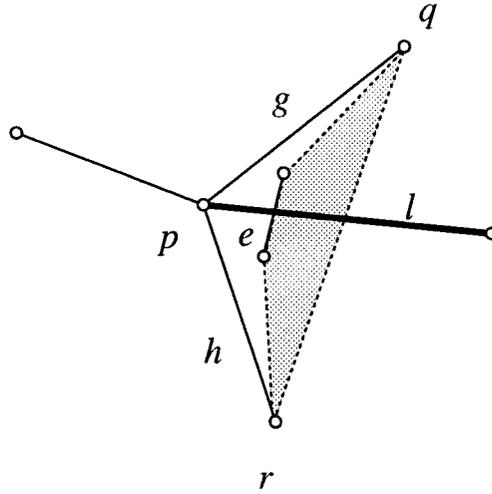


Figure 3. Expansion of a star.

Hence $|l| \leq |e|$, as e instead of l would be of level i , otherwise. Moreover, by the triangle inequality, $|e| \leq |g| + |h|$. We conclude $|l| \leq 2 \cdot w_{i-1}(p)$.

Similar arguments can be used to bound the length of edges stemming from a modification of the expanded star. Let g, l, x , and y be defined as in the modification step; cf. figure 2. Recall that g already appears in $\sigma_{i-1}(p)$. So $|g| \leq w_{i-1}(p)$. Furthermore, g cannot cross edges of level $\leq i - 1$ by invariant (3) but, by definition, it crosses the level- i edge l added in step i . This implies $|l| \leq |g|$. By construction of x and y , $\max\{|px|, |py|\} \leq |g| + |l|$. We conclude $|px|, |py| \leq 2 \cdot |g| \leq 2 \cdot w_{i-1}(p)$. \square

Finally, in order to bound the weight of $GT(S)$ by means of the weights of the points in S after step k , we utilize the following result proved in (Aichholzer et al., 1995).

Lemma 6. *Let T be an arbitrary triangulation of S . Then the edges of T can be oriented such that each point $p \in S$ has an in-degree of at most 3.*

The assertion below is now easy to prove.

Lemma 7. $|GT(S)| \leq 3 \cdot \sum_{p \in S} w_k(p)$.

Proof: According to Lemma 6, we orient the edges of $GT(S)$ such that each point $p \in S$ has at most three incident edges pointing at it. By invariant (1), these edges also have to be present in the final star $\sigma_k(p)$. Therefore, their total length is bounded by $3 \cdot w_k(p)$. Summing over all points in S gives the stated result. \square

Theorem 1, the main result of this paper, now follows from combining Lemmas 4, 5, and 7.

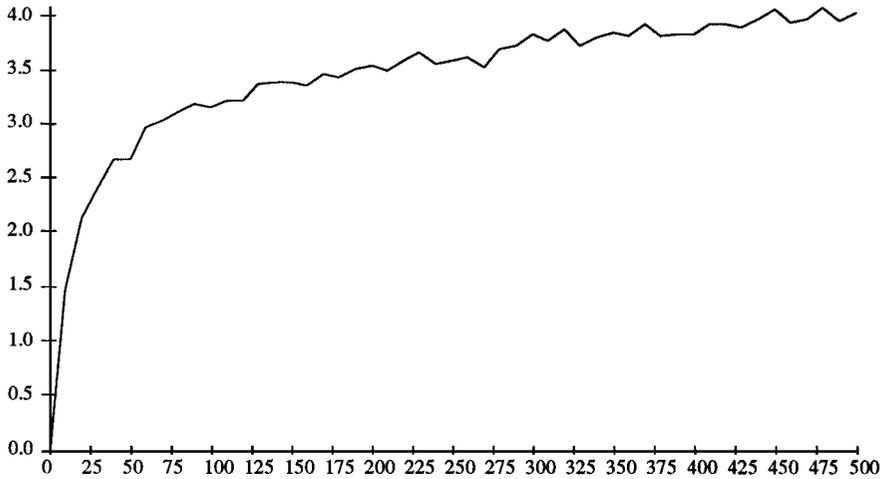


Figure 4. Expected number of levels.

4. On the number and length of levels

The quality of the bound on GT expressed in Theorem 1 depends on how a GT is structured into levels. Below we summarize several observations on the number of levels as well as on lower bounds on the weight of particular levels.

Let us first report on some experimental results. We have run a level decomposition algorithm for the GT of $n \leq 500$ points uniformly distributed in the unit square. Figure 4 displays the number k of levels in dependence of n . For each n , this number has been averaged over 100 sets of cardinality n . As can be observed, k shows a very slow increase already for moderately large n . Similar results have been obtained independently in (Jansson, 1995) in experiments with up to 1000 points.

For constant k , Theorem 1 implies that the GT is a constant approximation to the MWT. Hence, for uniformly distributed points, the observed behaviour of GT is in accordance with the theoretical result in (Levcopoulos and Lingas, 1992).

For specially constructed point sets, however, k can be up to linear in n . The points p_1, p_2, \dots, p_n in figure 5 are placed on a circle where, for $i = 2, \dots, n$, $|p_{i-1}p_i|$ increases by a fixed amount such that $|p_n p_1| = 2 \cdot |p_1 p_2|$. Aside from convex hull edges, only $p_1 p_3$ belongs to level L_1 . This is because, for $2 \leq i \leq n$, $|p_i p_j| > |p_{i-1} p_{i+1}|$ if $j \geq i + 2$. For similar reasons, only $p_3 p_5$ belongs to level L_2 . By repeating this type of argument, we see that in fact each inner edge of $GT(\{p_1, \dots, p_n\})$ constitutes a separate level. This gives a lower bound of $k \geq n - 3$.

As a curious fact, notice that $GT(\{p_1, \dots, p_n\})$ —though far from being a light triangulation—coincides with $MWT(\{p_1, \dots, p_n\})$ in this case.

From Lemma 2 the question arises whether each particular level of a GT can be bounded in weight by the corresponding MWT. (The proof of Theorem 1 shows that levels can at most double in weight in the worst case.) A result of this kind would immediately give a

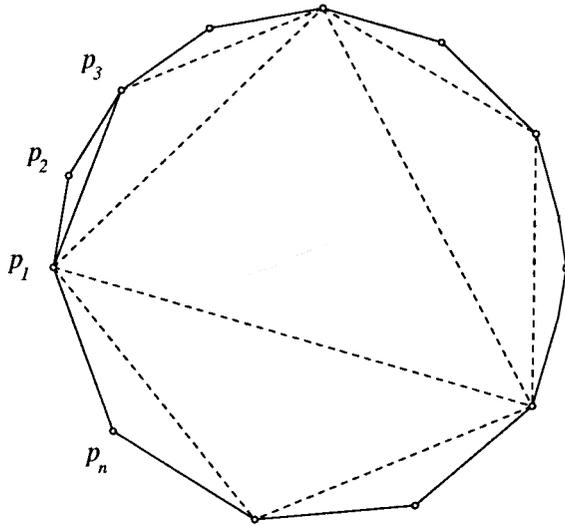


Figure 5. GT with many levels.

bound on GT that depends linearly on k rather than exponentially. Figure 6 exhibits a point set S where $|L_2| > |MWT(S)|$. The inner edge shown in solid at the left top of the picture belongs to level L_1 . This edge is shorter than all dashed edges it crosses and excludes them from $GT(S)$. Therefore, the bundle of solid edges emanating from the points on the bottom of the picture belongs to level L_2 . The dashed edges emanating from these points, however,

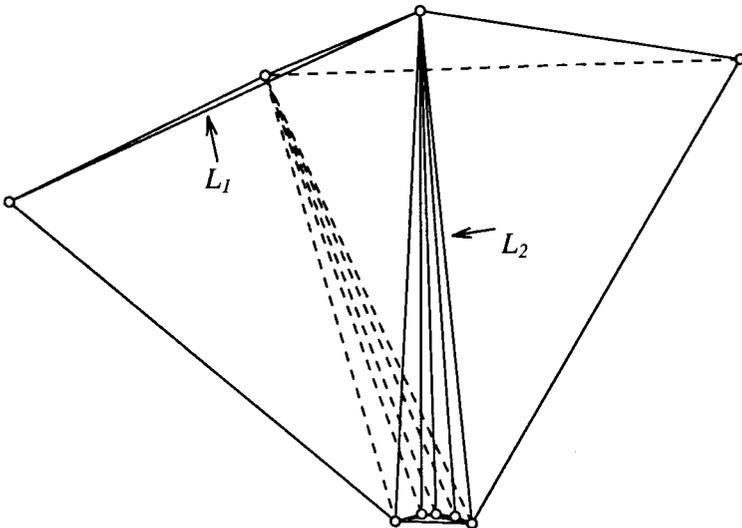


Figure 6. Level 2 (solid) exceeds the MWT (dashed).

are shorter than the solid ones and belong to $MWT(S)$. Thus, if the number of bottom points is sufficiently large, L_2 will dominate $MWT(S)$ in weight. A similar construction can be done for a higher level L_i , yielding $|L_i| > |MWT(S)|$ for a given value $i \leq k$.

Acknowledgments

Research was supported by the Spezialforschungsbereich F003, *Optimierung und Kontrolle*. Part of the work was done while Yin-Feng Xu visited the Institute for Theoretical Computer Science, Graz University of Technology.

Notes

1. By length of an edge we mean the Euclidean distance between its endpoints. Two distinct edges are said to cross if they intersect in their interiors. To ease the presentation, let us assume throughout that S is in general position and that no two edges have the same length.
2. We have recently learned that level decompositions have been used in (Jansson, 1995) for computing the greedy triangulation in parallel.

References

- O. Aichholzer, F. Aurenhammer, S.-W. Chen, N. Katoh, G. Rote, M. Taschwer, and Y.-F. Xu, "Triangulations intersect nicely," *Discrete & Computational Geometry*, vol. 16, pp. 339–359, 1996 (special issue).
- O. Aichholzer, F. Aurenhammer, and G. Rote, "Optimal graph orientation with storage applications," SFB-Report F003-51 (Optimierung und Kontrolle), TU Graz, Austria, 1995.
- M. Dickerson, R.L. Drysdale, S. McElfresh, and E. Welzl, "Fast greedy triangulation algorithms," *Proc. 10th Ann. ACM Symp., Computational Geometry*, pp. 211–220, 1994.
- J. Jansson, "Planar minimum-weight triangulations," MS Thesis, Rep. LU-CS-EX:95-16, Lund University, Sweden, 1995.
- C. Levcopoulos, "An $\Omega(\sqrt{n})$ lower bound for the nonoptimality of the greedy triangulation," *Information Processing Letters*, vol. 25, pp. 247–251, 1987.
- C. Levcopoulos and D. Krznaric, "Quasi-greedy triangulations approximating the minimum weight triangulation," *Proc. 7th ACM-SIAM Symp. Discrete Algorithms*, pp. 392–401, 1996.
- C. Levcopoulos and A. Lingas, "On approximation behaviour of the greedy triangulation for convex polygons," *Algorithmica*, vol. 2, pp. 175–193, 1987.
- C. Levcopoulos and A. Lingas, "Greedy triangulation approximates the minimum weight triangulation and can be computed in linear time in the average case," Report LU-CS-TR:92-105, Dept. of Computer Science, Lund University, 1992.