On Constrained Minimum Pseudotriangulations

Günter Rote[1*], Cao An Wang[2**], Lusheng Wang[3***], and Yinfeng Xu[4†]

[1] Institut für Informatik, Freie Universität Berlin, Takustraße 9, D-15195 Berlin, Germany, rote@inf.fu-berlin.de

[2] Department of Computer Science, Memorial University of Newfoundland, St. John's, NFLD, Canada A1B 3X8, wang@cs.mun.ca

[3] Department of Computer Science, City University of Hong Kong, Kowloon, Hong Kong, P.R. China, lwang@cs.cityu.edu.hk

[4] School of Management, Xi'an Jiaotong University, Xi'an, 710049, P.R. China, yfxu@xjtu.edu.cn

Abstract. In this paper, we show some properties of a pseudotriangle and present three combinatorial bounds: the ratio of the size of minimum pseudotriangulation of a point set S and the size of minimal pseudotriangulation contained in a triangulation T, the ratio of the size of the best minimal pseudotriangulation and the worst minimal pseudotriangulation both contained in a given triangulation T, and the maximum number of edges in any settings of S and T. We also present a linear-time algorithm for finding a minimal pseudotriangulation contained in a given triangulation. We finally study the minimum pseudotriangulation containing a given set of non-crossing line segments.

1 Introduction

A *pseudotriangle* is a simple polygon with exactly three vertices where the inner angle is less than π, see Figure 1. These three vertices are called *corners*. The boundary is composed of three pieces of nonconvex chains, where the nonconvex chain has either a reflex inner angle at each inner vertex or is a single edge (the degenerate case). A *pseudotriangulation* of a point set S is a partition of the interior of the convex hull of S into a set of pseudotriangles. This geometric structure plays an important role in planning collision-free paths among polyhedral obstacles [4] and in planning non-colliding robot arm motion [2, 5]. Previous

* Research by Günter Rote was partly supported by the Deutsche Forschungsgemeinschaft (DFG) under grant RO 2338/2-1.

** The work of Cao An Wang is supported by NSERC grant OPG0041629.

*** Lusheng Wang is fully supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China [Project No. CityU 1087/00E].

† The work of Yinfeng Xu is supported by NSFC(19731001) and NSFC(70121001) for excellent group project. The initial collaboration between Yinfeng Xu and Günter Rote on this paper was sponsored by the graduate program "Graduiertenkolleg *Algorithmische Diskrete Mathematik*" of the Deutsche Forschungsgemeinschaft (DFG), grant GRK219/3, when Yinfeng Xu visited Berlin.

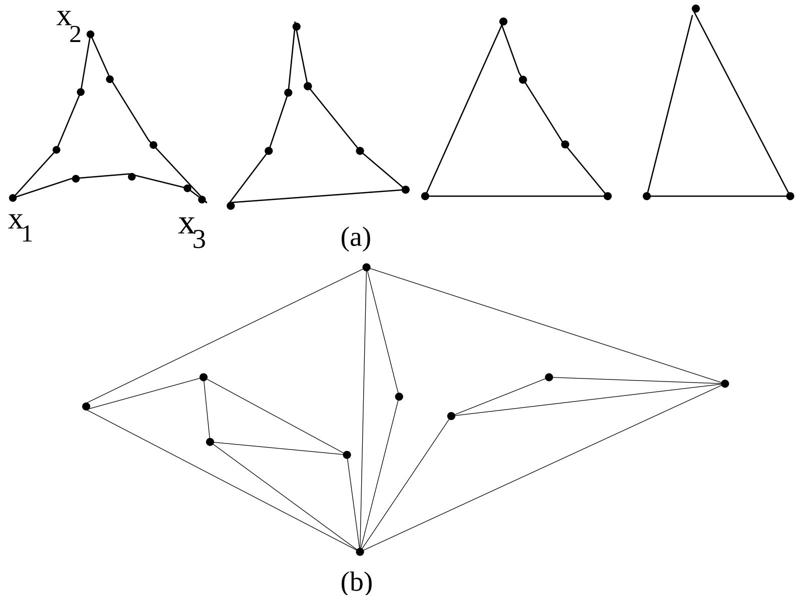

Fig. 1. (a) Some typical pseudotriangles. Vertices x_1, x_2, x_3 are corners, and the three right-hand side cases are degenerate cases of pseudotriangle. (b) A pseudotriangulation of 10 points.

research on this topic was mainly concentrated on the properties and algorithms for minimum pseudotriangulation of a given point set or a set of convex objects. In those cases, the edges of pseudotriangulations are chosen from the complete edge set of the point set.

It is natural to consider some constraint on the choice of edges. Our work mainly investigates the properties of the minimal pseudotriangulations constrained to be a *subset* of a given triangulation, the minimum pseudotriangulations constrained to be a *superset* of a given set of noncrossing line segments, and on algorithms to find these pseudotriangulations. This investigation is motivated in some applications that one may compromise a minimal pseudotriangulation by a faster construction algorithm, or the environment may be constrained by a set of disjoint obstacles. For example, the paper [1] investigates degree bounds for minimum pseudotriangulations which are constrained by some given subset of edges.

In order to find a minimal pseudotriangulation constrained in a given triangulation, one must be able to identify the edges to be removed. In Section 3, we show a structural property for these edges (Theorem 2). This property allows us to design a linear-time algorithm for finding a minimal pseudotriangulation, which is presented in Section 5.

In contrast to the pseudotriangulation of a set S of n points, where all minimum pseudotriangulations of S have the same cardinality, viz. $2n-3$ edges [5], the size of the minimum pseudotriangulation constrained in a given triangulation T depends not only on n, but on T.

We investigate the possible sizes of minimal and minimum pseudotriangulations in Section 4. We show that the ratio of the sizes of the best and the worst minimum pseudotriangulation constrained in some T against the size of the minimum pseudotriangulation triangulation of S can vary from 1 to $\frac{2}{3}$. The above bound is optimal asymptotically. Furthermore, the size of a 'minimal' pseudotriangulation constrained in a triangulation depends on the sequence of construction of pseudotriangles. (In a minimal pseudotriangulation, each pseudotriangle has been expanded into its limit, a further expansion will violate the definition of pseudotriangle. A minimal pseudotriangulation may not be minimum with respect to all possible pseudotriangulations constrained in that triangulation.) We show that the ratio of the size of the smallest minimal pseudotriangulation and the size of the largest minimal pseudotriangulation constrained in a same triangulation can vary from 1 to $\frac{2}{3}$. It is known that the size of minimum pseudotriangulation constrained on any setting of S and T is at least $2n-3$. We show that the *maximum* number of edges in such pseudotriangulations is bounded by $3n-8$.

In Section 6, we study the pseudotriangulations which *contain* a given set L of noncrossing line segments. Interestingly, we find that the size of a minimum pseudotriangulation for L depends only on the number of reflex vertices of L. The proof uses an algorithm for constructing such a minimum pseudotriangulation.

Finally, we discuss some open questions.

2 Preliminaries

We shall first give some definitions. A *triangulation* T of a planar point set S is a maximal planar straight-line graph with S as vertices. We assume throughout the paper that the points of S lie in general position, i.e., no three points lie on a line, and all angles are different from π.

Let T' be a subgraph of T. For a vertex $p \in S$ define $\alpha(p)$ be the largest angle at p between two neighboring edges incident to p. A vertex p in T' is called a *reflex* point if $\alpha(p) \geq \pi$ in T'.

A *minimum* pseudotriangulation of a point set is one with the smallest number of edges. It is known that the number of edges in any minimum pseudotriangulation of n points is $2n-3$, see [5].

We now prove some properties for a triangulated pseudotriangle.

Let p be a pseudotriangle, $T(p)$ be a triangulation of p. Let $T(p)-p$ denote the remainder of $T(p)$ after the removal of the edges of p. The *dual graph* of $T(p)$ is defined as usual: Each node in the graph corresponds to a triangle face in $T(p)$, and two nodes determine an edge of the graph if the corresponding triangles share an edge. A *star-chain* consists of three simple chains sharing a common end-node.

Lemma 1. *The dual of any triangulation of a pseudotriangle is a simple chain or a star-chain.*

Proof. See Figure 2a for an illustration. Each interior edge of the triangulation of a pseudotriangle must span on two different chains by the nonconvexity of its three chains. This implies that these interior edges can form at most one triangle. The lemma follows. □

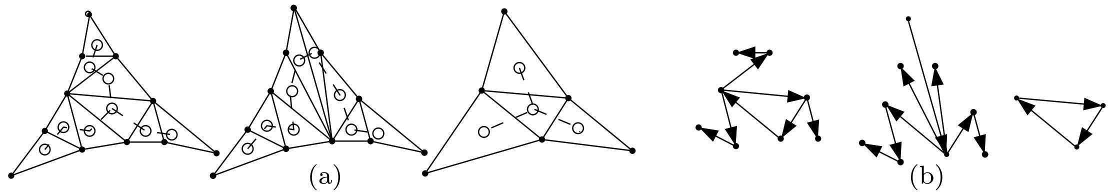

Fig. 2. (a) Different shapes of the dual graph in Lemma 1. (b) The edges of $T(p) - p$ in Lemma 2.

Lemma 2. *Let $T(p)$ be a triangulation of a pseudotriangle p. There is a perfect matching between the edges in $T(p)-p$ and the reflex vertices of p, which matches each edge to one of its vertices.*

Proof. By Lemma 1, the edges of $T(p) - p$ form either a tree which contains exactly one corner of p or a graph with a single cycle, which is formed by a triangle, see Figure 2b. In the first case, we choose the corner as a root and direct all edges of $T(p)-p$ away from the root. Then every reflex vertex will have one edge of the tree pointing towards it, thus establishing the desired one-to-one correspondence between the edges and the reflex vertices. If $T(p) - p$ contains a triangle, we orient the edges of the triangle cyclically, in any direction, and we orient all other edges away from the cycle. Again, every reflex vertex has one edge of the tree pointing towards it. (In fact, the matching between edges and reflex vertices is unique up to reorienting the central triangle.) □

We can extend the statement of the Lemma 2 from a single pseudotriangle to a pseudotriangulation.

Theorem 1. *Let T be a triangulation of a point set S, and let $P \subseteq T$ be a pseudotriangulation of S. Then there is a perfect matching between the edges in $T - P$ and the reflex vertices of P, which matches each edge to one of its two vertices.*

Proof. Every reflex vertex of P belongs to exactly one pseudotriangle in which it is a reflex vertex. Thus, we can simply apply Lemma 2 to each pseudotriangle of P separately. □

The following statement is important for our characterization of minimal pseudotriangulations in Theorem 2.

Lemma 3. *Let p be a pseudotriangle, and let E be a nonempty set of edges inside p which partition p into smaller pseudotriangles. Then one of the following two cases holds:*

(a) *E is a triangle.*
(b) *E contains an edge e such that $E - e$ still partitions p into smaller pseudotriangles.*

Proof. Every edge in E connects two different reflex side chains of p. If $|E| \geq 4$, then E contains at least two edges which connect the same pair of reflex side chains of p. We choose among all these edges the edge e which is incident with the pseudotriangle containing the common corner of these chains. Removing e will join two pseudotriangles into a new face which is bounded by portions of two reflex chains and a single edge between these chains. Hence this face is a pseudotriangle, and e is the desired edge for Case (b) of the lemma.

We are left with the case that E contains at most three edges. This case can be treated by an elementary case analysis. □

3 Minimal pseudotriangulations

Let T denote a triangulation of S and let P^T denote a pseudotriangulation constrained in T, i.e., $P^T \subseteq T$.

A pseudotriangulation P^T is *minimal* (denoted by P^T_{mal}) if no proper subset of P^T is a pseudotriangulation. P^T is called *minimum* (denoted by P^T_{mum}) if it contains the smallest number of edges over all possible pseudotriangulations constrained in T. For simplicity, we use 'constrained pseudotriangulation P^T' as pseudotriangulation constrained in a given triangulation T.

The definition of a minimal triangulation involves a statement about all subsets of edges. The following theorem shows that is is sufficient to check only a linear number of proper subsets to establish that a pseudotriangulation is minimal.

Theorem 2 (Characterization of minimal pseudotriangulations).
A pseudotriangulation P is minimal if and only if

- *there is no edge $e \in P$ such that $P - e$ is a pseudotriangulation, and*
- *there is no triangular face $\{e_1, e_2, e_3\} \in P$ such that $P - \{e_1, e_2, e_3\}$ is a pseudotriangulation.*

Proof. It is clear that the condition is necessary. Now, suppose that $P' \subset P$ is a pseudotriangulation which is a proper subset of P. We have to show that some edge or triangle of P can be removed. Let p be a pseudotriangle face of P' which contains some edges E of $P - P'$. These edges subdivide p into pseudotriangles, and we can apply Lemma 3 to p. We either get an edge whose removal yields a pseudotriangulation, or E is a triangle, whose removal merges 4 faces of P into p. □

4 Ratio of the sizes of pseudotriangulations

In this section, we show some relationships among the sizes of T, P^T (constrained pseudotriangulation), P^T_{mal} (minimal P^T in T), P^T_{mum} (minimum P^T in T), and $P_{\mathrm{mum}}(S)$ (minimum pseudotriangulation of the point set S).

Theorem 3. *Let S be a set of n points in general position and T be a triangulation of S. The number of edges in P^T_{mum} is at most $3n-8$, for $n \geq 5$. There are infinitely many values of n for which a triangulation exists where P^T_{mum} has $3n-12$ edges.*

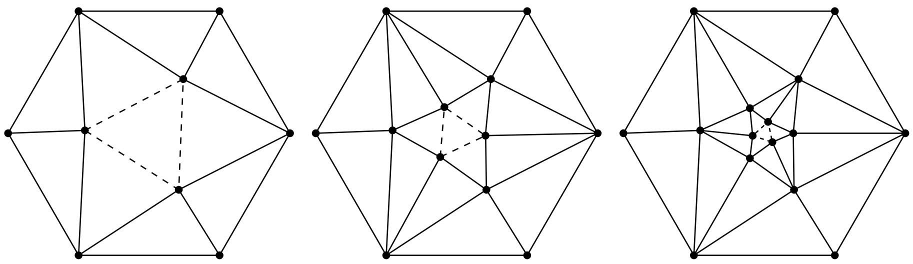

Fig. 3. Three steps of the inductive construction in Theorem 3. The three edges of the dotted central triangle can be removed.

Proof. Suppose that If k vertices lie on the convex hull of S, every triangulation T has $3n-k-3$ edges, and every pseudotriangulation P (in fact, any noncrossing set of edges) has at most $3n-k-3$ edges. This follows from Euler's relation. Thus, when $k \geq 5$, the upper bound follows. It is easy to check that when $n \geq 5$ and k is 3 or 4, we can always remove at least $5-k$ edges and still obtain a pseudotriangulation.

A family of triangulations which show the lower bound is given in Figure 3. The number of vertices is a multiple of 3 and $k=6$. The instances are constructed inductively, by removing the central triangle and subdividing the resulting pseudotriangle as shown in Figure 3. The new points are slightly twisted about the center in order to obtain a point set in general position, and to ensure that the "direct paths" which lead from the center to the vertices of the outer hexagon make zigzag turns. The only edge set which one can remove is the central triangle. The resulting pseudotriangulation has $3n-12$ edges. One can check by inspection, using Theorem 2, that it is a minimal pseudotriangulation. Since there was only one way to obtain a pseudotriangulation as a subgraph of T, it is the unique minimal pseudotriangulation. Hence, is is also a minimum pseudotriangulation. □

We have an example of a minimum pseudotriangulation with $n=41$ vertices and $3n-8$ edges. We believe that the upper bound of $3n-8$ is tight for infinitely many values of n.

Theorem 4. (a) *There are cases of S and T such that the size of T and P^T_{mum}, and all other pseudotriangulations P^T are the same.*

(b) *The ratio between the sizes of two different minimal constrained pseudotriangulations in a given triangulation is between $\frac{2}{3}$ and $\frac{3}{2}$. These bounds are asymptotically tight.*

(c) *The ratio of the size of the minimum pseudotriangulation of S and the minimum pseudotriangulation constrained in T is between 1 and $\frac{2}{3}$, which is asymptotically tight. The same bound holds for the size of the minimal constrained pseudotriangulation in T.*

Proof. The bounds on the ratios follow from the fact that a pseudotriangulation of n points has between $2n-3$ and $3n-6$ edges. We omit the detailed general proofs that the bounds are tight in this version of the paper, but we show some typical tight instances in Figure 4.

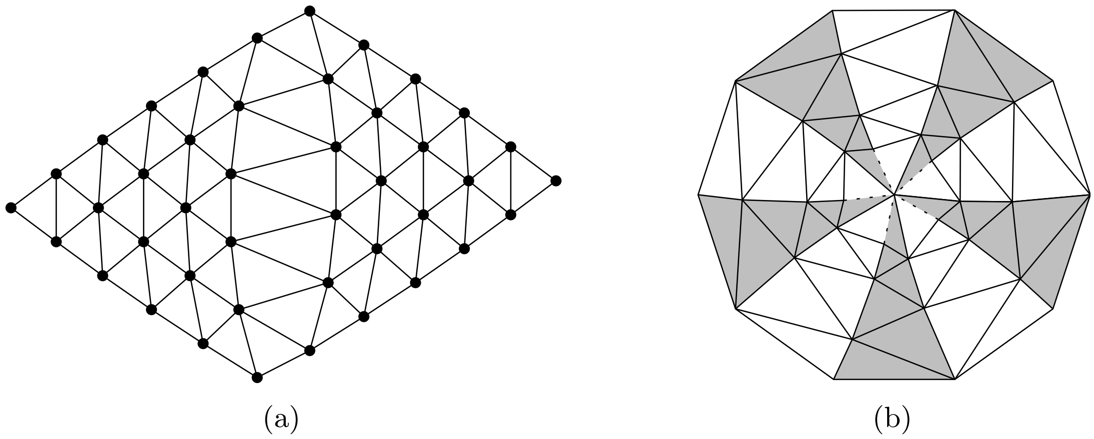

Fig. 4. Examples for the proof of Theorem 4

(a) The triangulation T in Figure 4(a) is obtained by perturbing a triangular grid so that the sides bulge. One can check by inspection, using Theorem 2, that it is a minimal pseudotriangulation, and hence also a minimum pseudotriangulation in T.

(b) In the triangulation of Figure 4(b) we can obtain a minimal triangulation with $3n-18$ edges by removing the five dotted edges in the center, or we can get another minimal triangulation with $2n-2$ edges by removing the edges in the shaded funnels.

(c) The example of Figure 3 in Theorem 3 is a minimum and minimal pseudotriangulation $P^T(S)$ with $3n-12$ edges. A minimum pseudotriangulation of S has always $2n-3$ edges. $\square$

5 Constructing a minimal pseudotriangulation in a triangulation

In the following, we shall present a linear-time greedy algorithm to construct a minimal pseudotriangulation in a given triangulation T. By Theorem 2, we just need to check whether we can remove a single edge or a triangle and keep a pseudotriangulation. If this is the case, we remove the edge or triangle and continue with the resulting pseudotriangulation. The following lemma explains how to carry out this test efficiently.

Lemma 4. (a) *Let P be a pseudotriangulation and $e \in P$ be an edge. Then $P - e$ is a pseudotriangulation if and only if the removal of e creates a new reflex vertex, in other words, if one endpoint of e is not reflex in P and reflex in $P - e$.*

(b) *Let P be a pseudotriangulation and $\{e_1, e_2, e_3\} \in P$ be a triangular face in P. Then $P - \{e_1, e_2, e_3\}$ is a pseudotriangulation if and only if the removal of the triangle makes all three vertices reflex, or more precisely, if the three vertices of $\{e_1, e_2, e_3\}$ are not reflex in P and reflex in $P - \{e_1, e_2, e_3\}$.*

Proof. Removing an edge or a triangle creates a new face from merging two or four pseudotriangles, respectively. We have to check whether this new face contains 3 convex vertices. The proof follows easily by counting the convex angles incident to the affected vertices, before and after removing the edge or the triangle. (It also follows that in case (a), only *one* endpoint of e can be a *new* reflex vertex in $P - e$.) □

Computationally, the conditions of Lemma 4 can be checked very easily. For example, let $e = ab$ be an edge in a pseudotriangulation P. Let α_1 and α_2 be the two angles incident to e at a, and let β_1 and β_2 be the two corresponding angles at b. Then $P - e$ is a pseudotriangulation if and only if $\alpha_1 < \pi$, $\alpha_2 < \pi$, and $\alpha_1 + \alpha_2 > \pi$, or if $\beta_1 < \pi$, $\beta_2 < \pi$, and $\beta_1 + \beta_2 > \pi$. The condition can be similarly formulated for the removal of a triangle (Lemma 4(b)). Thus, for a given edge or triangle, it can be checked in constant time whether it can be removed.

The algorithm for constructing a minimal pseudotriangulation now works as follows. We call an edge or a triangle *removable* if it satisfies the condition of Lemma 4(a) or (b), respectively. We start with the given triangulation. The algorithm maintains a list of all *removable edges*, which is initialized in linear time by scanning all edges. When a removable edge exists, we simply remove this edge, and update the list of removable edges. The removal of an edge $e = ab$ may affect the removability status of at most four edges of the current pseudotriangulation P (namely, the two neighboring edges at a and at b). These edges can be checked in constant time.

We repeat this procedure until the list of removable edges becomes empty. Now we check if there is any removable *triangle* according to the condition of Lemma 4(b), and we remove it. One can easily show that the removal of a triangle cannot create a new removable edge or a new removable triangle. Thus we can simply scan all faces of P sequentially, in linear time.

In the end we obtain a pseudotriangulation without removable edges or triangles, which is a minimal pseudotriangulation, by Theorem 2.

Theorem 5. *The algorithm produces a minimal pseudotriangulation P^T_{mal} of a given triangulation T in linear time.* □

6 Constructing a pseudotriangulation containing a given set of edges

In this section, we find a minimum pseudotriangulation which *contains* a given set L of non-crossing line segments. The basic idea is to maintain the set of reflex vertices of the given straight-line graph $G(S, L)$ as an invariant when we add extra edges to L to build the pseudotriangulation of L [3].

Theorem 6. *For any noncrossing set of line segments L, there is a pseudotriangulation $T'_L(S) \supseteq L$ which has the same set of reflex vertices as $G(L, S)$.*

Proof. We prove this by gradually adding edges to the set L until we get a pseudotriangulation. First we add all convex hull edges to L. This does not change the set of reflex vertices.

Then the edge set L partitions the interior of the convex hull into faces, which can be considered independently. So let us consider a single face F, see Figure 5. The boundary of F has one component B which is the exterior boundary of F, and possibly several other components inside F. Note that B is a single cycle of edges when we walk along the boundary of F inside B, although this cycle may visit the same edge twice (from two different sides) or it may visit a vertex several times. Nevertheless, we treat B as if it were a simple polygon.

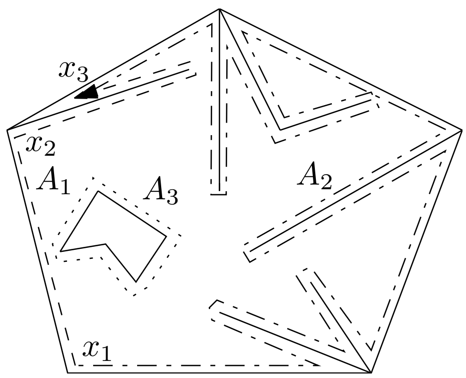

Fig. 5. Illustration for the proof of the Theorem 6.

We will subdivide F into pseudotriangles by repeatedly carrying out the following steps:

- Select a corner x_1 on B and walk clockwise along B until we find the next two corners x_2 and x_3 on B. (B must contain at least 3 corners.) We denote the path from x_1 via x_2 to x_3 along B by A_1, and the remaining part of B by A_2. By A_3, we denote the (possibly empty) set of interior components of the boundary of F, see Figure 5.
- Find the shortest path S from x_1 to x_3 in F which is homotopic to the path A_1 from x_1 to x_3. In other words, we put a string from x_1 along x_2 to x_3 and pull the string taut, regarding B and the components inside F

as obstacles. In other words, we take that shortest path from x_1 to x_3 in F which separates A_1 from $A_2 \cup A_3$.

It is clear that this path S consists of the following pieces:

(a) an initial piece following some part of B from x_1 towards x_2, turning left;
(b) a connecting line segment through the interior of F;
(c) some part of the boundary of the convex hull of $A_2 \cup A_3$, turning right;
(d) a connecting line segment through the interior of F; and
(e) a final piece following some part of B from x_2 to x_3, turning left.

Any of the pieces (a), (c), and (e) may be missing. If (c) is missing, then there is of course only one connecting segment instead of (b) and (d). It follows that the region that is cut off by this path (on the left side of S) is a pseudotriangle that contains no points inside. It may happen that S consists of a single reflex chain from x_1 to x_3 along B. In this case, F was an empty pseudotriangle, and we are done with F. Otherwise, we continue this procedure with the remaining part of F. It is also clear that no edge of S will destroy a reflex vertex. Being a geodesic path, S will only go through reflex vertices (besides the endpoints x_1 and x_3), and it will make left turns when passing around a component that is on the left side, and similarly for right turns. □

The following immediate consequence of the theorem extends the known results for $L = \emptyset$, where $r = n$.

Corollary 1. *Every minimum pseudotriangulation of a point set S with n points containing a given set L of edges with r reflex vertices has $2n - r - 2$ pseudotriangles and $3n - r - 3$ edges.* □

7 Conclusion

Several problems remain for further study.

- How to find the minimum pseudotriangulation constrained in T? Is this problem NP-hard?
- Study minimum pseudotriangulations subject to some other constraints.
- How to find the minimum-weight pseudotriangulation?

References

1. Aichholzer O., Hoffmann M., Speckmann B., and Tóth C. D., 'Degree bounds for constrained pseudo-triangulations', in preparation.
2. Kirkpatrick D., Snoeyink J., and Speckmann B., 'Kinetic collision for simple polygons', in *Proc. 16th Ann. Symp. on Computational Geometry*, 2000, pp. 322–329.
3. Pocchiola M. and Vegter G., 'Topologically sweeping visibility complexes via pseudotriangulations', *Discrete and Computational Geometry* **16** (1996), 419–453.
4. Pocchiola M. and Vegter G., 'The visibility complex', *International Journal on Computational Geometry and Applications* **6** (1996), 279–308.
5. Streinu I., 'A combinatorial approach to planar non-colliding robot arm motion planning', *Proc. 41st Ann. Symp. Found. Comput. Sci.* (FOCS), 2000, pp. 443–453.