A Quasi-Polynomial Time Approximation Scheme for Minimum Weight Triangulation

[Extended Abstract]

Jan Remy
Institut für Theoretische Informatik
ETH Zürich
CH-8092 Zürich
jremy@inf.ethz.ch

Angelika Steger
Institut für Theoretische Informatik
ETH Zürich
CH-8092 Zürich
steger@inf.ethz.ch

ABSTRACT

The MINIMUM WEIGHT TRIANGULATION problem is to find a triangulation $\mathcal{T}^*$ of minimum length for a given set of points P in the Euclidean plane. It was one of the few longstanding open problems from the famous list of twelve problems with unknown complexity status, published by Garey and Johnson [8] in 1979. Very recently the problem was shown to be $\mathcal{NP}$-hard by Mulzer and Rote. In this paper, we present a quasi-polynomial time approximation scheme for MINIMUM WEIGHT TRIANGULATION.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*; G.2.2 [**Discrete Mathematics**]: Graph Theory—*Graph algorithms*

General Terms

Algorithms, Theory

Keywords

Minimum weight triangulation, approximation algorithms

1. INTRODUCTION

Let $P \subset \mathscr{R}^2$ denote a set of points in the Euclidean plane of cardinality n. A triangulation $\mathcal{T}$ of P is a collection of non-intersecting edges or straight-line segments, dividing the interior of the convex hull of P into triangular regions. The *length* of a triangulation $\mathcal{T}$ is defined as $\ell(\mathcal{T}) = \sum_{\{u,v\}\in\mathcal{T}} d(u,v)$, i.e., the total length of the triangulation in the $\mathscr{L}_2$-metric. The MINIMUM WEIGHT TRIANGULATION problem is to find a triangulation $\mathcal{T}^*$ of minimum length for P. Note that $\mathcal{T}^*$ is not necessarily unique.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
STOC'06, May 21–23, 2006, Seattle, Washington, USA.
Copyright 2006 ACM 1-59593-134-1/06/0005 ...$5.00.

The problem appeared in the famous list of 12 problems with unknown complexity which was published by Garey and Johnson in 1979 [8]. Although the problem has received much attention, only comparably little progress was made for many years. For a long time neither a polynomial time algorithm nor a proof of $\mathcal{NP}$-hardness was known. Very recently, Mulzer and Rote [16] proved that MINIMUM WEIGHT TRIANGULATION is $\mathcal{NP}$-hard. Polynomial time algorithms are known for special cases, like for instance triangulating polygonal domains [9, 11].

The approximability of the problem was also subject of previous research. Unfortunately, natural triangulations like the Delauny or the greedy triangulation can be worse then the optimum by a factor of $\Omega(n)$. The excellent survey of Bern and Eppstein [6] contains several such examples. In 1987 Plaisted and Hong [17] proposed an algorithm which produces triangulations within a factor of $\mathcal{O}(\log n)$ of an optimum triangulation. Roughly a decade later Levcopoulos and Krznaric [12, 14] proved that a suitable variant of the greedy algorithm constructs triangulations that exceed the weight of an optimum triangulation by a factor of at most c for some (rather large) constant c. Therefore MINIMUM WEIGHT TRIANGULATION is in the class $\mathcal{APX}$.

However, it remained open whether there is a $(1+\varepsilon)$-approximation algorithm for every $\varepsilon > 0$. Recall that one classifies such $(1+\varepsilon)$-approximation algorithms according to their time complexity. If the complexity is $\text{poly}(n, 1/\varepsilon)$ then the problem has a *fully-polynomial time approximation scheme (FPTAS)*. If the complexity is just $\text{poly}(n)$ for each fixed $\epsilon > 0$ (but potentially exponential in $1/\varepsilon$), then the problem is said to have a *polynomial time approximation scheme (PTAS)*. Finally, if the algorithm has time complexity $n^{\text{polylog}(n)}$ for each fixed $\epsilon > 0$ we have a *quasi-polynomial time approximation scheme (QPTAS)*.

A decade ago Arora [1, 2] and independently Mitchell [15] introduced approximation schemes for several $\mathcal{NP}$-hard geometric optimization problems. This had risen the hope for the existence of a PTAS or a QPTAS for MINIMUM WEIGHT TRIANGULATION. Indeed, Arora conjectured several times [2, 5, 3, 4] that his technique may also apply to the related MINIMUM WEIGHT STEINER TRIANGULATION problem. In this variant, the triangulation may include additional points at arbitrary locations - the Steiner points. For various technical reasons this problem seemed to be more amenable to Arora's techniques than MINIMUM WEIGHT TRIANGU-

LATION. However, no such algorithm has been found till today. Our contribution is the following.

THEOREM 1. MINIMUM WEIGHT TRIANGULATION *admits a QPTAS. The algorithm computes for every $\varepsilon > 0$ a $(1+\varepsilon)$-approximation in $n^{\mathcal{O}\left(\log^8 n\right)}$ time.*

The existence of a QPTAS implies that the MINIMUM WEIGHT TRIANGULATION problem is not $\mathcal{APX}$-hard, unless SAT $\in$ DTIME$[n^{\text{polylog}(n)}]$. This seems to be widely disbelieved. Therefore, the existence of a QPTAS can be seen as a strong indication for the existence of a PTAS. On the other hand, Mulzer and Rote [16] also show that it is $\mathcal{NP}$-hard to approximate MINIMUM WEIGHT TRIANGULATION within a relative error of $\mathcal{O}\left(1/n^2\right)$. This excludes the existence of an FPTAS, unless $\mathcal{P} = \mathcal{NP}$.

Further Related Work. If the input points are in convex position a $(1+\varepsilon)$-approximation can be computed in time $\mathcal{O}(n \log n)$ [13]. By allowing Steiner points, the weight of the triangulation can decrease. In particular, there exist instances for which the weight decreases by a factor of $\Omega(n)$ if one allows Steiner points [7]. It is not known whether this problem is $\mathcal{NP}$-complete or not. Eppstein [7] presented an algorithm for MINIMUM STEINER TRIANGULATION which computes a 316-approximation in $\mathcal{O}(n \log n)$ time.

Another related problem is MINIMUM WEIGHT PSEUDO-TRIANGULATION which asks for a pseudo-triangulation of minimum weight. A pseudo-triangulation is a planar graph on the given point set such that each face is a *pseudo-triangle*, that is a planar polygon that has exactly three convex vertices with internal angle less than π. Like the preceding problems, the complexity status of this problem is still open. Recently, Gudmundsson and Levcopoulos gave a polynomial time $4(1+4\sqrt{2})$-approximation algorithm [10]. They also proved that the minimum weight pseudo-triangulation of the interior of a simple polygon can be computed in time $\mathcal{O}\left(n^3\right)$.

2. ALGORITHMIC STRATEGY

In this section we give a high-level overview on our general approach and the necessary concepts and tools. Assume for the moment that we have chosen an appropriate square $\mathfrak{Q}_\circ$ which contains all points from P. We construct a subdivision of $\mathfrak{Q}_\circ$, by first cutting $\mathfrak{Q}_\circ$ vertically, then cutting the two resulting rectangles horizontally, then cutting each of the four resulting squares again vertically, and so on until we obtain squares of side length at most one. This subdivision can be seen as a tree with rectangles as nodes (from now on we will use the term 'rectangle' to denote a rectangle or a square). Given a rectangle $\mathfrak{R}$ encountered in this subdivision and a triangulation $\mathcal{T}$ for the point set P, there is a natural way to define the restriction of $\mathcal{T}$ to $\mathfrak{R}$. Namely, this restriction contains exactly those edges from $\mathcal{T}$ for which both of its endpoints are contained in $\mathfrak{R}$.

From a very high level, the general approach of our algorithm can then be described as follows. We compute a triangulation $\mathcal{T}$ by dynamic programming in a bottom-up fashion (note that what we outline now is not yet a quasi-polynomial time algorithm).

- First, for all leafs of the subdivision of $\mathfrak{Q}_\circ$ we enumerate all restrictions of triangulations of P to these leafs.
- Then we recursively compute all restrictions of triangulations of a rectangle $\mathfrak{R}$ by considering all pairs of restrictions to its two children $\mathfrak{R}'$ and $\mathfrak{R}''$ and combining them optimally.

The main drawback of this approach is that there are far too many different restrictions of triangulations of P to a rectangle $\mathfrak{R}$. A first important observation is that within the above approach we actually do not really need to store all restrictions. Within this bottom-up approach all we need to know is how such a restriction looks from "outside", that is which edges of the restriction can still be connected by straight lines to points outside of $\mathfrak{R}$. This is what we will call the *local hull* of a restriction. A precise definition is deferred to Section 3.3. There we will also show properties of local hulls that will allow us to efficiently implement the recursion step. More precisely, we will show that given a rectangle $\mathfrak{R}$, its two children $\mathfrak{R}'$ and $\mathfrak{R}''$, a local hull $\mathcal{H}$ for $\mathfrak{R}$ and local hulls $\mathcal{H}'$ and $\mathcal{H}''$ for $\mathfrak{R}'$ and $\mathfrak{R}''$ together with triangulations of the interiors of the hulls $\mathcal{H}'$ and $\mathcal{H}''$, we can check in polynomial time whether the hulls $\mathcal{H}'$ and $\mathcal{H}''$ are 'consistent' with $\mathcal{H}$ and, if so, compute a triangulation of the interior of $\mathcal{H}$ that is an *optimum* continuation of the given triangulations for $\mathcal{H}'$ and $\mathcal{H}''$.

Note that, using these ideas the time complexity of the above approach is reduced polynomially to (a) the number of different local hulls we have to consider per rectangle and (b) the computation of triangulations of the interior of the local hulls of the *leafs* of the subdivision of $\mathfrak{Q}_\circ$. Let us first consider issue (b). Computing such triangulations optimally seems difficult, as it is essentially equivalent to the original problem of computing an optimum triangulation for P. However, as we will see in Section 4, the sizes of leaf rectangles can be made sufficiently small so that *any* triangulation of the interior of a given local hull for a leaf rectangle will suffice to get a $(1+\epsilon)$-approximation of the optimum triangulation for the point set P. It thus remains to consider issue (a). This is more difficult, as it is easy to construct examples for which there exist exponentially many different local hulls. Our approach will therefore be as follows. We define a polynomial time algorithm $\mathcal{L}(\mathfrak{Q}_\circ, P)$ that, given a square $\mathfrak{Q}_\circ$ and a point set P of size $|P| = n$, constructs for each rectangle $\mathfrak{R}$ of the subdivision of $\mathfrak{Q}_\circ$ a set $\mathbf{H}_{\mathcal{L}}(\mathfrak{R})$ of different local hulls such that $|\mathbf{H}_{\mathcal{L}}(\mathfrak{R})| = 2^{\text{polylog}(n)}$ for all rectangles $\mathfrak{R}$. Then the scheme which we outlined above will allow us to compute in time $2^{\text{polylog}(n)}$ a triangulation $\mathcal{T}_{\mathcal{L}}$ such that

$$\ell(\mathcal{T}_{\mathcal{L}}) \;\leq\; (1+\varepsilon) \cdot \min_{\mathcal{T} \in \mathcal{T}_{\mathcal{L}}(\mathfrak{Q}_\circ, P)} \ell(\mathcal{T}),$$

where $\mathcal{T}_{\mathcal{L}}(\mathfrak{Q}_\circ, P)$ denotes the set of all triangulations $\mathcal{T}$ such that for all rectangles $\mathfrak{R}$ the local hull of the restriction of $\mathcal{T}$ to $\mathfrak{R}$ is contained in $\mathbf{H}_{\mathcal{L}}(\mathfrak{R})$. It remains to be shown that we can define a set of polynomially many different squares $\mathfrak{Q}_i$ such that for at least one $\mathfrak{Q}_i$ we have

$$\min_{\mathcal{T} \in \mathcal{T}_{\mathcal{L}}(\mathfrak{Q}_i, P)} \ell(\mathcal{T}) \;\leq\; (1+\varepsilon) \cdot \ell(\mathcal{T}^*),$$

where $\mathcal{T}^*$ denotes an *optimum* triangulation for the given point set P. As $(1+\varepsilon)^2 \leq 1+3\epsilon$ for all $0 < \varepsilon \leq 1$ this algorithm therefore produces a $(1+3\varepsilon)$-approximation to $\mathcal{T}^*$.

Outline. This paper is organized as follows. In Section 3 we explain the subdivision of $\mathfrak{Q}_\circ$ in more detail. Moreover, we study fundamental properties of triangulations with respect to this subdivision. Among others we give a precise definition of the local hull. In Section 4 we introduce our QPTAS which is based on the concepts of Section 3. The remainder of the paper outlines the analysis of the algorithm. In Section 5 we define the algorithm $\mathcal{L}$ that computes the set of local hulls $(\mathbf{H}_{\mathcal{L}}(\mathfrak{R}))_{\mathfrak{R}}$. Due to space restrictions we only explain the main ideas behind the analysis. This is done in Section 6 and Section 7.

3. PRELIMINARIES

3.1 Notation and Conventions

A *planar straight-line graph (PSLG)* $\mathcal{G}$ on a point set P is a set of pairwise non-intersecting straight-line segments or *edges* with endpoints in P. For instance, a triangulation and all its subgraphs are PSLGs. For an edge $\{u, v\} \in \mathcal{G}$ we sometimes denote the corresponding line segment by uv. More precisely, uv denotes the set of points on the line segment connecting u and v. Our convention is that a PSLG is a set of edges. Only if the context is not clear, we use the notation $E(\mathcal{G})$ to refer to the edge set of $\mathcal{G}$. Similarly, $P(\mathcal{G})$ denotes the points of P that are incident to at least one edge in $E(\mathcal{G})$.

As usual, $\triangle_{xyz}$ denotes the triangle with vertices x, y, and z. Therefore triangles induce a PSLG and we use the notations $E(\triangle)$ and $P(\triangle)$ to refer to $\triangle$'s sides and vertices, respectively. Above we mentioned that a triangulation $\mathcal{T}$ can be seen as PSLG. For the sake of brevity we write $e \in \mathcal{T}$ instead of $e \in E(\mathcal{T})$ unless the context is unclear. Throughout, $\Delta(\mathcal{T})$ denotes the set of triangles (or triangular regions) induced by $\mathcal{T}$.

The halfspaces of a line H are denoted by H^+ and H^-, respectively. Our convention is that halfspaces are always open. We will not use any kind of global orientation, but instead we will define H^+ and H^- locally by specifying specific points that they contain. For any region $R \subset \mathscr{R}^2$, ∂R and $\overline{R}$ denotes its boundary and closure, respectively. Unless stated otherwise, $\log n$ denotes the logarithm of n to base 2.

3.2 Rectangles and their Generators

In the following we denote by $\mathfrak{Q}_\circ$ a square that contains all points from a given point set P. We recursively construct a subdivision of $\mathfrak{Q}_\circ$ that is similar to an ordinary quad tree as follows. A given rectangle $\mathfrak{R}$ is subdivided into two (!) *children* $\mathfrak{R}'$ and $\mathfrak{R}''$ by a *separator* C through the barycenter of $\mathfrak{R}$. The separator is vertical if $\mathfrak{R}$ is a square and horizontal otherwise. We stop the subdivision if $\mathfrak{R}$ is a square of side length at most 1. For definiteness let us note here, that we will later ensure that we only consider such squares $\mathfrak{Q}_\circ$ for which *no* point from P is contained on the *boundary* of a rectangle $\mathfrak{R}$.

A rectangle $\mathfrak{R}$ is obtained from $\mathfrak{Q}_\circ$ by a sequence of separators, say $C_0, \ldots, C_l$. We call l the *level of* $\mathfrak{R}$, denoted by $\mathrm{lev}(\mathfrak{R})$. A separator C has *level* l if it subdivides a rectangle at level $l-1$ into two rectangles at level l. By definition $\mathfrak{Q}_\circ$ has level -1. One easily checks inductively that a rectangle at level l has maximum side length at most $|\mathfrak{Q}_\circ|/2^{\lceil l/2 \rceil}$. Observe that, by construction, the height of this subdivision is thus

$$t := 2\lceil \log_2 |\mathfrak{Q}_\circ| \rceil,$$

where $|\mathfrak{Q}_\circ|$ denotes the side length of the square $\mathfrak{Q}_\circ$.

3.3 Local Hulls

We now study the so-called local hull, one of the key concepts of our paper. Local hulls are defined for a rectangle $\mathfrak{R}$ with respect to a triangulation $\mathcal{T}$.

Definition 1. Let $\mathfrak{R}$ be a rectangle and let $\mathcal{T}$ be a triangulation for a point set P. The *local hull* $\mathcal{H}(\mathfrak{R}, \mathcal{T})$ of $\mathfrak{R}$ with respect to $\mathcal{T}$ is defined as the set of edges in $\mathcal{T}$ for which both endpoints belong to $\mathfrak{R}$ and that satisfy in addition at least one of the following two properties: *(i)* the edge belongs to the convex hull $\mathrm{conv}(P)$ of the given point set, or *(ii)* the edge is a side of a triangle in $\mathcal{T}$ that has its third vertex outside of $\mathfrak{R}$.

In other words the local hull consists of all edges that lay entirely in $\mathfrak{R}$ and that either belong to $\mathrm{conv}(P)$ or are sides of triangles that cross the boundary of $\mathfrak{R}$. An example is given in Figure 1. As a local hull is defined as a set of edges, we can view it as a planar straight-line graph. The following lemma shows that, viewed as a planar graph, local hulls are quite simple objects in which all cycles correspond to faces.

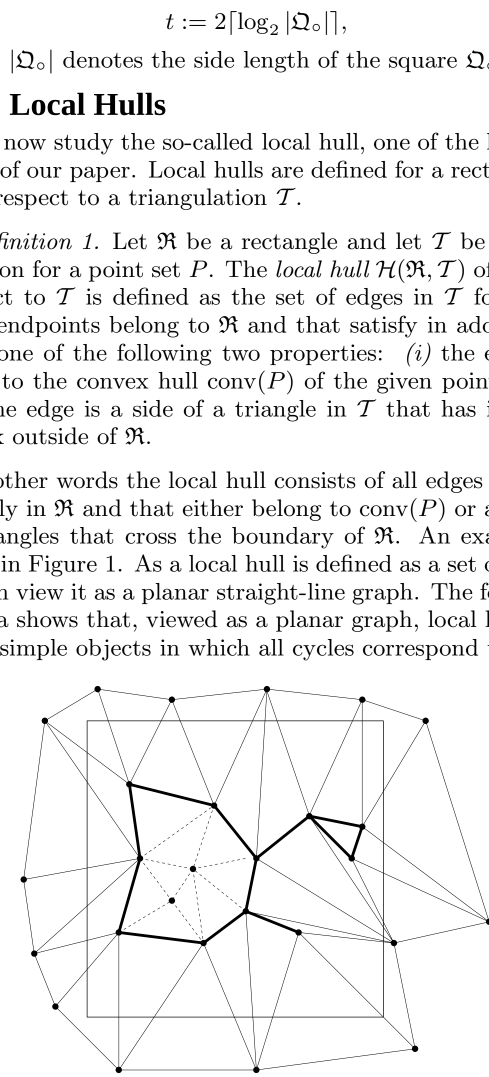

Figure 1: The local hull consists of the bold drawn edges.

LEMMA 1. *For all triangulations $\mathcal{T}$ and rectangles $\mathfrak{R}$, the local hull $\mathcal{H}(\mathfrak{R}, \mathcal{T})$ has the following property. If $\mathcal{C}$ is a cycle in $\mathcal{H}(\mathfrak{R}, \mathcal{T})$, then $\mathcal{C}$ induces a face in $\mathcal{H}(\mathfrak{R}, \mathcal{T})$.*

The *interior of* $\mathcal{H}(\mathfrak{R}, \mathcal{T})$ is defined as the union of the interior regions of all cycles contained in $\mathcal{H}(\mathfrak{R}, \mathcal{T})$ and is denoted by $\mathrm{int}(\mathcal{H}(\mathfrak{R}, \mathcal{T}))$. Note that by the previous lemma, two cycles of $\mathcal{H}(\mathfrak{R}, \mathcal{T})$ may overlap in a point, but not in an edge (as this would result in a cycle that contains this edge in its interior region). Hence we know that $\mathrm{int}(\mathcal{H}(\mathfrak{R}, \mathcal{T}))$ can be written as the union of simple polygons, such that any two of them overlap in at most one point. Throughout we use

$$\mathrm{area}(\mathcal{H}(\mathfrak{R}, \mathcal{T})) = \mathrm{int}(\mathcal{H}(\mathfrak{R}, \mathcal{T})) \cup \mathcal{H}(\mathfrak{R}, \mathcal{T}).$$

Note that $\mathrm{area}(\mathcal{H}(\mathfrak{R}, \mathcal{T}))$ contains all triangles of $\mathcal{T}$ that are entirely contained in $\mathfrak{R}$. Moreover, the local hull is in general not connected. Observe also that we have $\mathcal{H}(\mathfrak{Q}_\circ, \mathcal{T}) = \mathrm{conv}(P)$.

3.4 Border Triangulations

Within the recursive scheme outlined in Section 2 it is essential to understand how local hulls are related to the local hulls of their descendants. In this section we develop the necessary background. So let $\mathfrak{R}$ be a rectangle that is divided by a separator C into two rectangles $\mathfrak{R}'$ and $\mathfrak{R}''$. Furthermore, let $\mathcal{T}$ denote a triangulation for the point set P. We consider three different sets. Namely, the set of triangles within $\mathfrak{R}$, that cross the cut, that is,

$$\Delta_{\mathrm{B}}(\mathfrak{R},\mathcal{T}) := \{\triangle \in \Delta(\mathcal{T}) : (P(\triangle) \subset \mathfrak{R}) \wedge (P(\triangle) \cap \mathfrak{R}' \neq \emptyset) \\ \wedge (P(\triangle) \cap \mathfrak{R}'' \neq \emptyset)\}$$

and the set of edges in triangles from $\Delta_{\mathrm{B}}(\mathfrak{R},\mathcal{T})$ that cross the cut and those that are contained within one of the smaller rectangles. That is

$$\mathcal{C}(\mathfrak{R},\mathcal{T}) := \bigcup_{\triangle \in \Delta_{\mathrm{B}}(\mathfrak{R},\mathcal{T})} \{e \in E(\triangle) : e \cap \mathfrak{R}' \neq \emptyset \\ \wedge e \cap \mathfrak{R}'' \neq \emptyset\}$$

and

$$\mathcal{B}(\mathfrak{R},\mathcal{T}) := \bigcup_{\triangle \in \Delta_{\mathrm{B}}(\mathfrak{R},\mathcal{T})} \left\{e \in E(\triangle) : e \subset \mathfrak{R}' \vee e \subset \mathfrak{R}''\right\}.$$

We call $\mathcal{B}(\mathfrak{R},\mathcal{T})$ the *border* of $\mathfrak{R}$ with respect to its separator C and the triangulation $\mathcal{T}$. Similarly, we call $\mathcal{C}(\mathfrak{R},\mathcal{T})$ the *cut edges* of $\mathcal{T}$ in $\mathfrak{R}$. Figure 2 illustrates these definitions. The following lemma is an immediate consequence of the definition of the local hulls and the set of border edges.

LEMMA 2. *For all rectangles $\mathfrak{R} = \mathfrak{R}' \cup \mathfrak{R}''$, and triangulations $\mathcal{T}$ we have*

$$\mathcal{H}(\mathfrak{R}',\mathcal{T}) \;=\; \left\{e \in \mathcal{H}(\mathfrak{R},\mathcal{T}) \cup \mathcal{B}(\mathfrak{R},\mathcal{T}) : e \subseteq \mathfrak{R}'\right\}$$

and similarly for $\mathcal{H}(\mathfrak{R}'',\mathcal{T})$.

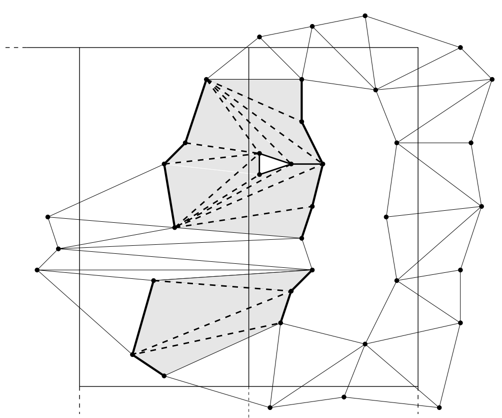

Figure 2: The border triangulation at a separator C, that subdivides $\mathfrak{R}$. The triangles belonging to $\Delta_{\mathrm{B}}(\mathfrak{R},\mathcal{T})$ are drawn shaded; the edges in $\mathcal{B}(\mathfrak{R},\mathcal{T})$ are drawn by bold lines, while those in $\mathcal{C}(\mathfrak{R},\mathcal{T})$ are drawn by dashed bold lines.

LEMMA 3. *Let $\mathfrak{R} = \mathfrak{R}' \cup \mathfrak{R}''$ and $\mathcal{T}$ a triangulation of P and let $R_1,\ldots,R_k$ denote the regions given by*

$$\mathrm{int}(\mathcal{H}(\mathfrak{R},\mathcal{T})) \setminus \left[\mathrm{area}(\mathcal{H}(\mathfrak{R}',\mathcal{T})) \cup \mathrm{area}(\mathcal{H}(\mathfrak{R}'',\mathcal{T}))\right]. \quad (1)$$

For all $i = 1,\ldots,k$ it holds that $int(R_i) \cap P = \emptyset$ and that R_i is bounded by a connected subgraph $\mathcal{W}_i$ of $\mathcal{T}$.

In order to better understand the meaning of Lemma 3 have another look at Figure 2. In this example we have just two regions R_1 and R_2, namely the two regions which are shaded in gray. Note that the interior of both of these two gray regions does not contain any point from P. Note also that even though the upper of the two regions seems to contain a region from $\mathrm{int}(\mathfrak{R}'')$ in its interior (i.e., the small white triangle), this is actually not true, as this triangle is connected by a bold line, i.e., an edge from $\mathcal{B}(\mathfrak{R},\mathcal{T})$, to the boundary of R_i. Lemma 3 says that this is always the case.

The importance of Lemma 3 stems from the fact that a minimum weight triangulation of the interior of a simple polygon is known to be computable in cubic time [9, 11] if it contains no points. The algorithm is a straight-forward dynamic program and it is well-known that it can also be used to triangulate a face of a PSLG. That is, we have

LEMMA 4. *Let R_i denote a region induced by (1) bounded by $\mathcal{W}_i$. The interior of $\mathcal{W}_i$ can be triangulated optimally in time $\mathcal{O}\left(|\mathcal{W}_i|^3\right)$.*

Coming back to the situation of Lemma 3 we thus see that if we denote by $n(\mathfrak{R}) = |\mathfrak{R} \cap P|$ the number of points in $\mathfrak{R}$ then we can find in time $\mathcal{O}(n(\mathfrak{R})^3)$ a minimum triangulation for all regions $R_1,\ldots,R_k$. This fact is one of the key features of local hulls that we will use in the next section.

4. THE APPROXIMATION SCHEME

In the previous section we defined the local hull of a rectangle with respect to a fixed triangulation $\mathcal{T}$. The aim of this section is to show that vice-versa we are also able to compute an (almost) optimal triangulation based on the knowledge of all local-hulls. For the remainder of this section we consider an arbitrary but fixed point set P such that $|P| = n$. We assume that the maximum distance between two points of P is at least sn for some $s \in \mathbb{N}$. That is, we have

$$\max_{p,q \in P} d(p,q) \geq s \cdot n.$$

Note that this can always be achieved by scaling the point set appropriately. Later we choose $s = \mathcal{O}\left(1/\varepsilon\right)$. Let $\mathfrak{Q}_\circ$ denote an arbitrary square of side length at most $6sn$ that contains all points of P. For definiteness, we assume that no point from P is contained on the boundary of any rectangle of the subdivision of $\mathfrak{Q}_\circ$ We show in Section 7 how to construct such a $\mathfrak{Q}_\circ$ appropriately. Let $\mathbf{H}_\mathfrak{R}$ denote the set of all possible local hulls of $\mathfrak{R}$, that is

$$\mathbf{H}[\mathfrak{R}] = \{\mathcal{H}(\mathfrak{R},\mathcal{T}) : \mathcal{T} \text{ is a triangulation of } P\}. \quad (2)$$

The idea behind our algorithm is to compute a triangulation $\mathcal{T}$ recursively in a bottom-up fashion. More precisely, we will recursively compute for each rectangle $\mathfrak{R}$ and each local hull $\mathcal{H} \in \mathbf{H}[\mathfrak{R}]$ a triangulation of $\mathrm{area}(\mathcal{H})$ which we store in an array $T[\mathfrak{R},\mathcal{H}]$. Our convention is that these triangulations contain all edges of $\mathcal{H}$. Recall that for $\mathfrak{R} = \mathfrak{Q}_\circ$ there exists only one local hull, namely $\mathcal{H} = \mathrm{conv}(P)$ and $T[\mathfrak{Q}_\circ, \mathrm{conv}(P)]$ will therefore contain a proper triangulation of the whole point set P. It remains to specify how we compute the triangulations $T[\mathfrak{R}, \mathcal{H}]$. If $\mathfrak{R}$ is a leaf, we triangulate the interior of $\mathcal{H}$ arbitrarily, e.g. using a greedy algorithm and store the triangulation in $T[\mathfrak{R}, \mathcal{H}]$.

Next assume that $\mathfrak{R}$ is a rectangle that contains a separator C that divides $\mathfrak{R}$ into two rectangles $\mathfrak{R}'$ and $\mathfrak{R}''$. The idea is to compute $T[\mathfrak{R}, \mathcal{H}]$ by considering all pairs of local hulls $(\mathcal{H}', \mathcal{H}'')$ in $\mathbf{H}[\mathfrak{R}'] \times \mathbf{H}[\mathfrak{R}'']$. For a given pair $(\mathcal{H}', \mathcal{H}'')$ we first check whether $\mathcal{H}'$ and $\mathcal{H}''$ are consistent with $\mathcal{H}$. Firstly, we check whether the edges of $(\mathcal{H}' \cup \mathcal{H}'') \setminus \mathcal{H}$ are contained in area$(\mathcal{H})$. In other words, we require that $\mathcal{H} \cup \mathcal{H}' \cup \mathcal{H}''$ looks from "outside" as $\mathcal{H}$. Secondly, we require that the regions obtained by (1) satisfy the properties stated in Lemma 3.

Finally, we use Lemma 4 in order to compute minimum triangulations for the regions $R_1, \ldots, R_k$ in polynomial time. Combined with the triangulations $T[\mathfrak{R}, \mathcal{H}']$ and $T[\mathfrak{R}, \mathcal{H}'']$ this yields a triangulation of the interior of $\mathcal{H}$. Let $\mathcal{T}(\mathcal{H}', \mathcal{H}'')$ denote this triangulation. From all triangulations $\mathcal{T}(\mathcal{H}', \mathcal{H}'')$ such that $(\mathcal{H}', \mathcal{H}'') \in \mathbf{H}[\mathfrak{R}'] \times \mathbf{H}[\mathfrak{R}'']$, we choose one with minimum length and store it in $T[\mathfrak{R}, \mathcal{H}]$. We claim that by construction the triangulations stored in $T[\mathfrak{R}, \mathcal{H}]$ satisfy the following property.

Lemma 5. *For all rectangles $\mathfrak{R}$ we have*

$$\ell\left(T[\mathfrak{R}, \mathcal{H}]\right) \leq \sum_{e \in \mathcal{T}^*(\mathcal{H}) \cap \mathfrak{R}} \ell(e) + 6n(\mathfrak{R}),$$

where $n(\mathfrak{R})$ denotes the number of points of P in $\mathfrak{R}$ and $\mathcal{T}^*(\mathcal{H})$ denotes a minimum triangulation for P that contains $\mathcal{H}$.

For $\mathfrak{R} = \mathfrak{Q}_\circ$ there exists only one local hull, namely $\mathcal{H} = \text{conv}(P)$. That is, we have $\mathbf{H}[\mathfrak{Q}_\circ] = \{\text{conv}(P)\}$, and every triangulation of area(conv(P)) corresponds thus to a triangulation of P. Lemma 5 thus implies

$$\ell\left(T[\mathfrak{Q}_\circ, \text{conv}(P)]\right) \leq \ell(\mathcal{T}^*) + 6n.$$

As we did assume that there are points in P which have distance at least $s \cdot n$, we know that $\ell(\mathcal{T}^*) \geq s \cdot n$ and we hence have

$$\ell\left(T[\mathfrak{Q}_\circ, \text{conv}(P)]\right) \leq (1 + 6/s) \cdot \ell(\mathcal{T}^*).$$

That is, $T[\mathfrak{Q}_\circ, \text{conv}(P)]$ is a $(1+\varepsilon)$-approximation, if we choose $s \geq 6/\varepsilon$.

Of course, the time complexity of this algorithm will in general not be very good, as firstly we have to compute the sets of local hulls $\mathbf{H}[\mathfrak{R}]$ (and it is a priori not clear how we can do that efficiently) and secondly the time complexity of the recursion can only be bounded by a polynomial in $\max_{\mathfrak{R}} |\mathbf{H}[\mathfrak{R}]|$ (which might be exponential in n). The key observation now is the following. We observe that the only point where we used that the sets $\mathbf{H}[\mathfrak{R}]$ were defined as in equation (2) was within the proof of Lemma 5. If we allow $\mathbf{H}[\mathfrak{R}]$ to denote arbitrary sets of local hulls, i.e. if we replace the "=" sign in equation (2) by a "$\subseteq$" sign everything goes through as before, except that we have to replace $\mathcal{T}^*(\mathcal{H})$ in the statement of Lemma 5 by $\mathcal{T}^*_{\mathbf{H}}(\mathcal{H})$ where $\mathcal{T}^*_{\mathbf{H}}(\mathcal{H})$ denotes a minimum triangulation of P such that

i) $\mathcal{H} \subseteq \mathcal{T}^*_{\mathbf{H}}(\mathcal{H})$ and,

ii) $\mathcal{H}(\mathfrak{R}, \mathcal{T}^*_{\mathbf{H}}(\mathcal{H})) \in \mathbf{H}[\mathfrak{R}]$ for all rectangles $\mathfrak{R}$.

In other words: A family of local hulls $(\mathbf{H}[\mathfrak{R}])_{\mathfrak{R}}$ induces a set of triangulations such that each element of this set has the property that the local hull of every rectangle $\mathfrak{R}$ is contained in $\mathbf{H}[\mathfrak{R}]$. Then $\mathcal{T}^*_{\mathbf{H}}(\mathcal{H})$ denotes the best such triangulation containing $\mathcal{H}$. Clearly, $\mathcal{T}^*_{\mathbf{H}}(\text{conv}(P))$ is a minimum amongst all triangulations in this set. We thus obtain the following theorem.

Theorem 2. *Let P be a point set and $\mathfrak{Q}_\circ$ be given as stated in the beginning of this section. Assume furthermore that $\mathcal{L}$ is an algorithm that computes for P and $\mathfrak{Q}_\circ$ a family of local hulls $(\mathbf{H}_{\mathcal{L}}[\mathfrak{R}])_{\mathfrak{R}}$ such that*

$$\max_{\mathfrak{R}} |\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]| = f(n)$$

for some function $f(n) \geq 1$. Then we can compute in time $\mathcal{O}\left(\text{poly}(n) \cdot f(n)^2\right)$ a triangulation $\mathcal{T}$ such that

$$\ell(\mathcal{T}) \;\leq\; (1 + 6/s) \cdot \ell\left(\mathcal{T}^*_{\mathcal{L}}\right),$$

where $\mathcal{T}^*_{\mathcal{L}}$ is a minimum triangulation such that $\mathcal{H}(\mathfrak{R}, \mathcal{T}^*_{\mathcal{L}}) \in \mathbf{H}_{\mathcal{L}}[\mathfrak{R}]$ for all $\mathfrak{R}$.

In essence, this theorem says that it suffices to choose the family $(\mathbf{H}_{\mathcal{L}}[\mathfrak{R}])_{\mathfrak{R}}$ such that firstly $f(n)$ is quasi-polynomial and secondly $\mathcal{T}^*_{\mathcal{L}}$ is nearly optimal. This intuition is reflected in the organization of the paper. In Section 5 we will describe how one can define an appropriate algorithm $\mathcal{L}$ for computing $\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]$. One can show that

$$|\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]| = n^{\mathcal{O}\left(\log^8 n\right)}.$$

Therefore, the running time of our algorithm is quasi-polynomial in n. In Section 7 we explain how to choose $\mathfrak{Q}_\circ$ such that

$$\ell(\mathcal{T}^*_{\mathcal{L}}) \leq (1 + \mathcal{O}\left(1/s\right)) \cdot \ell(\mathcal{T}^*).$$

Combining both results and choosing $s = \mathcal{O}\left(1/\varepsilon\right)$ we obtain an $(1+\varepsilon)$-approximation algorithm with time complexity $n^{\mathcal{O}\left(\log^8 n\right)}$.

Remark 1. If one can show that there exists a family $(\mathbf{H}_{\mathcal{L}}[\mathfrak{R}])_{\mathfrak{R}}$ with the aforesaid properties such that $f(n) = \text{poly}(n)$, then Theorem 2 implies the existence of a PTAS.

5. SMOOTH HULLS

In this section we construct the algorithm $\mathcal{L}$ and show that $\mathcal{L}$ computes $\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]$ in quasi-polynomial time. This requires some preliminary work. Our strategy is the following. We consider a fixed triangulation $\mathcal{T}$ and rectangle $\mathfrak{R}$. The idea is to extract several features from $\mathcal{T}$ which can be used to describe the local hull of $\mathfrak{R}$ up to some level of detail. Based on these features we construct the local hulls in $\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]$. Therefore the set of features should be comparably small.

First, we define a small sample of the border triangulation. However, this sample turns out to be dense enough such that we only miss very small triangles. So, this sample provides portions of the local hull with the exception of small "gaps". Therefore, our next goal is to derive several properties of those gaps. We then show that it is possible to roughly reconstruct the local hulls within the gaps based only on a small amount of additional information. This finally yields a set of local hulls for $\mathfrak{R}$ which we call *smooth hulls*.

5.1 Characteristic Triangles

Let $s = \mathcal{O}(1/\varepsilon)$ as defined in the previous section and let t denote the depth of the subdivision of a square $\mathfrak{Q}_\circ$. We also still assume that $|\mathfrak{Q}_\circ| \leq 6sn$ and thus $t = \mathcal{O}(\log n)$. Our first step is to define a suitable small sample of the border triangulation. We consider a rectangle $\widehat{\mathfrak{R}}$ that is subdivided by C into its children $\mathfrak{R}$ and $\mathfrak{R}'$. Note that in comparison to the previous sections we slightly changed the notation. As we will mostly work with one of the children, it makes the notation more concise if we call this rectangle $\mathfrak{R}$. Fix any triangulation $\mathcal{T}$ and let $\widehat{\mathcal{H}}$ denote the local hull of $\widehat{\mathfrak{R}}$ with respect to $\mathcal{T}$.

Let $m = 2^{\lceil \log(st) \rceil}$. We partition both $\mathfrak{R}$ and $\mathfrak{R}'$ into at most m^2 *cells* using a regular grid of granularity r/m, where r denotes the maximum side length of $\mathfrak{R}$. Let $\mathcal{M}[\mathfrak{R}]$ and $\mathcal{M}[\mathfrak{R}']$ denote the set of cells in $\mathfrak{R}$ and $\mathfrak{R}'$, respectively. We assume that $\mathfrak{Q}_\circ$ is chosen such that no point of P falls on the boundary of such a cell. Throughout c_l denotes the side length of cells at level l. One can easily see that c_l decreases exponentially with l.

Now, consider a tuple $(\mathfrak{C}, \mathfrak{C}')$ with $\mathfrak{C} \in \mathcal{M}[\mathfrak{R}]$ and $\mathfrak{C}' \in \mathcal{M}[\mathfrak{R}']$ and let

$$\mathcal{T}(\mathfrak{C}, \mathfrak{C}') := \{\triangle \in \Delta_{\mathrm{B}}(\widehat{\mathfrak{R}}, \mathcal{T}) : P(\triangle) \cap \mathfrak{C} \neq \emptyset \\ \wedge P(\triangle) \cap \mathfrak{C}' \neq \emptyset\}.$$

That is, $\mathcal{T}(\mathfrak{C}, \mathfrak{C}')$ is the set of triangles in $\Delta_{\mathrm{B}}(\widehat{\mathfrak{R}}, \mathcal{T})$ that have at least one vertex in each of $\mathfrak{C}$ and $\mathfrak{C}'$. We walk along the line through C from left to right if C is horizontal and from top to bottom if C is vertical. From now on we always assume that a separator is oriented in this way. By $\triangle_{\mathrm{first}}(\mathfrak{C}, \mathfrak{C}')$ we denote the first triangle of $\mathcal{T}(\mathfrak{C}, \mathfrak{C}')$ we transverse on this walk. Similarly, $\triangle_{\mathrm{last}}(\mathfrak{C}, \mathfrak{C}')$ denotes the last triangle, we transverse.

Proceeding similarly for all tuples $(\mathfrak{C}, \mathfrak{C}')$, we obtain a set of triangles

$$\mathcal{T}_{\mathrm{char}}(C, \mathcal{T}) := \{\triangle_{\mathrm{first}}(\mathfrak{C}, \mathfrak{C}'), \triangle_{\mathrm{last}}(\mathfrak{C}, \mathfrak{C}') : \\ \mathfrak{C} \in \mathcal{M}[\mathfrak{R}], \mathfrak{C}' \in \mathcal{M}[\mathfrak{R}']\}$$

which we call the *characteristic triangles of* $\Delta_{\mathrm{B}}(C, \mathcal{T})$. By the definition of m we immediately obtain

LEMMA 6. $|\mathcal{T}_{\mathrm{char}}(C, \mathcal{T})| = \mathcal{O}(\log^4 n)$.

5.2 Gaps

Using the same notation as above, we next study fundamental properties of $\mathcal{T}_{\mathrm{char}} = \mathcal{T}_{\mathrm{char}}(C, \mathcal{T})$. As mentioned above $\mathcal{T}_{\mathrm{char}}$ is in general just some sparse sample of the corresponding border triangulation. So, we first discuss which portions of the border triangulation are not contained in $\mathcal{T}_{\mathrm{char}}$. Let E_{cut} denote the set of edges that are cut by C and are either sides of triangles in $\mathcal{T}_{\mathrm{char}}$ or edges of $\widehat{\mathcal{H}}$. Along C we index the edges in E_{cut} in the order in which we transverse them according to the orientation of C defined above. Let $e = e_j$ and $\tilde{e} = e_{j+1}$ denote two subsequent edges in E_{cut}. If both belong to the same triangle in $\mathcal{T}_{\mathrm{char}}$, then the border triangulation between e and $\tilde{e}$ is completely described by $\mathcal{T}_{\mathrm{char}}$ as it just consists of this triangle. Otherwise, the line segment of C bounded by its intersection points with e and $\tilde{e}$ is either inside $\mathrm{int}(\widehat{\mathcal{H}})$ or outside. We only miss triangles of the border triangulation if the line segment is inside of $\mathrm{int}(\widehat{\mathcal{H}})$. In this case we say that e and $\tilde{e}$ defines a *gap along* C which we denote by $\Gamma = \Gamma(e, \tilde{e})$. The line segment of C which is bounded by e and $\tilde{e}$ is called the *window of* Γ and denoted by $W(\Gamma)$.

Between e and $\tilde{e}$ triangles of $\Delta_{\mathrm{B}}(\widehat{\mathfrak{R}}, \mathcal{T}) \setminus \mathcal{T}_{\mathrm{char}}$ cross the separator C. Those triangles induce portions $\mathcal{S} = \mathcal{S}(\Gamma, \mathcal{T})$ and $\mathcal{S}' = \mathcal{S}'(\Gamma, \mathcal{T})$ of $\mathcal{B}(\widehat{\mathfrak{R}}, \mathcal{T}) \cap \mathfrak{R}$ and $\mathcal{B}(\widehat{\mathfrak{R}}, \mathcal{T}) \cap \mathfrak{R}'$, respectively. Although $\mathcal{T}_{\mathrm{char}}$ provides no edges of $\mathcal{S}$ or $\mathcal{S}'$, we will derive several properties of both. For the sake of simplicity we only consider $\mathcal{S}$. Due to symmetry, similar definitions and statements are straightforward for $\mathcal{S}'$.

Note that there is a region R_i of (1) such that $\mathcal{S} \subseteq \mathcal{W}_i$ and that $\mathcal{S}$ is connected. The next lemma states one of the key properties of gaps that will subsequently be used to construct the algorithm $\mathcal{L}$ that computes our desired restricted class of local hulls. The lemma essentially says that while we cannot say anything specific about the location of gaps (e.g. we cannot guarantee any bound on the number of points on the missing portions $\mathcal{S}$ and $\mathcal{S}'$ nor can we bound their lengths), we nevertheless know that $\mathcal{S}$ and $\mathcal{S}'$ are contained within a strip of width at most $4c_l$, where c_l denotes the side length of a cell $\mathfrak{C} \in \mathcal{M}[\mathfrak{R}]$.

LEMMA 7. *Let $\Gamma = \Gamma(e, \tilde{e})$ denote a gap along C and let $\mathcal{S} = \mathcal{S}(\Gamma)$. There exists a line H with the following properties. Either H is orthogonal to C and contains the barycenter of $W(\Gamma)$ or there are 2 points of P such that H contains those points and crosses $W(\Gamma)$. Moreover, for all $u \in P(\mathcal{S}) \cup \{y, \tilde{y}\}$ we have $d(u, H) \leq 2c_l$. Here, y and $\tilde{y}$ denote the intersections points of C with e and $\tilde{e}$, respectively.*

We now choose a line B orthogonal to H that contains y or $\tilde{y}$ (or both, in case H and C are orthogonal) in such a way that (if H is not orthogonal to C) the halfspace B^- that contains $H \cap \mathfrak{R}'$ also contains the other point of y resp. $\tilde{y}$. If H is orthogonal to C we have $\mathfrak{R} \subset B \cup B^+$. Note that $W(\Gamma)$ is contained in B^- if H and C are not orthogonal. Note also that we cannot exclude that B^- contains points from $P(\mathcal{S})$. The next lemma, however, states that if such points exist then they are all contained in a small strip along the line B.

LEMMA 8. *For all $u \in P(\mathcal{S}) \cap B^-$ we have $d(u, B) \leq c_l$.*

5.3 The Skeleton

So far we have studied properties of a gap with respect to a given triangulation. Within our algorithmic framework we need to proceed in the opposite way: given Γ we need to construct a "good" $\mathcal{S}$. In order to do this in an efficient way we will provide a set of points $\mathbf{K}$, which we call a *skeleton*, that has at most logarithmic size and is dense enough to describe $\mathcal{S}$ reasonably precisely.

By Lemma 7 we know that there exists a line H such that the points on $\mathcal{S}$ are enclosed within a very narrow strip around H. However, we have no information about the behavior of $\mathcal{S}$ "along" H. Our approach is to sample $\mathbf{K}$ from $\mathcal{S}$ along H in steps of roughly size c_l. This is explained in detail in Section 6. In this section we first collect some properties that a skeleton should have. Note that our aim is to define a skeleton with respect to a given gap Γ and the point set $P \cap \mathfrak{R}$, but not with respect to a triangulation.

Let $\Gamma = \Gamma(e, \tilde{e})$ denote a gap and let x, $\tilde{x}$, y, and $\tilde{y}$ be defined as above. Moreover, let H denote an *arbitrarily* chosen line such that H is either the line orthogonal to C through the barycenter of $W(\Gamma)$ or some line through two points of P that intersects with $W(\Gamma)$ in exactly one point. We construct the line B as above, that is, B is orthogonal to H and contains y or $\tilde{y}$. The halfspaces B^- and B^+ are oriented as usual. Throughout this section B_p denotes the line parallel to B such that $p \in B_p$ if $p \in B^+$ and $B_p = B$ otherwise. The halfspace of B_p containing B^- is denoted by B_p^-.

A skeleton is a subset $\mathbf{K}$ of $P \cap \mathfrak{R}$ that satisfies certain properties. In fact, we will usually view $\mathbf{K}$ as a disjoint union of three sets

$$\mathbf{K} = \{\hat{p}\} \uplus \mathbf{K}^+ \uplus \mathbf{K}^-,$$

where $\hat{p}$ is a sort of extreme point and $\mathbf{K}^+$ and $\mathbf{K}^-$ define the two flanks of the skeleton, cf. Figure 3. A set $\mathbf{K} = \{\hat{p}\} \uplus \mathbf{K}^+ \uplus \mathbf{K}^-$ is called a skeleton, if it satisfies the following four properties:

(P1) If there are points of $\mathbf{K}$ in B^- then $\hat{p} \in \mathbf{K}$ is the point in B^- which has maximum distance to B. Otherwise $\hat{p}$ has minimum distance to B.

(P2) $x \in \mathbf{K}^+$, $\tilde{x} \in \mathbf{K}^-$, $B_x^+ \cap \mathbf{K}^+ = \emptyset$ and $B_{\tilde{x}}^+ \cap \mathbf{K}^- = \emptyset$.

(P3) $|\mathbf{K}^+ \cap B^-| \leq 1$ and $|\mathbf{K}^- \cap B^-| \leq 1$.

(P4) $d(B_p, B_q) > c_l$ for all $p, q \in \mathbf{K}^+$ $(p, q \in \mathbf{K}^-)$, $p \neq q$.

The intuition behind those properties is quite simple. (P1) basically states that $\hat{p}$ is that point in $\mathbf{K}$ that is "closest" to $W(\Gamma)$, where closeness is measured with respect to the hyperplane B. Property (P2) says that the two endpoints of the edges which define the gap Γ have to belong to $\mathbf{K}^+$ respectively $\mathbf{K}^-$ and that all other points have to be closer to $W(\Gamma)$, where closeness is again measured with respect to B. In the sections to come the points in $B^- \cap \mathfrak{R}$ will always play some special role. Property (P3) simply says that $\mathbf{K}^+$ and $\mathbf{K}^-$ both contain at most one such special point. Finally, property (P4) says that the points in $\mathbf{K}^+$ and $\mathbf{K}^-$ should be reasonably scattered along H.

Our next step is to define the notion of a *feasible skeleton*, which has the property that it uniquely defines as "standardized" hull segment. In order to characterize a feasible skeleton we first define a polyline $\mathcal{S}^\sim = \mathcal{S}^\sim(\Gamma, \mathbf{K})$ that connects x with $\tilde{x}$ via the points in $\mathbf{K}$. As we will define $\mathcal{S}^\sim$ as a union of piecewise convex segments we refer to $\mathcal{S}^\sim$ as a *convex continuation* (of the points in $\mathbf{K}$). Assume that the points in $\mathbf{K}^+ \cup \{\hat{p}\}$ and $\mathbf{K}^- \cup \{\hat{p}\}$ are ordered by increasing distance to B_x and $B_{\tilde{x}}$, respectively. Henceforth we write $p \to q$ if $p, q \in \mathbf{K}^+ \cup \{\hat{p}\}$ $(p, q \in \mathbf{K}^+ \cup \{\hat{p}\})$ are subsequent (!) in this order.

We construct $\mathcal{S}^\sim$ piecewise as follows. For $p, q \in \mathbf{K}^+$, $p \to q$, consider the (open) region D_{pq} that is bounded by e, the line segment pq, B_p and B_q. We let

$$\mathcal{S}^\sim[p, q] = E(\mathrm{conv}((P \cap D_{pq}) \cup \{p, q\})) \setminus \{p, q\}.$$

In other words $\mathcal{S}^\sim[p, q]$ is the convex polyline that separates all points of P that lie in the region D_{pq} from e. Observe that we add the two points p and q to the points in the open region D_{pq} in order to guarantee that the edge $\{p, q\}$ will be one side of the constructed convex hull.

If $p, q \in \mathbf{K}^-$, then $\mathcal{S}^\sim[p, q]$ is defined similarly; we just replace the edge e by the edge $\tilde{e}$ in the above definitions. Finally, consider the case $p \to q$ where $p \in \mathbf{K}^+ \cup \mathbf{K}^-$ and $q = \hat{p}$. Assume that $p \in \mathbf{K}^+$; the case $p \in \mathbf{K}^-$ is again symmetric. This time we let D_{pq} denote the open region bounded by e, pq, B_p, the shortest line segment connecting $\hat{p}$ to $W(\Gamma)$ and (if $p \in B^-$) by $\mathcal{S}^\sim[p', p]$, where $p' \in \mathbf{K}^+$ such that $p' \to p$. (Recall that if $p \in B^-$ then $p \notin B_p$. We thus need the segment $\mathcal{S}^\sim[p', p]$ in order to guarantee that D_{pq} is a finite region.) As before we let

$$\mathcal{S}^\sim[p, q] = E(\mathrm{conv}((P \cap D_{pq}) \cup \{p, q\})) \setminus \{p, q\}.$$

The convex continuation $\mathcal{S}^\sim$ is defined as the union of these pieces, that is

$$\mathcal{S}^\sim = \bigcup_{p,q \in \mathbf{K}: p \to q} \mathcal{S}^\sim[p, q].$$

This definition is illustrated in Figure 3. Intuitively, $\mathcal{S}^\sim$ consists of convex segments that connect two subsequent points in $\mathbf{K}$ such that $\hat{p}$ can be seen as the "peak" of $\mathcal{S}^\sim$. Let $p_0 \in \mathbf{K}^+$ and $q_0 \in \mathbf{K}^-$ such that $p_0 \to \hat{p}$ and $q_0 \to \hat{p}$. Observe that

$$\widehat{\mathcal{S}}^\sim := \mathcal{S}^\sim[p_0, q_0] = \mathcal{S}^\sim[p_0, \hat{p}] \cup \mathcal{S}^\sim[\hat{p}, q_0]$$

is a convex polyline, since $\hat{p}$ is an extreme point.

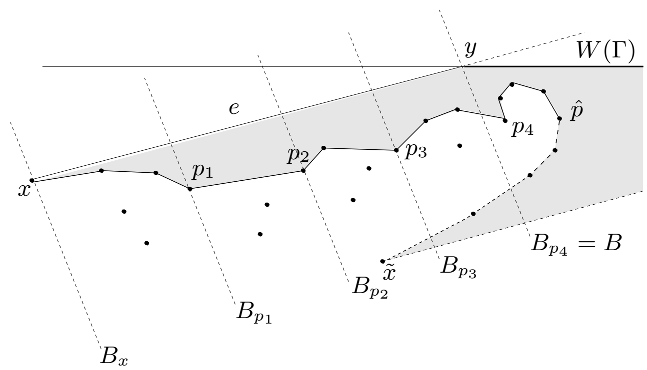

Figure 3: $\mathcal{S}^\sim$ restricted to the portion defined by $\mathbf{K}^+$ (solid lines). The polygonal region $D(\mathcal{S}^\sim)$ is shaded.

Throughout $D(\mathcal{S}^\sim)$ denotes the region bounded by $\mathcal{S}^\sim$, e, $\tilde{e}$ and $W(\Gamma)$. For a feasible skeleton we require that this region exists (cf. property (P5) below). Note that we view $D(\mathcal{S}^\sim)$ as an open region. In particular it therefore contains no points from $P(\mathcal{S}^\sim)$. A skeleton $\mathbf{K}$ is said to be *feasible* if it satisfies the following properties:

(P5) $\mathcal{S}^\sim$ is neither self-intersecting nor intersects with e or $\tilde{e}$.

(P6) The polygonal region $D(\mathcal{S}^\sim)$ does not contain points of P. (Note that by definition of the convex continuation this is equivalent to saying that the line segments $B_p \cap D(\mathcal{S}^\sim)$ contain no point.)

(P7) Let $p, q \in \mathbf{K} \setminus \{\hat{p}\}$ such that $p \to q$. Let X_q and Y_q denote the lines through q that are parallel to C and orthogonal, respectively. Let X_q^- denote the halfspace of X_q containing $W(\Gamma)$ and let $v \in X_q^- \cap P(\mathcal{S}^\sim[p, q])$ such that $d(v, X_q)$ is maximum. If $v \notin \{p, q\}$ then $d(B_v, B_q) > c_l$. Similarly, if $W(\Gamma) \cap Y_q = \emptyset$ the following is true. Let Y_q^- denote the halfspace containing $W(\Gamma)$ and let $v \in Y_q^- \cap P(\mathcal{S}^\sim[p, q])$ such that $d(v, Y_q)$ is maximum. If $v \notin \{p, q\}$ then $d(B_v, B_q) > c_l$.

(P8) For all $u \in P(\mathcal{S}^{\sim})$ there is an edge $\{u, w\}$ connecting u and some $w \in W(\Gamma)$ such that $uw \subset D(\mathcal{S}^{\sim})$. In specific $\{u, w\}$ crosses neither $\mathcal{S}^{\sim}$ nor e and $\tilde{e}$.

Too get some intuition for the last condition observe that whenever we consider the segment $\mathcal{S} = \mathcal{S}(\Gamma)$ defined by an underlying triangulation, then every edge of $\mathcal{S}$ is contained in a triangle that crosses $W(\Gamma)$. In particular, $\mathcal{S}$ satisfies property (P8) and we require that $\mathcal{S}^{\sim}$ does so, too. We close this subsection with the observation that any skeleton contains at most logarithmically many points.

LEMMA 9. *For fixed $s \in \mathbb{N}$ we have $|\mathbf{K}| = \mathcal{O}(\log n)$ for every skeleton $\mathbf{K}$.*

5.4 Smooth Hulls and Triangulations

For fixed $\mathfrak{Q}_\circ$ we have the following definitions. A triangulation $\mathcal{T}$ *patches* Γ *smoothly*, if and only if there exist hyperplanes H and H' through $W(\Gamma)$ and feasible skeletons $\mathbf{K} = \mathbf{K}(\Gamma, H)$ and $\mathbf{K}' = \mathbf{K}'(\Gamma, H')$ such that

$$\mathcal{S}(\Gamma, \mathcal{T}) = \mathcal{S}^{\sim}(\mathbf{K}) \quad \text{and} \quad \mathcal{S}'(\Gamma, \mathcal{T}) = \mathcal{S}^{\sim\prime}(\mathbf{K}').$$

Similarly, a local hull of $\mathfrak{R}$ is called *smooth*, if all gaps are patched smoothly and if the hull of its parent $\widehat{\mathfrak{R}}$ is either smooth or identical to $\text{conv}(P)$. Finally, $\mathcal{T}$ is said to be *smooth with respect to* $\mathfrak{Q}_\circ$ if and only if every rectangle has a smooth local hull.

5.5 The Algorithm $\mathcal{L}$

Our next goal is to construct a quasi-polynomial time algorithm $\mathcal{L}$ that generates for all $\mathfrak{R}$ a set of hulls $\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]$ such that $\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]$ contains all smooth hulls of $\mathfrak{R}$.

We initialize $\mathbf{H}_{\mathcal{L}}[\mathfrak{Q}_\circ] = \{\text{conv}(P)\}$ and proceed top-down in the sense of Lemma 2. So assume we already have computed $\mathbf{H}_{\mathcal{L}}[\widehat{\mathfrak{R}}]$ and that C subdivides $\widehat{\mathfrak{R}}$ into $\mathfrak{R}$ and $\mathfrak{R}'$. We compute $\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]$ as follows.

1. For each $\widehat{\mathcal{H}} \in \mathbf{H}_{\mathcal{L}}[\widehat{\mathfrak{R}}]$ enumerate all feasible sets $\mathcal{T}_{\text{char}}$ of characteristic triangles.

2. For all such choices of $\widehat{\mathcal{H}}$ and $\mathcal{T}_{\text{char}}$, compute all gaps, say $\Gamma_1, \ldots, \Gamma_r$. For each gap Γ_i, choose a line H_i and a feasible skeleton $\mathbf{K}_i$. Enumerate all such choices.

3. For each such choice $[(H_1, \mathbf{K}_1), \ldots, (H_r, \mathbf{K}_r)]$ we compute $\mathcal{H}$ as follows. Check whether the convex continuations $\mathcal{S}^{\sim}(\mathbf{K}_i)$ are pairwise disjoint. If yes, combine them with $\widehat{\mathcal{H}} \cap \mathfrak{R}$ and the sides of triangles in $\mathcal{T}_{\text{char}}$ that are contained in $\mathfrak{R}$. This yields $\mathcal{H}$ which we add to $\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]$.

One easily checks that $\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]$ contains all smooth hulls of $\mathfrak{R}$. So it remains to bound $|\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]|$ and the time complexity of $\mathcal{L}$. This however, is far beyond the scope of this extended abstract. The hard part is to show that only $\mathcal{O}\left(\log^6 n\right)$ gaps occur in step (2) of the algorithm. Once this is achieved, we obtain using basic counting arguments

LEMMA 10. *$\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]$ contains all smooth hulls of $\mathfrak{R}$ and for fixed $s \in \mathbb{N}$, $|\mathbf{H}_{\mathcal{L}}[\mathfrak{R}]| \leq n^{\mathcal{O}\left(\log^8 n\right)}$. Moreover, the family $(\mathbf{H}_{\mathcal{L}}[\mathfrak{R}])_{\mathfrak{R}}$ can be computed in time $n^{\mathcal{O}\left(\log^8 n\right)}$.*

6. THE SMOOTHING LEMMA

Our QPTAS computes an almost optimal smooth triangulation. In order to show that this also approximates an optimum triangulation, we will show in the next section how to change a general triangulation into a smooth triangulation without much increase in length. In this section we introduce a technical statement needed for this transformation, the so-called Smoothing Lemma. Consider a separator C and a gap $\Gamma = \Gamma(e, \tilde{e})$ on C. Let

$$\mathcal{W}(\mathcal{T}, e, \tilde{e}) = \mathcal{S} \cup \mathcal{S}' \cup \{e, \tilde{e}\}\,.$$

We call $\mathcal{W}(\mathcal{T}, e, \tilde{e})$ the *enclosing walk of* $\Gamma(e, \tilde{e})$ (cf. Figure 4).

LEMMA 11 (SMOOTHING LEMMA). *Let $\mathcal{T}$ denote a triangulation and let C denote a separator at level l subdividing $\widehat{\mathfrak{R}}$ into $\mathfrak{R}$ and $\mathfrak{R}'$. Furthermore, let $\Gamma(e, \tilde{e})$ be a gap on C. Then there exist feasible skeletons $\mathbf{K}$ and $\mathbf{K}'$ as well as a triangulation $\mathcal{T}^{\sim}$ such that*

1. *The edges in the symmetric difference $\mathcal{T}^{\sim} \Delta \mathcal{T}$ are contained in the interior of $\mathcal{W} = \mathcal{W}(\mathcal{T}, e, \tilde{e})$. In other words, $\mathcal{T}$ and $\mathcal{T}^{\sim}$ are identical outside $\mathcal{W}$.*

2. *$\mathcal{T}^{\sim}$ patches $\Gamma(e, \tilde{e})$ smoothly with respect to $\mathbf{K}$ and $\mathbf{K}'$.*

3. *$\ell(\mathcal{T}^{\sim}) - \ell(\mathcal{T}) \leq 18c_l \cdot h$, where h denotes the number of edges of $\mathcal{T}$ in the interior of $\mathcal{W}$.*

In the remainder of this section we outline the proof of this statement. We start by fixing some notation. As usual we use $x = e \cap \mathfrak{R}$ and $\tilde{x} = \tilde{e} \cap \mathfrak{R}$ to denote the endpoints of e and $\tilde{e}$ that lie in $\mathfrak{R}$. Similarly, x' and $\tilde{x}'$ denote the endpoints in $\mathfrak{R}'$. Recall that Lemma 7 guarantees the existence of lines H and H' such that $\mathcal{S}$ and $\mathcal{S}'$ are enclosed within a strip of width $4c_l$ around H and H', respectively. The lines B and B' denote appropriate hyperplanes orthogonal to H. Moreover, B_p denotes the line parallel to B through p if $p \in B^+$, and $B_p = B$ otherwise. The halfspace of B_p containing B^- is denoted by B_p^- (this is unique as $B_p \subset B^+ \cup B$). Throughout $d_\perp(u, v)$ and $d_\parallel(u, v)$ denote the distance between u and v measured orthogonal to H and parallel, respectively. We have

$$d(u, v) \leq d_\perp(u, v) + d_\parallel(u, v)$$

due to the triangle inequality.

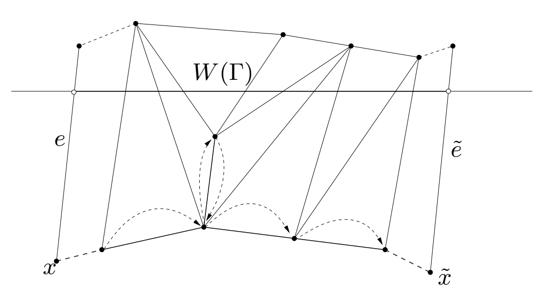

Figure 4: The walk $\mathcal{S}$ traverses the edges in the same order as we traverse the corresponding triangles along C.

Recall that every edge in $\mathcal{S}$ is a side of at most two triangles of $\mathcal{T}$ that lie within $\mathcal{W}$ and that every edge that crosses

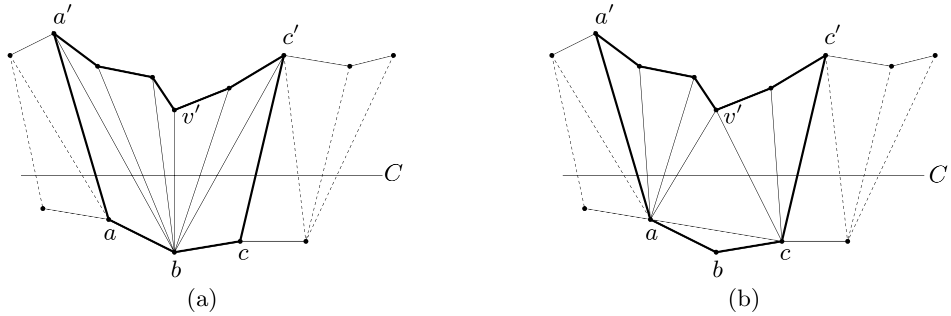

Figure 5: We identify $\mathcal{W}_{abc}$ (bold) and modify the triangulation in its interior.

$W(\Gamma)$ belongs to such a triangle. We thus get a natural walk along the edges of $\mathcal{S}$, if we traverse the edges of $\mathcal{S}$ in the order in which the corresponding triangles are encountered if we walk along $W(\Gamma)$ from e to $\tilde{e}$. This is illustrated in Figure 4. Observe, this walk can be written as a set of directed edges (arcs). Henceforth, we slightly abuse notation by assuming that $\mathcal{S}$ is indeed a set of such arcs rather then a PSLG.

Note that in general $\mathcal{S}$ will be a walk (and not a path). That is, some points can be visited several times. With $\mathcal{S}[x,p]$ we always denote the walk on $\mathcal{S}$ from x until the *first* occurrence of p on $\mathcal{S}$. Similarly, $\mathcal{S}[\tilde{x},p]$ denotes the walk which we obtain by traversing $\mathcal{S}$ in opposite direction, starting in $\tilde{x}$, until the *first* occurrence of p.

LEMMA 12. *Let L denote a line parallel to B such that $L \subset B^+$ and $z := e \cap L$ exists. Let $X = \mathcal{S} \cap L$ denote the intersection points of $\mathcal{S}$ with L and let z_0 denote the first intersection of $\mathcal{S}$ with L if we traverse $\mathcal{S}$ starting in x. Then the line segment zz_0 contains no other point from X.*

A similar statement holds for $z := \tilde{e} \cap L$ and the last intersection point of L and $\mathcal{S}$, as everything is symmetric if we walk backwards starting at $\tilde{x}$.

6.1 Sampling the Skeleton

One can show that it is possible to sample a feasible skeleton $\mathbf{K}$ from $\mathcal{S}$ which satisfies in addition the following properties:

(A1) $\mathbf{K} \subseteq P(\mathcal{S})$.

(A2) If $p \in \mathbf{K}^+$ then $P(\mathcal{S}[x,p]) \setminus \{p\} \subset B_p^+$ and similarly if $p \in \mathbf{K}^-$ then $P(\mathcal{S}[\tilde{x},p]) \setminus \{p\} \subset B_p^+$.

(A3) For all $p,q \in \mathbf{K}^+$ s.t. $p \to q$, all points in $P(\mathcal{S}[x,q]) \setminus \{q\}$ are either contained in B_p^+ or satisfy $d(u,B_p) \le 4c_l$. For $p,q \in \mathbf{K}^-$ a similar statement holds with respect to $\mathcal{S}[\tilde{x},q]$.

As $\mathbf{K}$ is feasible, $\mathcal{S}^\sim = \mathcal{S}^\sim(\mathbf{K})$ exists. Due to symmetry, we can clearly sample a feasible skeleton $\mathbf{K}'$ with similar properties from $\mathcal{S}'$.

6.2 Smoothable Detours

Let X denote the set of edges of the triangulation that cross C within the enclosing walk $\mathcal{W}(\mathcal{T},e,\tilde{e})$ of Γ, that is,

$$X = (E(\mathcal{T}) \setminus \{e,\tilde{e}\}) \cap \big(P(\mathcal{S}) \times P(\mathcal{S}')\big)\,. \tag{3}$$

Note that X includes neither e nor $\tilde{e}$. For the sake of exposition, we consider only $\mathcal{S}$ in the sequel. Throughout it can be easily checked that similar definitions and statements can be derived for $\mathcal{S}'$. Observe that every arc of $\mathcal{S}$ is contained in exactly one of the following subwalks.

- $\widehat{\mathcal{S}}$. The subwalk of $\mathcal{S}$ from the *first* occurrence of $p \in \mathbf{K}^+$ and the *last* one of $q \in \mathbf{K}^-$, where $p \to \hat{p}$ and $q \to \hat{p}$.
- $\mathcal{S}[p,q]$. Defined for $p,q \in \mathbf{K}^+$, $p \to q$ as the subwalk of $\mathcal{S}$ from the *first* occurrence of p to the *first* occurrence of q. For $p,q \in \mathbf{K}^-$ by considering $\mathcal{S}$ in reverse direction.

In other words, we decompose $\mathcal{S}$ into $\widehat{\mathcal{S}}$ and a family of sets $\mathcal{S}[p,q]$. Lemma 12 and (A2) imply that $\mathcal{S}[p,q] = \mathcal{S}^\sim[p,q]$ if $\mathcal{S}[p,q]$ is convex. Intuitively, this says that there is no substantial structural difference between $\mathcal{S}^\sim[p,q]$ and $\mathcal{S}[p,q]$ and it suffices to make the latter convex without increasing the length of the triangulation too much. This is achieved by iteratively "shortcutting" $\mathcal{S}[p,q]$. Similarly, we proceed for $\widehat{\mathcal{S}}$ which we aim to transform into $\widehat{\mathcal{S}}^\sim$. This motivates the following definition.

Let a, b, and c denote three subsequent points on $\mathcal{S}$, i.e., $\{(a,b),(b,c)\} \subseteq \mathcal{S}$. We say that (a,b,c) is a *detour*, if there is at least one edge in X incident to b that intersects with the line segment ac that connects a and c. Let $X_b \subseteq X$ denote the set of all edges in X that are incident to b. We call (a,b,c) *simple* if $|X_b| = 1$. Moreover, (a,b,c) is called *smoothable* if either $\{(a,b),(b,c)\} \subset \widehat{\mathcal{S}}$ or $\{(a,b),(b,c)\} \subset \mathcal{S}[p,q]$ for an appropriate choice of p and q.

6.3 The Smoothing Action

A *smoothing action* is a single step of the process that iteratively "shortcuts" $\mathcal{S}$ such that we obtain $\mathcal{S}^\sim$. We choose a smoothable detour (a,b,c) on $\mathcal{S}$ or $\mathcal{S}'$ according to a carefully chosen order of preference (this requires more technical background then this extended abstract can provide). We assume that

$$\{(a,b),(b,c)\} \subseteq \mathcal{S}$$

as the case $\{(a,b),(b,c)\} \subseteq \mathcal{S}'$ is just symmetric. Once the detour is chosen, we locally modify $\mathcal{T}$ such that we obtain a triangulation $\mathcal{T}_<$ with

$$\mathcal{S}_< = (\mathcal{S} \cup \{(a,c)\}) \setminus \{(a,b),(b,c)\}\,.$$

This is achieved as follows. Let $a',c' \in P(\mathcal{S}')$ such that $\triangle_{aba'}$ and $\triangle_{cbc'}$ are the triangles corresponding to (a,b) and (b,c). The edges $\{a,a'\}$ and $\{c,c'\}$ so to say enclose

X_b. Note that $a' = c'$ is possible. Let $\mathcal{W}_{abc}$ denote the walk that encloses all triangles in $\mathcal{T}$ that are traversed along C between $\triangle_{aba'}$ and $\triangle_{cbc'}$. In Figure 5 this walk $\mathcal{W}_{abc}$ is indicated by a bold polygon. All the modifications take place within the area enclosed by $\mathcal{W}_{abc}$.

We triangulate the interior of $\mathcal{W}_{abc}$ with a triangulation that contains $\{a, c\}$. Clearly, $\mathcal{S}'_{abc} = \mathcal{S}' \cap \mathcal{W}_{abc}$ is a subwalk of $\mathcal{S}'$. Observe furthermore that $\mathcal{S}'_{abc}$ visits each of its points exactly once, since every point is connected to b. Thus the subwalk $\mathcal{S}'_{abc}[u'_1, u'_2]$ from u'_1 to u'_2 is well-defined.

We choose a point v' as follows. If there are points in $B \cup B^-$ we choose v' such that the distance to C is minimized. Otherwise, we choose v' among all points (in B^+) having minimum distance to B. The new triangulation of $\mathcal{W}_{abc}$ contains the triangle $\triangle_{av'c}$ and has the property that all triangles but $\triangle_{abc}$ cross C. This is illustrated in Figure 5 b). Our order of preference guarantees that this triangulation always exists.

It remains to bound the increase of length due to this sequence of smoothing actions. This is quite elaborate, so we just outline the main ideas. Observe that every smoothing action replaces one edge crossing C by one that lies inside $\mathfrak{R}$ or $\mathfrak{R}'$ (namely the edge $\{a, c\}$). As $d_\perp(a, c) \leq 4c_l$ it remains to show that $d_\parallel(a, c)$ is bounded by the length of the replaced edges plus $14c_l$. We roughly argue as follows. The smoothing actions' order of preference guarantees that we first consider detours with $a, c \in B_b \cup B_b^-$ (the bound on $d_\parallel(a, c)$ is quite obvious for this type of detours). If no such detour is left then (A2) and Lemma 12 assure that $a, c \in B_p^+$, where $p \in \mathbf{K}$ is chosen such that b is on the subwalk between p and $\hat{p}$ and this subwalk is shortest possible (very roughly speaking). In this case, the bound on $d_\parallel(a, c)$ follows by (A3). On the other hand, a smoothing action "redirects" some of the remaining edges in X. The order of preference guarantees that the endpoints of those edges "tend" to B over the sequence of smoothing actions. This allows us to apply (A3) similarly as above.

7. ON THE CHOICE OF $\mathfrak{Q}_\circ$

It remains to be shown that $\mathcal{T}^*_{\mathcal{L}}$ is a good approximation to $\mathcal{T}^*$, that is, $\ell(\mathcal{T}^*_{\mathcal{L}}) \leq (1+\varepsilon)\ell(\mathcal{T}^*)$. This is in general not the case for arbitrary fixed $\mathfrak{Q}_\circ$. However, we can show that there exists

$$\mathbf{Q} = \{\mathfrak{Q}_1, \mathfrak{Q}_2, \ldots, \mathfrak{Q}_k\}$$

such that $\mathcal{T}^*_{\mathcal{L}} \leq (1+\varepsilon)\ell(\mathcal{T}^*)$ for at least one $\mathfrak{Q}_\circ \in \mathbf{Q}$, where $k = \text{poly}(n)$. Due to space restrictions, we just sketch the main ideas of our proofs.

In order to comply with the assumptions of the previous sections we require that $\mathfrak{Q}_\circ$ (and thus every $\mathfrak{Q}_i \in \mathbf{Q}$) has to satisfy the following properties.

– $P \subset \mathfrak{Q}_\circ$.

– $\max_{p,q \in P} d(p, q) \geq s \cdot n$.

– $|\mathfrak{Q}_\circ| = 2^\alpha$, for some $\alpha \in \mathbb{N}$ s.t. $2^\alpha \in [3sn, 6sn]$.

– For all rectangles $\mathfrak{R}$ in the subdivision, neither the boundary of $\mathfrak{R}$ nor in the boundary of any cell $\mathfrak{C} \in \mathcal{M}[\mathfrak{R}]$ contains a point in P.

Indeed the properties listed here are somewhat stronger than actually required. By scaling and slightly shifting the point set one can check that those assumptions do not mean loss of generality. It remains to show that there exists a "good" $\mathfrak{Q}_\circ \in \mathfrak{R}$. For a fixed square $\mathfrak{Q}_\circ$ we proceed in top-down fashion through the dissection tree and use Lemma 11 to modify the triangulation appropriately, that is, we assure iteratively that all gaps are patched smoothly. Using an averaging argument similar to that of Arora [2] we can indeed show that there exists a "good" $\mathfrak{Q}_\circ \in \mathfrak{R}$ for which the increase of length of the triangulation in the aforesaid process is at most $\varepsilon\ell(\mathcal{T}^*)$ for this choice of $\mathfrak{Q}_\circ$. As $\mathbf{Q}$ has polynomial size, this (or a better) $\mathfrak{Q}_\circ$ can be identified by running our approximation algorithm for each square in $\mathbf{Q}$.

8. CONCLUDING REMARKS

In this paper we introduced a quasi-polynomial time approximation scheme for Minimum Weight Triangulation. Due to the recent result of Mulzer and Rote [16], the most natural aim of future research is to prove the existence of a PTAS. This seems to be promising, as in fact our result can be seen as an indication that such an algorithm exists.

9. REFERENCES

[1] S. Arora. Polynomial time approximation schemes for Euclidean TSP and other geometric problems. In *37th Annual Symposium on Foundations of Computer Science*, pages 2–11, 1996.

[2] S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.

[3] S. Arora. Approximation schemes for NP-hard geometric optimization problems: A survey. *Mathematical Programming, Series B*, 97:43–69, 2003.

[4] S. Arora and K. Chang. Approximation schemes for degree-restricted MST and red-blue separation problems. *Algorithmica*, 40(3):189–210, 2004.

[5] S. Arora, P. Raghavan, and S. Rao. Approximation schemes for Euclidean k-medians and related problems. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, pages 106–113, 1998.

[6] M. Bern and D. Eppstein. Approximation algorithms for geometric problems. In D. Hochbaum, editor, *Approximation Algorithms for NP-hard Problems*. PWS Publishing, 1997.

[7] D. Eppstein. Approximating the minimum weight Steiner triangulation. *Discrete & Computational Geometry*, 11:163–191, 1994.

[8] M. R. Garey and D. S. Johnson. *Computers and Intractability - A Guide to the Theory of NP-Completness*. 1979.

[9] P. Gilbert. New results on planar triangulations. Master's thesis, University of Illinois, 1979.

[10] J. Gudmundsson and C. Levcopoulos. Minimum weight pseudo-triangulations. In *Proceedings of the 20th European Workshop on Computational Geometry*, 2004.

[11] G. Klincsek. Minimal triangulations of polygonal domains. *Annals of Discrete Mathemathics*, 9:121–123, 1980.

[12] C. Levcopoulos and D. Krznaric. Quasi-greedy triangulations approximating the minimum weight triangulation. In *Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 392–401, 1996.

[13] C. Levcopoulos and D. Krznaric. A near-optimal heuristic for minimum weight triangulation of convex polytops. In *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 519–527, 1997.

[14] C. Levcopoulos and D. Krznaric. Quasi-greedy triangulations approximating the minimum weight triangulation. *Journal of Algorithms*, 27:303–338, 1998.

[15] J. S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple new method for the geometric k-MST problem. In *Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 402–408, 1996.

[16] W. Mulzer and G. Rote. Minimum weight triangulation is NP-hard. In *Proceedings of the 22nd Annual ACM Symposium on Computational Geometry*, 2006 (to appear).

[17] D. A. Plaisted and J. Hong. A heuristic triangulation algorithm. *Journal of Algorithms*, 8:465–437, 1987.