Quasi-Greedy Triangulations Approximating the Minimum Weight Triangulation^{*}

Christos Levcopoulos Drago Krznaric Department of Computer Science Lund University, Box 118, S-221 00 Lund, Sweden. christos@dna.lth.se drago@dna.lth.se

Abstract

This paper settles the following two longstanding open problems:

- 1. What is the worst-case approximation ratio between the greedy and the minimum weight triangulation?
- 2. Is there a polynomial time algorithm that always produces a triangulation whose length is within a constant factor from the minimum?

The answer to the first question is that the known $\Omega(\sqrt{n})$ lower bound is tight. The second question is answered in the affirmative by using a slight modification of an $O(n \log n)$ algorithm for the greedy triangulation. We also derive some other interesting results. For example, we show that a constant-factor approximation of the minimum weight convex partition can be obtained within the same time bounds.

1 Introduction

Let S be any set of n points in the plane. A *triangulation* of S is a maximal straight-line graph whose vertices are the points in S. Any triangulation of S partitions the convex hull of S into empty triangles.

A triangulation that has received special attention is the *minimum weight triangulation*, in which the optimization criteria is to minimize the total edge length. This triangulation has some good properties [2] and is e.g. useful for numerical approximation of bivariate data [20].

Gilbert [4] showed that it can be computed for simple polygons in time $O(n^3)$ by dynamic programming. For general point sets, however, it is not known whether a minimum weight triangulation can be found in polynomial time, nor is it known whether this is an NP-hard problem. In consequence of this, heuristics for approximating it have been considered. There are two well-known heuristics: the greedy triangulation and the Delaunay triangulation, both being computable in $O(n \log n)$ time [9, 18].

Lloyd [15] showed that, in general, none of these two heuristics produce a minimum weight triangulation. For an arbitrary large n, Manacher and Zobrist [16] showed that one

^{*}This paper was partially supported by TFR

can place n points so that the Delaunay is $\Omega(n/\log n)$ from the optimum. They also showed that n points can be placed so that the greedy is $\Omega(n^{1/3})$ from the optimum. Kirkpatrick [6] strengthened the former bound by exhibiting, for each n, a set of n points for which the Delaunay triangulation has length $\Omega(n)$ times the optimum. The lower bound for the greedy was improved by Levcopoulos [8] who constructed, for each n, a set of n points for which the greedy triangulation has length $\Omega(\sqrt{n})$ times the optimum.

In this paper we prove that the latter bound is tight. That is, for any set S of n points in the plane, we show that the greedy triangulation of S has length no more than $O(\sqrt{n})$ times the optimum. This also gives an $O(\sqrt{n})$ upper bound for the heuristic recently proposed by Heath and Pemmaraju [5], because their triangulation is obtained by optimally triangulating a subgraph of the greedy triangulation. A similar idea was also used earlier by Lingas [14], but with the difference that he triangulated a subgraph of the Delaunay triangulation. It can be shown, however, that these two heuristics can produce triangulations that are $\Omega(\sqrt{n})$, respectively $\Omega(n)$, times longer than the optimum [10].

Although the greedy and the Delaunay can in general yield "bad" approximations, there are some special cases for which they have been proved to perform well. For example, if the points are uniformly distributed, both the Delaunay and the greedy triangulation are expected to be within a constant factor from the optimum [1, 12]. Also, if the points lie on their convex hull, then the greedy triangulation approximates the optimum [11].

In this paper we show that there is in fact only one certain configuration of points for which the greedy does not approximate the optimum. This is, loosely speaking, when the points form two concave chains facing each other, and the greedy heuristic produces edges that connect all points on a long piece of one of the chains with points lying on a short piece of the other chain (for a precise description, see the proof of Lemma 3.2). It suffices therefore to impose a small modification on the greedy heuristic in order to avoid this worst-case scenario: whenever the greedy heuristic is about to produce a diagonal d, we check whether there is a diagonal d', crossing d and only slightly longer than d (this is why we call it quasi-greedy), and fulfilling some other local conditions that can be checked in constant time. If so, then we produce d' instead of d. In this way we obtain a fast algorithm that for any set of points produces a triangulation whose length is within a constant factor from the optimum.

The heuristic that achieved the best proved worst-case approximation ratio was proposed by Plaisted and Hong [17]. They started by finding a convex partition that is at most 12 times longer than the minimum weight convex partition. Then they triangulated each of the convex polygons induced by their convex partition using the so-called ring heuristic. In this way they produced a triangulation whose length is within a factor $O(\log n)$ from the optimum. Their method was implemented by Smith [19] to run in time $O(n^2 \log n)$. Interestingly, it follows from Lemma 4.2 in this paper that any convex partition which is *c*-sensitive (Definition 2.3) and approximates the minimum can be used to obtain a constant-factor approximation of the minimum weight triangulation by e.g. greedy triangulating each induced convex polygon (assuming *c* is constant and no three vertices are collinear).

The paper is organized as follows. In Section 2 we give some basic definitions and preliminaries. In particular, we define the so-called *greedy convex partition* which is a subgraph of the greedy triangulation that partitions the convex hull into empty convex polygons. Then we prove in Section 3 that the length of the greedy convex partition is no more than $O(\sqrt{n})$ times the length of the minimum weight convex partition. In Section 4 we show that the length of the greedy triangulation is asymptotically no more than the length of the greedy convex partition plus the length of the minimum weight triangulation. In Section 5 we use the ideas developed in previous sections to show that a constant-factor approximation of the minimum weight triangulation can be computed in polynomial time $(O(n \log n))$ time by the recent results in [9]). Finally we show in Section 6 that a constant-factor approximation of the minimum weight convex partition can be obtained within the same time bounds. In that section we also discuss briefly how we can generalize our results in order to compute a constant-factor approximation for the minimum weight triangulation of a set of line segments, and how to deal with degenerate cases.

2 Definitions and preliminaries

A planar straight-line graph (PSLG) is a graph G which consists of a finite vertex set S and an edge set L. The vertices in S correspond to distinct points in the plane. The edges in Lcorrespond to straight-line segments with endpoints in S, such that no edge in L properly includes any vertex in S nor properly intersects any edge in L. A diagonal of G is a straightline segment with endpoints in S such that it together with G forms another PSLG. The length of a G, which we denote by |G|, is the total edge length in G. To simplify the presentation, we assume that S is in general position in the sense that no three vertices are collinear (we describe briefly in Section 6.3 how our results can be extended to degenerate cases).

A triangulation of S is a PSLG with vertex set S and with a maximum number of edges. A minimum weight triangulation of S (in brief, MT(S)) is a triangulation T of S such that |T| is minimized. The greedy triangulation of S (GT(S) for short) is obtained by repeatedly producing a shortest possible edge that does not properly intersect any of the previously generated edges.

By a convex polygon we mean a PSLG P such that P is a simple cycle and all vertices of P lie on their convex hull. A convex partition of S is a connected PSLG with vertex set S such that its edges partition the interior of the convex hull of S into bounded convex regions. A minimum weight convex partition of S, which we abbreviate as MC(S), is a convex partition of S such that its length is minimized. For a vertex v in S, we denote by max(v) the length of a longest edge incident to v in MC(S). The following observation is easy to show.

Observation 2.1 For any set S of vertices, $\sum_{v \in S} \max(v) \le 2|\operatorname{MC}(S)|$.

Next we define a subgraph of GT(S) such that it is a convex partition of S. It is obtained by selecting for each vertex v at most three edges of GT(S), which we call *spokes*. If v lies on the convex hull of S, then the spokes of v are the two convex hull edges incident to v. Otherwise, let e be a shortest greedy edge incident to v such that there exists greedy edges e' and e'' incident to v with the following three properties.

- (i) e' and e'' are not longer than e,
- (ii) e, e' and e'' partition the infinitesimal vicinity of v into (three) convex regions, and
- (iii) within the two of these regions bounded by e no greedy edge incident to v is shorter than e.

Then the spokes of v are the edges e, e' and e''. The union of all spokes which we select in this way, by considering all vertices in S, form our greedy convex partition, which we shall abbreviate as GC(S). For a vertex v in S, let v_G stand for the length of a longest spoke that was selected for v. The following observation is straightforward.

Observation 2.2 For any set S of vertices, $|GC(S)| \leq 3 \sum_{v \in S} v_G$.

In [13] the following notion was introduced.

Definition 2.3 Let G be a planar straight-line graph with vertex set S and let r be a real number greater than zero. An edge e of G is said to be r-sensitive if for any diagonal d of S that properly intersects e, the distance from any endpoint of d to the closest endpoint of e is not greater than r times the length of d. We say that G is r-sensitive if and only if all its edges are r-sensitive.

We shall use the following results from [13] and [11], respectively.

Fact 2.4 (Theorem 3.1 in [13]) For any vertex set S, the greedy triangulation of S is 4-sensitive.

Fact 2.5 (Theorem 2.1 in [11]) For any convex polygon P, |GT(P)| = O(|MT(P)|).

3 A tight bound for the greedy convex partition

This section is entirely devoted to the proof of the following theorem.

Theorem 3.1 For any set S of n vertices (in general position),

$$\frac{|\operatorname{GC}(S)|}{|\operatorname{MC}(S)|} = O(\sqrt{n}).$$

To begin the proof, let r be a number such that the following holds.

$$\frac{|\operatorname{GC}(S)|}{|\operatorname{MC}(S)|} = 7r.$$

We assume w.l.o.g. that r > 1000 (this assumption will not be used until the proof of Lemma 3.2 below). Let S' be the set of all vertices v in S such that $v_G > r \cdot \max(v)$. By Observations 2.1 and 2.2 it follows that S' is nonempty. Indeed, if v_G would be $\leq r \cdot \max(v)$ for each $v \in S$, then $|\operatorname{GC}(S)| \leq 3 \sum_{v \in S} v_G \leq 3r \sum_{v \in S} \max(v) \leq 6r |\operatorname{MC}(S)|$.

Let a be a vertex in S' such that for any other vertex v in S' it holds that $v_G \leq a_G$. By the following calculations (using Observations 2.1 and 2.2)

$$\begin{aligned} |\operatorname{GC}(S)| &\leq 3\sum_{v\in S'} v_G + 3\sum_{v\in S-S'} v_G \\ &\leq 3na_G + 3r\sum_{v\in S-S'} \max(v) \end{aligned}$$

$$\leq 3na_G + 6r |\mathrm{MC}(S)|$$

$$\leq 3na_G + (6/7)|\mathrm{GC}(S)|$$

it follows that $|\operatorname{GC}(S)| \leq 21 n a_G$. Combining this with Lemma 3.2 below we get that

$$\frac{|\operatorname{GC}(S)|}{|\operatorname{MC}(S)|} = O\left(\frac{na_G}{a_G^2/\max(a)}\right) = O\left(\frac{n}{r}\right).$$

Hence, 7r = O(n/r), which implies that $r = O(\sqrt{n})$.

Lemma 3.2

$$|\mathrm{MC}(S)| = \Omega\left(\frac{a_G^2}{\max(a)}\right).$$

Proof We first observe that a does not lie on the convex hull of S, because a_G would otherwise be the length of a convex hull edge which also belongs to MC(S), and so the property $a_G > r \cdot \max(a) > 10^3 \max(a)$ would not hold. Let e be a spoke of a whose length equals a_G , and let e' and e'' be the two other spokes that were together with e selected for a when defining GC(S). We assume w.l.o.g. that e' is horizontal and that a is its right endpoint, and that the endpoint of e'' which is different from a lies lower than a (as depicted in Figure 1). Among those vertices that lie higher than a, let u be the one which is closest to a, and let l be the distance between a and u. From the assumption that $a_G > 10^3 \max(a)$ it follows easily that $a_G > 10^3 l$ (recall that no three vertices are collinear). Further, by the definition of GC(S), the greedy algorithm will not produce a greedy edge which connects a with a vertex lying higher than a until it starts to produce greedy edges of length $\geq a_G > 1000l$ (since a_G is the length of a shortest such greedy edge).

We shall begin with considering the situation just after having produced all greedy edges of length $\leq 2l$ (we now have a PSLG where all diagonals have length > 2l and all (greedy) edges have length $\leq 2l$). Let e_1 be the first greedy edge (among those produced so far) which is crossed by a straight-line walk from a to u (we must cross some edge since we would otherwise have produced a greedy edge that connects a and u). Let p be the point in which (a, u) and e_1 intersect, and denote the endpoints of e_1 by v_1 and u_1 . Clearly |a, p| < l and at least one endpoint of e_1 , say u_1 , lies higher than a.

Imagine a straight-line segment h whose one endpoint is anchored at a, and consider the sweep performed by h as we move its other endpoint along e_1 from p towards u_1 . It is easy to see that h will not cross any vertex as long as the length of h is $\leq 2l$, since there is no diagonal connecting a and a vertex within distance $\leq 2l$ from a, and so $|a, u_1| > 2l$. By triangle inequality on the triangle $T' = (a, p, u_1)$ we thus get that $|p, u_1| > l$, which implies that $|v_1, p| < |e_1| - l \leq l$. So if we consider the triangle $T = (a, p, v_1)$, we see that $|v_1, a| < |e_1| \leq 2l$ (again by triangle inequality).

Now, if the straight-line segment (v_1, a) is not properly intersected by any greedy edge (among those produced so far), then (v_1, a) has to coincide with a produced greedy edge. Otherwise, if some greedy edge properly intersects (v_1, a) , then such an edge has to have (exactly) one endpoint properly in T, because it cannot intersect (v_1, p) nor (p, a) (recall that e_1 was the first edge crossed by a straight-line walk from a to u). In this case we can consider the convex hull of the vertex set consisting of a, v_1 , and all vertices properly in T. For each



Figure 1: The situation just after having produced all greedy edges of length $\leq 2l$.

edge of this convex hull, except the one that coincides with (v_1, a) , it holds that the edge has length < 2l and is not properly intersected by any produced greedy edge, which implies that the edge coincides with a produced greedy edge. Thus we can conclude that either (v_1, a) is a produced greedy edge, or a and v_1 belongs to a concave chain formed by produced greedy edges in T. In the former case we define the *hub* as the set $\{v_1\}$, whereas we in the latter case let the *hub* be the set of vertices except a on the mentioned concave chain. In the remainder we shall assume w.l.o.g. that the vertices of the hub lie south-west of a.

Observation 3.3 Vertex u_1 is visible from a.

Proof If we would hypothesize that some greedy edge properly intersects (a, u_1) , then we can first observe that such an edge must have one endpoint properly inside the triangle $T' = (a, p, u_1)$ (it cannot intersect the other sides of T'). See Figure 2. So there would be a



Figure 2: If u_1 was not visible from a (a dotted line illustrates a bisector of an edge).

vertex u' lying properly inside T' and on the convex hull of the vertices in T'. Now, let v' be

a vertex of the hub, and suppose that u' and v' have been chosen so that they are at least as close to each other as any other such pair. It is easy to show that no greedy edge (among those produced so far) may properly intersect the straight-line segment (v', u'). But if D is the disk centered at v_1 and of radius $|e_1|$, then T' is contained in D (recall that $|v_1, a| < |e_1|$). Thus v_1 is closer to any vertex properly inside T' than to u_1 , which implies that u' and v' are within distance $< |e_1|$ from each other. Therefore a greedy edge that connects u' and v' must have been produced, and so we would have crossed that edge before e_1 by the straight-line walk from a to u, which is a contradiction. \Box

By Observation 3.3 it follows that u_1 lies to the left of a (u_1 would otherwise be closer to a than to v_1 and we would thus have produced (a, u_1)). We also observe that u_1 lies within distance < 3l from a (triangle inequality on T').

Let us now consider the situation just after having produced all greedy edges of length $\leq 3l$. By the facts that u_1 was visible from a and $|a, u_1| < 3l$, it follows that some greedy edge e_2 has been produced which blocks the visibility between a and u_1 and which has length > 2l but not greater than $|a, u_1| < 3l$. Since e_2 properly intersects (a, u_1) we get that e_2 is incident to a vertex of the hub, because there is no vertex properly in T'. We can therefore make calculations for e_2 which are analogous to those we did above for e_1 . In this way we infer that e_2 has an endpoint, say u_2 , which is visible from a, lying north-west of a, and within distance < 4l from a. Consequently, before having produced all greedy edges of length $\leq 4l$, the greedy algorithm must produce a greedy edge e_3 which blocks the visibility between a and u_2 , and so on. Indeed, we can repeat this scenario till the greedy algorithm start to produce greedy edges of length $\geq a_G > 1000l$. (This is roughly illustrated in Figure 3). Thus we realize that the greedy algorithm will for some $m \geq a_G/l - 2 > 998$ produce a sequence e_1, e_2, \ldots, e_m of greedy edges such that each $e_i, i > 1$, has the following properties:

- (i) e_i properly intersects (a, u_{i-1}) ,
- (ii) $il < |e_i| < (i+1)l$, and
- (iii) e_i connects a vertex v_i of the hub with a vertex u_i that lies north-west of a.



Figure 3: The sequence e_1, e_2, \ldots, e_m of greedy edges.

For later purpose, we note that the distance between any pair in $\{a, v_1, v_2, \ldots, v_m\}$ is < 2l by property (iii), since all these vertices are contained in the triangle T with side-lengths < 2l.

In this paragraph we consider the final triangulation, i.e. the situation after having produced all greedy edges in GT(S). If we walk along the straight-line segment (u, a), starting at its intersection with e_{70} and ending at its intersection with e_m , we will cross a set of produced greedy edges. Let E' be a maximal subset of this set of greedy edges such that the following holds: (1) for any two edges in E', their endpoints that do not belong to the hub are different, and (2) the greedy edges $e_{70}, e_{71}, \ldots, e_m$ are included in E'. Let $e'_1, e'_2, \ldots, e'_{m'}$ be an enumeration of E' of the same relative order as they were crossed by the walk. Finally, let $u'_1, u'_2, \ldots, u'_{m'}$ be the endpoints of $e'_1, e'_2, \ldots, e'_{m'}$ which do not belong to the hub, respectively.

Sublemma 3.4 The sequence $u'_1, u'_2, \ldots, u'_{m'}$ forms one or two concave chains.

Proof Let k be any integer in the interval [70, m-2], and let s' be the subsequence of $u'_1, u'_2, \ldots, u'_{m'}$ which starts at u_k and ends at u_{k+2} . It suffices to show that s' forms one concave chain if the angle $\alpha = \angle u_k, a, u_{k+2}$ is ≤ 45 degrees (recall that each vertex u_i lies north-west of a). Suppose therefore that $\alpha \leq 45$ degrees. We shall first show that $|u_k, u_{k+2}| < (k-5)l$. (In the continuation it may help to consult Figure 4.) By properties (ii) and (iii) we have that $kl < |v_k, u_k| < (k+1)l$ and $|v_k, a| < 2l$, from which it follows that $(k-2)l < |a, u_k| < (k+3)l$ (using triangle inequality). In an analogous manner it follows that $kl < |a, u_{k+2}| < (k+5)l$. It is now straightforward to show that $|u_k, u_{k+2}| < (k-5)l$ by using traditional trigonometry on the triangle (u_k, a, u_{k+2}) (recall that $k \ge 70$ and $\alpha \le 45$ degrees).



Figure 4: s' forms a concave chain in Δ .

Let q be the point on e_{k+2} such that the distance between v_k and q equals kl, and let Δ be the triangle (u_k, q, u_{k+2}) . (If such a point q does not exist then we let u_{k+2} play the roll of q, although it will later turn out that q is properly included in e_{k+2}). By property (iii) we have that $|v_k, v_{k+2}| < 2l$, and so $|v_{k+2}, q| > (k-2)l$ (by triangle inequality). Thus, since $|e_{k+2}| < (k+3)l$ by property (ii), it follows that $|q, u_{k+2}| < 5l$, and we thus infer that each side in the triangle Δ has length < kl.

In the remainder of the proof, including the observation below, we consider the situation just after having produced the greedy edge e_k .

Observation 3.5 No greedy edge (among those produced so far) can properly intersect the side (u_k, q) of the triangle Δ .

Proof Imagine a straight-line segment h which always has one endpoint on e_k and one endpoint on e_{k+2} (we refer to e_{k+2} although it has not yet been produced). If h is collinear with (a, u) then no greedy edge intersects h, and any point of h is visible from at least one vertex of the hub (recall that e_k is the first greedy edge crossed by a straight-line walk from a to u). Thus if we start from that position and move h continuously further away from the hub, we realize that h cannot cross any vertex as long as all points of h are within distance $< |e_k|$ from v_k (such a vertex would be visible and within distance $< |e_k|$ from some vertex on the hub and we would thus have produced a greedy edge crossed by a straight-line walk from a to u). This observation tells us that no vertex is properly contained in the region bounded by $e_k, e_{k+2}, (u_k, q)$, and the hub. So it is easy to see that no greedy edge (among those produced so far) can properly intersect (u_k, q) (such an edge would be incident to a vertex of the hub, and so it would be crossed before e_k by a straight-line walk from a to u).

We claim that at least one vertex is properly contained in Δ . Indeed, if this was not the case then there would be no greedy edge properly intersecting the side (u_k, u_{k+2}) of Δ (such an edge must leave one endpoint properly inside Δ since it cannot intersect its other sides by Observation 3.5). Consequently, the side (u_k, u_{k+2}) would coincide with a produced greedy edge (it has length $< lk < |e_k|$), and so we would not be able to produce the greedy edge e_{k+1} as the vertex u_{k+1} is not longer visible from v_{k+1} (recall that no two u_i 's are identical by property (i)). Thus we can consider the convex hull of the vertex set consisting of u_k, u_{k+2} , and all vertices properly inside the triangle Δ . For each edge of that convex hull, except the one which coincides with (u_k, u_{k+2}) , it holds that the edge has length $< kl < |e_k|$ and is not properly intersected by any produced greedy edge, which implies that the edge coincides with some produced greedy edge. Hence, u_k and u_{k+2} belongs to a concave chain formed by produced greedy edges in the triangle Δ . Moreover, no vertex is properly contained in the region bounded by this concave chain, e_k , e_{k+2} , and the hub. Thus the vertex sequence of this concave chain equals the sequence s'.

By Sublemma 3.4, also the sequence $u_{70}, u_{71}, \ldots, u_m$ forms one or two concave chains. Let C stand for any of those possible concave chains. Further, let u_{low} be the endpoint of C that is closest to a, and let u_{upp} be the endpoint of C which is farthest from a.

Sublemma 3.6 For any vertex u_i of the concave chain C, in the minimum weight convex partition of S, the vertex u_i is incident to an edge of length greater than (i-2)l/4 if $(4low + 10)/3 \le i \le (upp - 1)/2$.

Proof We begin the proof by making the following three observations.

Observation 3.7 For any vertex u_i of the concave chain C it holds that $|u_i, u_{upp}| > (upp - i - 3)l$ and $|u_{low}, u_i| > (i - low - 3)l$.

Proof By property (ii) we have that $|e_i| < (i+1)l$ and $|e_{upp}| > upp \cdot l$. By property (iii) we have that the distance between v_i and v_{upp} is < 2l. Thus the distance between u_i and

 u_{upp} is $> |e_{upp}| - |e_i| - |v_i, v_{upp}| > (upp - i - 3)l$. In an analogous manner it follows that $|u_{low}, u_i| > (i - low - 3)l$.

Observation 3.8 For any vertex u_i of the concave chain C, the distance from u_i to the closest endpoint of e_{upp} is greater than (i-2)l if $i \leq (upp-1)/2$.

Proof By properties (ii) and (iii) we have that $|e_i| > il$ and $|v_i, v_{upp}| < 2l$, from which it follows that the distance between u_i and v_{upp} is > (i-2)l (by triangle inequality). Further, by Observation 3.7, we have that the distance between u_i and u_{upp} is > (upp - i - 3)l. It is now straightforward to show that $(upp - i - 3)l \ge (i - 2)l$ for $i \le (upp - 1)/2$.

Observation 3.9 For any vertex u_i of the concave chain C, the distance from u_i to any endpoint of C is greater than (i-2)l/4 if $(4low+10)/3 \le i \le (upp-1)/2$.

Proof By Observation 3.8 it follows that the distance between u_i and u_{upp} is > (i-2)l/4 for $i \leq (upp-1)/2$. By Observation 3.7 we have that the distance between u_{low} and u_i is > (i - low - 3)l, and it is straightforward to show that $(i - low - 3)l \geq (i - 2)l/4$ for $i \geq (4low + 10)/3$.

Let D_i be a disk centered at u_i and of radius (i-2)l/4 (we consider an arbitrary integer i in the interval [(4low + 10)/3, (upp - 1)/2]. Let C' be the subsequence of $u'_1, u'_2, \ldots, u'_{m'}$ which starts at u_{low} and ends at u_{upp} . See Figure 5. By Observation 3.9 the concave chain C' partitions D_i into one convex and one concave region, of which the latter is labeled D'_i . It suffices to show that D'_i can only contain vertices of C', because it follows then that u_i must be incident to an edge of MC(S) which has length > (i-2)l/4 (otherwise there would be a concave angle at u_i in MC(S), and that is only possible if u_i lies on the convex hull of S, which it obviously does not).

Now, if the greedy edge e_{upp} does not intersect D'_i then it follows trivially that D'_i may only contain vertices of C' (consider the simple polygon induced after removing from GT(S)all edges that intersect (u, a) between e_{low} and e_{upp}). We can therefore assume in the continuation that e_{upp} intersects D'_i . So the greedy edge e_{upp} partitions D'_i into two regions, one containing the vertex u_i and the other not containing any vertex of C'. Again it is easy to see that the former region may only contain vertices of C'. Thus it remains to show that the latter region, which we call R, does not contain any vertex. Indeed, if we would hypothesize that there is a vertex in R, then the straight-line segment, call it b, connecting that vertex with u_i would properly intersect the greedy edge e_{upp} . But by Observation 3.8 the distance between u_i and the closest endpoint of e_{upp} is > (i-2)l. Thus it follows from Fact 2.4 (the greedy triangulation is 4-sensitive) that b would have length > (i-2)l/4, which is a contradiction since any point in R is within distance $\leq (i-2)l/4$ from u_i .

Continuation of the proof of Lemma 3.2: It now follows from Sublemma 3.6 that |MC(S)| is greater than

$$\sum_{i=\lceil (4low+10)/3\rceil}^{\lfloor (upp-1)/2 \rfloor} (i-2)l/4$$



Figure 5: An illustration for the proof of Sublemma 3.6.

which is easily shown to be greater than (recall that $low \ge 70$)

$${upp/2-5low/3} \\ (l/4) \sum_{i=0}^{upp/2} i.$$

If $u_{70}, u_{71}, \ldots, u_m$ form only one concave chain, then low = 70 and upp = m, and so

$$|\mathrm{MC}(S)| > (l/4) \sum_{i=0}^{m/2 - 5 \cdot 70/3} i = \Omega(lm^2).$$

In case $u_{70}, u_{71}, \ldots, u_m$ form two concave chains, say u_{70}, \ldots, u_x and u_{x+1}, \ldots, u_m , we consider the former chain if x > m/4, otherwise we consider the latter chain. So if x > m/4 we have that |MC(S)| is greater than

$$\binom{m/8 - 5 \cdot 70/3}{(l/4) \sum_{i=0}^{m/8} i = \Omega(lm^2)}$$

and for $x \leq m/4$ we get

$$|\mathrm{MC}(S)| > (l/4) \sum_{i=0}^{m/2-5(m/4+1)/3} i = \Omega(lm^2).$$

Finally we recall that $m \ge a_G/l - 2$ and $l \le \max(a)$, and so we have shown that $|MC(S)| = \Omega(a_G^2/\max(a))$.

4 A tight bound for the greedy triangulation

This section is entirely devoted to the proof of the following theorem.

Theorem 4.1 For any set S of n vertices (in general position),

$$\frac{|\mathrm{GT}(S)|}{|\mathrm{MT}(S)|} = O(\sqrt{n}).$$

The above theorem is obtained immediately by combining Theorem 3.1, Facts 2.4 and 2.5, and the following lemma.

Lemma 4.2 For any real number r, r > 0, let CP(S) be an arbitrary r-sensitive convex partition of S. Then CP(S) can be triangulated by adding diagonals of total length O(r|CP(S)| + r|MT(S)|).

We shall first prove another lemma (a proof of this lemma can also be found in [7]) which is used in the proof of Lemma 4.2. In order to state this lemma we need the following definitions. A polygon P is called *q*-bent if and only if it has the following three properties: (1) P is convex, (2) P can be drawn within a circle whose diameter equals the length of the longest side, called the *base*, of P, and (3) the sum of degrees of the two interior angles of Pat the endpoints of its base is not greater than 2q degrees. A side of a q-bent polygon which is not the base shall be termed *top-side*. By p(P) we denote the length of the perimeter of a simple polygon P.

Lemma 4.3 For any real number $q, 0 < q \le 45$, there exists some constant c depending on q, such that for any convex polygon P there are edges in GT(P) of total length $\le c \cdot p(P)$ which partition P into triangles and/or q-bent polygons.

Proof Let q be an arbitrary real number, $0 < q \le 45$. The boundary of P can be partitioned into $\le \lceil 180/q \rceil q$ -bent consecutive pieces as follows (we call a piece q-bent if it corresponds to the top-sides of some q-bent polygon). Start with an arbitrary vertex of P and mark it. Then traverse the boundary of P in, say, clockwise order and mark the last vertex for which it holds that the piece of P's boundary between this vertex and the previously marked vertex is q-bent. In this way, when the whole boundary of P is traversed, each marked vertex separates two q-bent pieces, and there are $\le \lceil 180/q \rceil$ such marked vertices. (To prove the lemma it suffices to show that there are O(180/q) marked vertices.)

Let q_1 be any one of these pieces. Let E be a set of edges in GT(P) with minimum cardinality, such that all polygons in the partition of P induced by E are triangles and/or q-bent polygons. To prove this lemma it suffices to show that the total length of all edges in E which are incident to q_1 is O(p(P)) (because then the same can be proved analogously for the other q-bent pieces).

Let v_0, v_1, \ldots, v_m be the vertices of q_1 in clockwise order from one end of q_1 to the other. For easier reference, we assume w.l.o.g. that the straight-line passing through v_0 and v_m is vertical, and that v_m lies above v_0 . From the definition of E it follows that if there is an edge (v_i, v_j) in $E, 0 \le i \le j \le m$, then for any integer l, i < l < j, no edge in E is incident to v_l (because if an edge in E, say e, would be incident to v_l , then $E - \{e\}$ would also partition P into triangles and/or q-bent polygons, and thus E would not be minimal, a contradiction). From this it follows easily that

- (i) the total length of all edges in E that have both their endpoints in q_1 is no more than length of q_1 .
 - To make the proof shorter, we use the following fact.

Proposition 4.4 (Corollary 2.1 in [14]) For any vertex v of P, let e be a longest edge in GT(P) that is incident to v. Then it holds that the total length of all edges in GT(P) which are incident to v is O(|e|).

By Proposition 4.4 we obtain

(ii) the total length of all edges in GT(P) which are incident to v_0 or v_m is O(p(P)).

It remains to show that the total length of those edges in E with only one endpoint in q_1 , except for v_0 and v_m , is O(p(P)). Let E_1 be the set of these edges. Let V' be the set of vertices in q_1 which are endpoints of some edge in E_1 . Let us call these vertices $v'_1, v'_2, \ldots, v'_{k-1}$ in clockwise order. Moreover, let us denote v_1 by v'_0 and v_m by v'_k . For each vertex v in V', denote by E(v) the set of edges in E_1 which are incident to v. By len(L) we denote the total segment length in a set L of straight-line segments. To show that $\sum_{v \in V'} len(E(v)) = O(p(P))$, we associate to each vertex v in V' a unique part of the boundary of P which has length $\Omega(len(E(v)))$ as follows.

Let *i* be the integer such that $v = v'_i$. Let the *root* of *v* be the consecutive piece of q_1 which includes *v*, such that the piece of the root of *v* which is below, respectively above, *v* has length equal to one half of the piece of q_1 which lies between v'_{i-1} , respectively v'_{i+1} , and *v*. (The root of *v* is depicted in Figure 6.) By this definition it is clear that roots of vertices do not overlap.



Figure 6: The root and the fan of v.

If there are at least two edges in E(v), then we associate to v an additional piece of P's boundary, which we call the fan of v, as follows. Let u_0, u_1, \ldots, u_l be the vertices of P in clockwise order from v such that (v, u_i) is in E(v) for $i = 0, 1, \ldots, l$. The fan of v is the piece

of P's boundary which lies between u_0 and u_l (see Figure 6). By this definition it is also clear that no fans of vertices overlap.

To complete the proof of this lemma, it remains to prove that the length of the root of v plus the length of the fan of v, if there is any, is $\Omega(\operatorname{len}(E(v)))$. First we observe that the distance between v'_{i-1} and v'_{i+1} is shorter than two times the length of v's root. On the other hand, from the definition of the greedy triangulation it follows that a shortest edge in E(v) is not longer than the distance between v'_{i-1} and v'_{i+1} , because otherwise the edge (v_{i-1}, v_{i+1}) would be in $\operatorname{GT}(P)$ instead of that edge in E(v). Hence,

(iii) the length of a shortest edge in E(v) is smaller than twice the length of v's root.

It remains to consider the case when there is more than one edge in E(v). Let e be a shortest edge and e' a longest edge of E(v). By Proposition 4.4 the total edge length in E(v) is within a constant factor from the length of e'. Next we observe that the fan of v includes the endpoints of the edges e and e', different from v, and thus the length of the fan is greater than |e'| - |e|. Combining these arguments with (iii), we easily get that the total edge length in E(v) is within a constant factor of the sum of the lengths of the root and the fan of v. \Box

Proof of Lemma 4.2 By Lemma 4.3 there exists a set D of diagonals, each diagonal in D belonging to the greedy triangulation of some convex polygon bounded by CP(S), such that the diagonals in D have total length O(|CP(S)|) and partition CP(S) into triangles and/or 1-bent polygons. Let \mathcal{P} be the set of all 1-bent polygons induced by adding the diagonals in D. We first give the following technical observation.

Sublemma 4.5 For any diagonal d in D it holds that d is (9r + 4)-sensitive.

Proof Let P be the convex polygon bounded by CP(S) such that d belongs to GT(P), and suppose that there is a diagonal a of S that properly intersects d. We shall first consider the case when none of a's endpoints is a vertex of P, that is when a properly intersects two sides of P. For an illustration, see Figure 7.



Figure 7: d is (9r + 4)-sensitive.

Let v be any endpoint of a, and let l be a shortest straight-line segment connecting vand any endpoint of the side of P which is first crossed by a straight-line walk from v to the other end of a. For the other endpoint, call it v', of a we define in a symmetrical way the segment l'. Further, let b be the straight-line segment which connects those endpoints of l and l' that belong to P. Finally, for the endpoint of b which is shared by l, let t be a shortest straight-line segment connecting this endpoint and any endpoint, say u, of d.

Now, since d is 4-sensitive with respect to the vertices of P by Fact 2.4, we have that $|t| \leq 4|b|$. Also, since the sides of P are r-sensitive, we have that $|l| \leq r|a|$ and $|l'| \leq r|a|$. Let h be the distance between v and u. By using triangle-inequality and the above facts, we get that $h \leq |l| + |t| \leq |l| + 4|b| \leq |l| + 4(|l'| + |a| + |l|) \leq r|a| + 4(r|a| + |a| + r|a|) = (9r + 4)|a|$. Thus, if we let d' be the shortest distance between v and any endpoint of d, since $h \geq d'$, it holds that $|a| \geq d'/(9r + 4)$. We can in a symmetrical manner treat the endpoint v' of a, thus concluding that d is (9r + 4)-sensitive with respect to those diagonals whose endpoints are not vertices of P.

For diagonals that have an endpoint in P, we observe that this is a degenerate case of the one considered above. For example, if v is a vertex of P, the segment l becomes a point of zero length, and similarly for v'. Thus we obtain that d is (9r + 4)-sensitive with respect to all diagonals of S.

In the remainder, for any 1-bent polygon in \mathcal{P} and any side s of this polygon, we can by Sublemma 4.5 assume that s is r'-sensitive with r' = 9r + 4 (note that if a diagonal is r-sensitive then it is also r''-sensitive for any $r'' \ge r$).

Given an edge e in MT(S) and a 1-bent polygon P such that two of its sides, say s and s', properly intersect e, we define the *transposal* of e within P as the shortest straight-line segment connecting an endpoint of s and an endpoint of s'. (Note that any other edge in MT(S) that properly intersects s and s' induces the same transposal as e.) Let T be the set of all distinct transposals obtained by considering all edges in MT(S) and all 1-bent polygons in \mathcal{P} .

Sublemma 4.6 The total segment length in T is no more than O(r'|CP(S)| + r'|MT(S)|).

Proof Consider a transposal t in T. Let e be an edge in MT(S) and P the 1-bent polygon such that t is a transposal of e within P. If the following holds:

$$e$$
 does not intersect the base of P , (1)

then, by the shape of a 1-bent polygon, t is not longer than the piece of e that is wholly internal to P. Thus the total length of all transposals in T for which (1) holds is no more than |MT(S)|.

Suppose now that (1) does not hold for t. So e properly intersects the base of P and one top-side, say s, of P. Comment: it is easy to show that any other edge in MT(S) that induces t must also properly intersect s. If the following holds:

$$|t| < 3r'|s|, \tag{2}$$

then we say that t is top-paid. It is easily seen that the total length of all top-paid transposals is no more than 3r'|CP(S)|. We say that a transposal is *mwt-paid* if neither (1) nor (2) holds for this transposal. To complete the proof of this sublemma it suffices now to prove the following claim.

Claim. The total length of all mwt-paid transposals induced by any edge e in MT(S) is no more than O(r'|e|).

Let e be an arbitrary edge in MT(S), and assume w.l.o.g. that e is vertical. We say that a 1-bent polygon P has orientation $k, 0 \le k < 360$, if a clockwise rotation of P by α degrees, $k \le \alpha < k + 1$, results in that the base becomes horizontal and below the top-sides of P. For an arbitrary integer $k, 0 \le k < 360$, let P_1, P_2, \ldots, P_m be the sequence of 1-bent polygons which are crossed by walking along e from its uppermost endpoint to its lowermost endpoint, and such that the following holds for each P_i :

- (a) e induces an mwt-paid transposal t_i within P_i , and
- (b) P_i has orientation k.

To prove the above claim it suffices to show that the total length of e's transposals within P_1, P_2, \ldots, P_m is O(r'|e|), since we consider an arbitrary orientation among 360 possible. We may also assume w.l.o.g. that the following holds for each P_i :

(c) a top-side of P_i , denoted by s_i , was crossed before the base of P_i by our top-down walk,

because we can otherwise consider the sequence in reverse order (i.e. as if we started at the lowermost endpoint of e and walked to its uppermost endpoint). By h_i we denote the piece of e between s_i and s_{i+1} (by h_m we mean the piece of e between s_m and the lowermost endpoint of e). Now, let P_a, \ldots, P_b be any maximal subsequence of P_1, P_2, \ldots, P_m such that the following holds for each $i, a \leq i < b$:

$$|t_i| > 3r'|h_i| \tag{3}$$

Indeed, it suffices to bound the transposals of e within the P_i 's for which (3) holds, because the total length of e's transposals within the other P_i 's is no more than 3r'|e| (since the h_i 's correspond to distinct pieces of e).

Observation 4.7 For any integer $i, a \leq i < b$, it holds that $2|t_i| < |t_{i+1}|$.

Proof Let us first describe the situation. By our top-down walk, we first crossed the top-side s_i of P_i , then the base of P_i , and after that the top-side s_{i+1} of P_{i+1} (in a degenerate case the base of P_i may coincide with s_{i+1}). Since (2) does not hold for t_i , but (3) does, we have that $|s_i| \leq |t_i|/(3r')$ and $|h_i| < |t_i|/(3r')$.

Let l be one of the two straight-line segments which connect an endpoint of s_i with an endpoint of s_{i+1} but do not intersect h_i . An illustration is given in Figure 8. If l does not properly intersect the base of P_i , we realize that l has length $\geq |t_i|$ (recall that the slopes of s_i and s_{i+1} differ by at most 5 degrees). On the other hand, if l properly intersects the base of P_i , since the base is r'-sensitive, the length of l is $\geq |t_i|/r'$ (t_i is a shortest straight-line segment connecting an endpoint of s_i with an endpoint of the base). Thus the length of l is always $\geq |t_i|/r'$ (recall that r' > 4).

Now, if we use triangle-inequality twice on the quadrangle bounded by l, h_i, s_i and s_{i+1} , we get that the piece of s_{i+1} in this quadrangle has length $> |t_i|/(r'-2|t_i|/(3r')) = |t_i|/(3r')$. In an analogous manner we can show that the remaining piece of s_{i+1} has length $> |t_i|/(3r')$.



Figure 8: The t_i 's grow exponentially.

and so the length of s_{i+1} is greater than $2|t_i|/(3r')$. Finally, by combining this observation with the fact that (2) does not hold for t_{i+1} , we conclude that $2|t_i| < |t_{i+1}|$.

By Observation 4.7 it follows that

$$\sum_{i=a}^{b-1} |t_i| < |t_b|.$$

If b < m, since (3) does not hold for t_b in this case, we get that the above sum is $< 3r'|h_b|$. It remains therefore to show that (3) does not hold for t_m when b = m. Indeed, let l be a straight-line segment connecting an endpoint of s_m with the lowermost endpoint of e. As in the proof of Observation 4.7 we can use the r'-sensitivity of P_m 's base to conclude that $|l| \ge |t_m|/r'$. Thus, since $|s_m| \le |t_m|/(3r')$, it follows that $|t_m| \le 1.5r'|h_m|$ (triangle inequality on the triangle bounded by l, s_m and h_m).

Continuation of the proof of Lemma 4.2: Next, given an edge e in MT(S) and a 1-bent polygon P such that e is incident to exactly one vertex of P, we define the *end-transposal* of ewithin P as follows. First we observe that e properly intersects exactly one side, say s, of P. Let u be the endpoint of e which belongs to P, and let v be an endpoint of s which is closest to u. The end-transposal of e within P is the straight-line segment which connects u and v. Since the side s is r'-sensitive, we immediately get that the length of this end-transposal is $\leq r'|e|$. Let T' be the set of all distinct end-transposals obtained by considering all edges in MT(S) and all 1-bent polygons in \mathcal{P} .

Observation 4.8 The total segment length in T' is no more than 2r'|MT(S)|.

Proof Each edge, say e, in MT(S) contributes with at most two end-transposals, both having length $\leq r'|e|$.

Finally, given an edge e in MT(S) and a 1-bent polygon P such that e is incident to exactly two vertices of P, we say that e is *stationary* within P. Let T'' be the set of all stationary edges obtained by considering all edges in MT(S) and all 1-bent polygons in \mathcal{P} . Clearly the total segment length in T'' is no more than |MT(S)|. The following observation is straightforward to show by using the convexity of 1-bent polygons and the planarity of MT(S).

Observation 4.9 No two segments in $T \cup T' \cup T''$ can properly intersect each other.

Thus the segments in $T \cup T' \cup T''$ partition the 1-bent polygons in \mathcal{P} into triangles and/or 1-bent polygons. Let P be an arbitrary 1-bent polygon which remain after this partitioning. If P has a constant number of vertices, then P can clearly be triangulated by adding diagonals of total length O(p(P)). Therefore, to complete the proof of this lemma, it suffices to prove the following claim.

Claim. P is a k-gon for some $k, 3 \le k \le 6$.

Assume contrary that P has more than 6 vertices. We may also assume w.l.o.g. that the base, call it b, of P is horizontal, and that b lies below the top-sides of P. First we consider the case when b coincides with an edge in MT(S) (this case is depicted in Figure 9). Let e and e' be the two edges in MT(S) that form together with b a triangle, such that the endpoint, call it v, that e and e' have in common lies higher than b. If v is a vertex of P, then we see that both e and e' are stationary within P, and so P is a triangle in this case. On the other hand, if v is not a vertex of P, e and e' properly intersect the same top-side of P (because the triangle they form with b may not have any vertex properly inside). Hence, in this case, the end-transposals of e and e' within P force P to be either a triangle or quadrangle. In any case we obtain a contradiction. Thus we can assume in the continuation that b does not coincide with an edge in MT(S).



Figure 9: v is a vertex of P and not.

Now, since b is not in MT(S), there must be an edge e in MT(S) that properly intersects b. Clearly e partitions P into two regions. If each of these two regions has more than two vertices of P on its boundary, we see that the transposal of e within P, or the end-transposal in case e is incident to a vertex of P, is properly inside P, which is a contradiction. Thus among the regions in the partition of P induced by all edges in MT(S) that properly intersect b, there is one region, call it R, which has all vertices of P on its boundary, except for maybe the four endpoints of the two top-sides incident to b. Let u be an endpoint of b which is not in R (u must exist because of e), and let u' be the other endpoint of b. We assume w.l.o.g. that u lies to the left of u'.

Let d be the edge in MT(S) that bounds the leftmost side of R. Now, since the infinitesimal vicinity to the right of d at its intersection with b lies in a triangle Δ of MT(S), and u'may not be properly inside this triangle, we realize that either (1) u' is an endpoint of Δ , or (2) there is an edge d' properly intersecting b such that d' also bounds Δ (and the rightmost side of R). These two cases are illustrated in Figure 10.

In case (1) we see that there is a side d'' of Δ such that d'' connects u' and the endpoint of d that lies higher than b. By the 1-bent shape of P, it is easy to see that the segment in $T \cup T' \cup T''$ which is induced within P by d'' lies on or above d'' (above if it is an end-transposal and on if d'' is stationary). Hence, since there are at least three vertices of P on the boundary of R, we infer that this segment is properly inside P, which is a contradiction. So it remains



Figure 10: The two possibilities for Δ .

only to show that case (2) also yields a contradiction. We first observe that, in case (2), d and d' may not have an endpoint in common that lies higher than b, because the region R would then have at most one vertex of P on its boundary. Thus there is a side s of Δ such that s connects the endpoints of d and d' that lie higher than b. Again, by the 1-bent shape of P and because of the at least three vertices on R's boundary, we realize that s induces a segment in $T \cup T' \cup T''$ that is properly inside P, a contradiction which completes the proof of Lemma 4.2.

5 A new heuristic for the minimum weight triangulation

In this section we give a triangulation algorithm, which is a small modification of the greedy one, and show that it produces a triangulation whose length is within a constant factor from the optimum. The algorithm is as follows (using an ALGOL-like language).

Algorithm: Quasi-Greedy(S)

 $G \leftarrow \{S, \emptyset\}$ **while** G is not a triangulation **do** let (v_1, u_1) be a shortest diagonal of G **if** all the following 6 conditions hold

- 1. the diagonal (v_1, u_1) forms an empty triangle (v_1, u_0, u_1) with two edges in G,
- 2. there is a diagonal (v_0, u_0) properly intersecting (v_1, u_1) and forming an empty triangle (v_0, v_1, u_0) with two edges in G,
- 3. the angle $\angle v_1, u_0, u_1$ is > 135 degrees in triangle (v_1, u_0, u_1) ,
- 4. $|v_0, u_0| < 1.1 |v_1, u_1|,$
- 5. $|v_0, p| < 0.5 |u_0, p|$, where p is the intersection of the straight-line extensions of (v_0, v_1) and (u_0, u_1) , and
- 6. there is an edge (u_1, u_2) in G such that (v_1, u_0, u_1, u_2) forms an empty quadrangle and the angle $\angle u_0, u_1, u_2$ in that quadrangle is > 180 degrees.

then add the edge (v_0, u_0) to G

else add the edge (v_1, u_1) to G end if end while return the triangulation G end Quasi-Greedy



Figure 11: An example of configuration for which all 6 conditions hold.

We call the triangulation of S produced by the above algorithm for the quasi greedy triangulation (because of the 4th condition), and we shall abbreviate it as QGT(S). The quasi greedy convex partition of S, abbreviated QGC(S), is defined in the same way as the greedy convex partition, but for each vertex we select its spokes from QGT(S) instead of GT(S). For a vertex v in S, let v_Q stand for the length of a longest spoke in QGC(S) that was selected for v. The following is the analogy of Observation 2.2.

Observation 5.1 For any set S of vertices, $|QGC(S)| \leq 3 \sum_{v \in S} v_Q$.

Lemma 5.2 Let S be any set of vertices (in general position) and let a be an arbitrary vertex in S. Then $a_Q = O(\max(a))$.

Proof To obtain a contradiction we hypothesize that $a_Q > c \cdot \max(a)$, where c is some sufficiently large constant. Let e be a spoke of a whose length equals a_Q , and let e' and e'' be the two other spokes that were selected for a when defining QGC(S) (a cannot lie on the convex hull of S because of the above hypothesis). We assume w.l.o.g. that e' is horizontal and that a is its right endpoint, and the endpoint of e'' which is different from a lies lower than a. Among those vertices that lie higher than a, let u be the one which is closest to a, and let l be the distance between a and u. We observe that the quasi-greedy algorithm will not produce an edge that connects a with a vertex lying higher than a until it starts to produce edges of length $\geq a_Q > c \cdot \max(a) \geq c \cdot l$ (by our initial hypothesis and because a_Q is the length of a shortest such edge).

Let us denote by G(t) the PSLG whose vertex set equals S and whose edge set consists of all quasi-greedy edges of length $\leq t$. Now, let e_1 be the first edge crossed in G(100l)by a straight-line walk from a to u. Similarly, let e_m be the first edge crossed in $G(c \cdot l)$ by a straight-line walk from a to u. Let e_1, e_2, \ldots, e_m be the sequence of edges crossed in QGT(S) by walking from e_1 to e_m along the line passing through u and a. As in the proof of Lemma 3.2 we can show that each e_i has an endpoint u_i lying, say, north-west of a, and an endpoint v_i lying south-west of a and within distance, say, < 3l from a. We can also show that there is a concave chain between v_1 and a, the so-called hub, such that all the v_i 's belong to the hub.

Now, following the proof of Lemma 3.2, we obtain in the same way that u_1, u_2, \ldots, u_m forms c' concave chains in G(100l), where c' is some constant independent of c.

It is not hard to to show that e_m has length greater than $c \cdot l/1.1 - l$ (the divisor 1.1 comes from condition (4) in the quasi-greedy algorithm). Thus we realize that there are two integers r and t, $1 \le r < t \le m$, such that $u_r, u_{r+1}, \ldots, u_t$ forms a concave chain and $|e_t|/|e_r| = \Omega(c^{1/c'})$.

Let us denote by $N(u_i)$ the first vertex in $u_i, u_{i+1}, \ldots, u_m$ which is not identical to u_i (i.e. the next vertex from u_i on the concave chain). Similarly, define $N(v_i)$ as the first vertex in v_i, \ldots, v_m which is not identical to v_i . Further, denote by p_i the intersection point between the two lines that are collinear with $(u_i, N(u_i))$ and $(v_i, N(v_i))$, respectively.



Figure 12: The angle $a, u_i, N(u_i)$ is > 135 degrees for $s \le i < t$.

Next we observe that, for some sufficiently large constant $s, r \leq s < t$, it holds that the angle $a, u_s, N(u_s)$ is greater than 135 degrees (we choose the smallest possible s with this property, and get in this way that $|e_t|/|e_s| = \Omega(c^{1/c'})$). Thus the point p_s must be within distance, say, $< 2|e_s|$ from $N(v_s)$ (see Figure 12). Moreover, since u_s, \ldots, u_t is concave, it follows that p_i is within distance $< 2|e_s|$ from $N(v_i)$ and the angle $a, u_i, N(u_i)$ is > 135 degrees for all $i, s \leq i < t$. We further observe that there is some edge $e_k, s \leq k < t$, such that e_{k+1} connects vertex v_k with vertex $N(u_k)$, because an edge connecting u_k and a would otherwise have been produced (see also Figure 13). If we choose e_k so that it is, say, at least 10 times longer than e_s , then we have that $|u_i, p_i| < 0.5|N(v_i), p_i|$ for all $i, k \leq i < t$ (such an edge e_k clearly exists if the ratio between $|e_t|$ and $|e_s|$ is sufficiently large, i.e. if the constant c is sufficiently large).



Figure 13: It holds that $|u_i, p_i|$ is less than $0.5|N(v_i), p_i|$ for $k \le i < t$.

Let us now consider the situation when the quasi-greedy algorithm is just about to produce the edge e_{k+1} , that is when e_{k+1} is a shortest diagonal and the quasi-greedy algorithm checks the six conditions for e_{k+1} . If we let $N(v_k), v_k, u_k, u_{k+1}, N(u_{k+1})$ play the role of v_0, v_1, u_0, u_1, u_2 in the quasi-greedy algorithm, then we see that it only remains to show that the fourth condition holds, because the other conditions are already fulfilled by the way we have chosen e_k . Indeed, if we show this, we arrive at the contradiction that the quasi-greedy algorithm will produce $(N(v_k), u_k)$ instead of e_{k+1} .

Let d be the diagonal that connects $N(v_k)$ and u_k . First we note that d has length $< |e_k| + l$ (recall that both d and e_k are within distance < l from a). It suffices therefore to show that $|e_k| + l < 1.1|e_k|$ (since e_k is shorter than e_{k+1}). But this is easily seen to be true from the fact that e_k has length greater than 100l.

Combining the above lemma with Observations 5.1 and 2.1 we obtain

Corollary 5.3 For any set S of vertices (in general position), |QGC(S)| = O(|MC(S)|).

By condition (6) in the quasi-greedy algorithm we see that it will for convex polygons produce the same triangulation as the ordinary greedy algorithm. Hence, by Fact 2.5, we obtain

Observation 5.4 For any convex polygon P, |QGT(P)| = O(|MT(P)|).

In [13] it was shown that the greedy triangulation is 4-sensitive. A similar result was also obtained in [7] and [3]. We shall next generalize this result so that it can be adapted to quasi-greedy triangulations. First we need the following definition. Given a real number r, we say that a triangulation is r-greedy if it can be produced by repeatedly adding a diagonal, such that its length is not greater than r times the length of a shortest diagonal in the partial

triangulation. Notice that the quasi greedy triangulation is 1.1-greedy, whereas the ordinary greedy triangulation is 1-greedy.

Lemma 5.5 Let S be any set of vertices and let ϵ be any real number such that $0 < \epsilon \leq 1$. Then any $(2 - \epsilon)$ -greedy triangulation T of S is $(4/\epsilon)$ -sensitive.

Proof We consider an arbitrary edge (v'_1, v'_2) of T, and an arbitrary diagonal (v_1, v_2) of S which properly intersects (v'_1, v'_2) . It suffices to show that $\min(|v_1, v'_1|, |v_1, v'_2|) \leq (4/\epsilon)|v_1, v_2|$.

Let l' be the straight-line segment connecting v'_1 and v'_2 . For easier reference, we assume w.l.o.g. that l' is vertical, that v_1 lies to the left of l', and that v'_1 lies higher than v'_2 . See Figure 14. Let d be the distance between v_1 and v_2 . To prove the theorem, we hypothesize throughout the proof that both $|v_1, v'_1|$ and $|v_1, v'_2|$ are greater than $4d/\epsilon$, and derive a contradiction from this hypothesis.



Figure 14: An example of configuration.

Let us denote by G the PSLG whose vertex set equals S and whose edge set consists of all edges in T of length $\leq 4d/\epsilon - 2d$. By the definition of $(2 - \epsilon)$ -greedy triangulations, it follows that all diagonals of G have length $> 2d/\epsilon$.

To go on we need the following definition. A path in the plane is called *free* if it does not include any vertex nor any edge of G. A path P is called *collision-free* if for each real r greater than zero there exists a free path P' whose length is not greater than r plus the length of P, and every point in P is within distance r from some point in P'.

Sublemma 5.6 Let u and v be any two vertices on either side of l', and let P be any collision-free path connecting u and v. If all points on P have greater y-coordinate than that of v'_2 but smaller than that of v'_1 , then P has length $> 2d/\epsilon$.

Proof Let P' be the shortest collision-free path connecting u and v such that no vertices of S lie in any open region bounded by P and P' (if there is any such region).

A way to imagine P and P' is to think about the vertices and edges in G as being obstacles ("pins and walls") and about P as being a wire. Then P' has the shape of the wire which would be obtained if it would be "stretched" between u and v (without being able to cross any obstacles).

Now, if all points on P would lie higher than v'_2 and lower than v'_1 , since the endpoints of P lie on either side of l', it is easy to see that P' would properly intersect l'. Thus there would be two vertices u' and v' on P' such that no edge of G properly intersects (u', v') and (u', v') properly intersects l'. So (u', v') would be a diagonal of G and thus of length $> 2d/\epsilon$ (recall that l' coincides with an edge in T, so (u', v') cannot be an edge).

Let I be the intersection point of l' with (v_1, v_2) . Notice that the distance between I and any endpoint of l' is $> 4d/\epsilon - d$ (by triangle inequality and our initial hypothesis). Hence, since the edges in G have length $\leq 4d/\epsilon - 2d$ and $|v_1, v_2| = d$, we obtain

Observation 5.7 If e is an edge of G that intersects (v_1, v_2) , then all points on e lie higher than v'_2 , lower than v'_1 , and on the same side of l'.

Sublemma 5.8 In G there is an edge properly intersecting (v_1, I) and an edge properly intersecting (v_2, I) .

Proof We prove the statement for (v_2, I) . The proof for (v_1, I) is symmetrical. Let us hypothesize that there is no edge in G that properly intersects (v_2, I) . Thus there must be an edge in G which properly intersects (v_1, I) (since (v_1, v_2) is too short to be a diagonal of G). Let e_1 be the edge in G with the rightmost intersection with (v_1, I) , and let I_1 be that intersection. Next, let v_{e1} be an endpoint of e_1 which is closest to I_1 . By Observation 5.7, both endpoints of e_1 lie to the left of l'. Thus there is a collision-free path from v_{e1} to v_2 of length $< 2d/\epsilon$ consisting of the segment (v_{e1}, I_1) , of length $\le 2d/\epsilon - d$, and (I_1, v_2) , of length < d. By Observation 5.7 all points on this path lie higher than v'_2 and lower than v'_1 , and so we obtain a contradiction to Sublemma 5.6.

Define e_1 as in the proof of Sublemma 5.8, and let e_2 be the edge in G with the leftmost intersection with (I, v_2) . Let v_l and v_r be the uppermost endpoint of e_1 and e_2 , respectively. We assume w.l.o.g. that the straight-line extensions of e_1 and e_2 are parallel or intersect each other above (v_1, v_2) .

Now, let h be the (unique) straight-line segment parallel with (v_1, v_2) , such that h connects v_l or v_r with some point p on e_1 or e_2 (p may possibly coincide with either v_l or v_r). Finally let Q be the quadrangle bounded by $h, (v_1, v_2), e_1$ and e_2 .

Sublemma 5.9 There is a vertex of S properly in Q.

Proof Indeed, if we assume that there is no vertex properly in Q, then there is no edge of G properly intersecting h (such an edge would leave an endpoint properly in Q since it cannot intersect the other sides of Q). Thus if v_p is an endpoint closest to p of the edge (either e_1 or e_2) that p lies on, then there is a collision-free path of length $\leq 2d/\epsilon$ consisting of the segment h, of length $\leq d$, and (p, v_p) , of length $\leq 2d/\epsilon - d$. Hence, by Observation 5.7, we obtain a contradiction to Sublemma 5.6.

Among the vertices in Q (at least one must exist by Sublemma 5.9), let q' be a vertex which is closest to the line (v_1, v_2) . Further, let h' be the straight-line segment parallel with (v_1, v_2) and crossing l', such that h' connects q' with some point q on either e_1 or e_2 . Finally, let v_q be an endpoint closest to q of the edge (either e_1 or e_2) that q lies on. Consider now the collision-free path of length $\leq 2d/\epsilon$ consisting of the segments h' (of length $\leq d$) and (q, v_q) (of length $\leq 2d/\epsilon - d$). By using Observation 5.7, it is straightforward to realize that all points on this path lie higher than v'_1 and lower than v'_2 . Moreover, since q' and v_q lie on either side of l', we obtain, again, a contradiction to Sublemma 5.6.

Finally, combining Corollary 5.3, Lemmata 5.5 and 4.2, and Observation 5.4, we obtain

Theorem 5.10 For any set S of vertices (in general position), |QGT(S)| = O(|MT(S)|).

6 Generalizations and extensions

6.1 Constrained cases

Consider the following problem. Given a PSLG G with vertex set S and edge set E, find a triangulation T of S whose edge set contains E and such that |T| is minimized. The triangulation T is called the *constrained minimum weight triangulation* of G. The quasigreedy algorithm can also be used in order to approximate T, by adding the edges in E to the partial triangulation before the while-loop is executed.

To show that the quasi-greedy algorithm gives a constant-factor approximation also in the constrained case, we can prove the analogy of Lemma 5.2 roughly as follows. Again we consider a half-plane H bounded by a line passing through a vertex a, and we hypothesize that there is a vertex u lying in H, visible from a (i.e. no edge in E intersects (a, u)), and within distance l from a, although the quasi-greedy algorithm connects a with vertices in Honly by edges much longer than l.

As in the proof of Lemma 3.2 we can show that the quasi-greedy algorithm has to produce an edge e_1 which blocks the visibility between a and u, and that e_1 has an endpoint u_1 lying in H and an endpoint v_1 lying in the complement of H. To see that u_1 is visible from a, let pbe the intersection point between e_1 and (a, u). If we choose e_1 sufficiently long we get that the angle p, u_1, a is quite small, and so u_1 is the point in the triangle (p, u_1, a) farthest from v_1 . Now, if any edge (either in E or produced by the algorithm) would properly intersect (a, u_1) , then it would leave an endpoint in the above mentioned triangle, which cannot be the case since the quasi-greedy algorithm would connect v_1 with that endpoint rather than with u_1 . Thus u_1 is visible from a, which means that its visibility from a has to be blocked by an edge e_2 which is only slightly longer than e_1 , and so on. Proceeding in this way, we arrive at edges e_i and e_{i+1} of fairly great length, such that e_i and u_{i+1} lying in H. We can now show that e_{i+1} fulfills the six conditions in the quasi-greedy algorithm, and so we obtain the contradiction that we will produce an edge that properly intersects e_{i+1} .

A problem that still remains is to adapt Lemma 4.2 to the constrained case (since there does not necessarily exist a constant c such that the edges of G are c-sensitive with respect to diagonals of S, we also have that the constrained quasi greedy convex partition might not be c-sensitive in this sense). This can be done roughly as follows. First we observe that

we can, for some constant c, still prove that all edges in our triangulation are c-sensitive with respect to diagonals of G (i.e. diagonals of S that do not properly intersect any edge in E) by following the lines of proof of Lemma 5.5 (this was shown for the ordinary greedy triangulation in [13]). Let CP(G) be the (constrained) quasi greedy convex partition of G. If CP(G) is r-sensitive with respect to diagonals of S, we have by Lemma 4.2 that CP(G) can be triangulated by adding diagonals of total length O(r|CP(G)| + r|MT(S)|). However, in the proof of Lemma 4.2, the r-sensitivity was used on an edge of CP(G) only with respect to minimum weight edges properly intersecting it. Therefore, since the edges of the constrained minimum weight triangulation cannot properly intersect edges of G, it follows that when we use the r-sensitivity of an edge with respect to some diagonal, we have that this edge is c-sensitive with respect to that diagonal. Thus we can still show that CP(G) can be triangulated by adding diagonals whose total length is proportional to |CP(G)| plus the length of the constrained minimum weigh triangulation.

6.2 Approximating the minimum weight convex partition

The quasi-greedy algorithm can also be used in order to compute the quasi greedy convex partition, thus obtaining a constant-factor approximation of the minimum weight convex partition. Indeed, it suffices to find for each vertex v in our input vertex set S the spokes of v in QGT(S).

Let m be the number of quasi-greedy edges that are incident to v, and let L be a list consisting of these edges. To simplify the exposition, we first show that the spokes of vcan be found in O(m) time under the assumption that the edges in L are sorted according to their lengths. In addition, we assume throughout this subsection that all diagonals of Shave distinct lengths. We will use the following observation, which follows directly from the definition of spokes.

Observation 6.1 Let e and e' be any two edges incident to v, and let R be the convex region which is bounded by e, e' and an infinitesimal circle centered at v. Then any edge intersecting R cannot be a spoke of v if it is longer than both e and e'.

For any set Y consisting of three edges incident to v, the edges in Y partition the infinitesimal vicinity of v into three open regions, of which we denote by $\max(Y)$ one whose interior angle at v is largest. The initial step is to remove the first (shortest) three edges from L. Let Y be a set consisting of these three edges. If $\max(Y)$ is a convex region, then it is not hard to see that the edges in Y constitute the spokes of v. Otherwise, it follows from Observation 6.1 that any spoke of v has to intersect or bound $\max(Y)$. Thus we can find the spokes of v by executing the following two steps as long as $\max(Y)$ is not convex.

- 1. Remove from Y the edge that does not bound $\max(Y)$.
- 2. Remove the shortest edge from L and insert it into Y.

Instead of using the sorted list L (which can be constructed in $O(m \log m)$ time) we can use any sorted list L' with the property that it contains the spokes of v. Such a list L' can be constructed in O(m) time as follows. Let e_1 be the shortest edge incident to v, and let e_2, e_3, \ldots, e_k be the maximal sequence of edges encountered (in this order) by scanning the edges around v in clockwise direction from e_1 , such that the clockwise angle from e_1 to e_k is < 180 degrees. Given any two edges e_i and e_j , $1 \le i < j \le k$, it follows from Observation 6.1 that any edge e_t , i < t < j, which is longer than both e_i and e_j cannot possibly be a spoke of v. By scanning e_1, e_2, \ldots, e_k we can remove all such edges, thus creating a list such that the edges in this list are sorted with respect to their lengths. In an analogous manner we can create a sorted list for the counter clockwise case. Then, by merging these two lists we obtain one sorted list L' containing all spokes of v. The total time for constructing L' in this way is clearly O(m) (we assume that QGT(S) is represented as a doubly connected edge list). We summarize this subsection in the following theorem.

Theorem 6.2 Let S be any set of n vertices (in general position). Given QGT(S), we can compute QGC(S) in O(n) time.

6.3 Degenerate cases

Throughout the paper we have assumed that no three vertices in S are collinear. In order to remove this assumption we have to add some details, of which the main one concerns the convex partitions that we use. In a degenerate case, we can define GC(S) and QGC(S) in such a way that the angle between two spokes of any vertex, say v, that lie next to each other around v is < 180 degrees (thus it may happen that we for some vertices have to select four spokes). In this way, the statement in Lemma 4.2 still holds for GC(S) and QGC(S). Similarly we can require that all angles in MC(S) are strictly less than 180 degrees. Redefining MC(S), GC(S) and QGC(S) in this way, the proofs of Theorems 4.1 and 5.10 can be adapted to degenerate cases by adding some non-difficult details.

However, if the purpose is to approximate the minimum weight convex partition by using the quasi greedy convex partition, then we must allow angles in QGC(S) to be exactly 180 degrees in order to obtain a constant approximation ratio (for example, in the case when MC(S) forms one triangle with n - 1 vertices on a side which is quite short compared to the other sides). For this reason, we select for each vertex v the spokes as described in Section 2 unless the following holds: There exist two edges e and e' incident to v, collinear with each other, such that in one of the half-planes bounded by them, no edge incident to v is shorter than both e and e'. In this case we select only two spokes for v, namely e and e'. Defining QGC(S) in this way, the statement in Lemma 5.2 still holds. Indeed, under the assumption that a_Q is greater than max(a), it still holds that one of a's spokes (partially) bounds an open half-plane H such that no edge in H that is incident to a has length $\langle a_Q$, although there is a vertex in H which is within distance max(a) from a.

7 Conclusion

Most greedy algorithms can be modified so that they compute QGT(S), since the six conditions can be checked locally and in constant time. The best worst-case time would be attained by modifying the algorithm in [9]. In this way we would obtain an algorithm that computes the quasi greedy triangulation in $O(n \log n)$ time (or in O(n) time if the Delaunay triangulation is given). **Acknowledgment** We are grateful to Prof. David Eppstein for pointing out an error in an earlier draft.

References

- R. C. Chang and R. C. T. Lee. On the average length of Delaunay triangulations. BIT, 24(3):269-273, 1984.
- [2] G. Das and D. Joseph. Which triangulations approximate the complete graph. In Proceedings of the International Symposium on Optimal Algorithms, volume 401 of Lecture Notes in Computer Science, pages 168-183. Springer-Verlag, 1989.
- [3] M. T. Dickerson, R. L. S. Drysdale, S. A. McElfresh, and E. Welzl. Fast greedy triangulation algorithms. In *Proceedings of the Tenth Annual ACM Symposium on Computational Geometry*, pages 211–220, 1994.
- [4] P. D. Gilbert. New results in planar triangulations. Master's thesis, University of Illinois, Urbana, Illinois, 1979.
- [5] L. Heath and S. Pemmaraju. New results for the minimum weight triangulation problem. Algorithmica, 12(6):533-552, 1994.
- [6] D. G. Kirkpatrick. A note on Delaunay and optimal triangulations. Information Processing Letters, 10(3):127-128, 1980.
- [7] C. Levcopoulos. Heuristics for Minimum Decomposition of Polygons. PhD thesis, Linköping University, Linköping, Sweden, 1987.
- [8] C. Levcopoulos. An $\Omega(\sqrt{n})$ lower bound for the nonoptimality of the greedy triangulation. Information Processing Letters, 25(4):247-251, 1987.
- [9] C. Levcopoulos and D. Krznaric. The greedy triangulation can be computed from the Delaunay in linear time. Technical Report LU-CS-TR:94-136, Department of Computer Science, Lund University, Lund, Sweden, 1994.
- [10] C. Levcopoulos and D. Krznaric. Tight lower bounds for minimum weight triangulation heuristics. Technical Report LU-CS-TR:95-157, Department of Computer Science, Lund University, Lund, Sweden, 1995.
- [11] C. Levcopoulos and A. Lingas. On approximating behavior of the greedy triangulation for convex polygons. *Algorithmica*, 2(2):175-193, 1987.
- [12] C. Levcopoulos and A. Lingas. The greedy triangulation approximates the minimum weight triangulation and can be computed in linear time in the average case. Technical Report LU-CS-TR:92-105, Department of Computer Science, Lund University, Lund, Sweden, 1992. A preliminary version appeared in Proc. ICCI '91, LNCS 497.

- [13] C. Levcopoulos and A. Lingas. C-sensitive triangulations approximate the minmax length triangulation. Technical Report LU-CS-TR:93-118, Department of Computer Science, Lund University, Lund, Sweden, 1993. A preliminary version appeared in Proc. FST-TCS '92, LNCS 652.
- [14] A. Lingas. A new heuristic for the minimum weight triangulation. SIAM Journal on Algebraic and Discrete Methods, 8:646-658, 1987.
- [15] E. L. Lloyd. On triangulations of a set of points in the plane. In Proceedings of the 18th Annual IEEE Symposium on Foundations of Computer Science, pages 228-240, 1977.
- [16] G. K. Manacher and A. L. Zobrist. Neither the greedy nor the Delaunay triangulation of a planar point set approximates the optimal triangulation. *Information Processing Letters*, 9(1):31-34, 1979.
- [17] D. A. Plaisted and J. Hong. A heuristic triangulation algorithm. Journal of Algorithms, 8(3):405-437, 1987.
- [18] F. P. Preparata and M. I. Shamos. Computational Geometry: An Introduction. Springer-Verlag, New York, 1985.
- [19] W. D. Smith. Studies in Computational Geometry Motivated by Mesh Generation. PhD thesis, Princeton University, 1989.
- [20] P. Yoeli. Compilation of data for computer-assisted relief cartography. In J. C. Davis and M. J. McCullagh, editors, *Display and Analysis of Spatial Data*. Wiley, 1975.