ELSEVIER

Computational Geometry 8 (1997) 195–218

Computational Geometry
Theory and Applications

Angle-Restricted Tours in the plane

Sándor P. Fekete [a,*], Gerhard J. Woeginger [b,1]

[a] *Center for Parallel Computing, Universität zu Köln, D-50923 Köln, Germany*
[b] *Department of Mathematics, Graz University of Technology, Steyrergasse 30, A-8010 Graz, Austria*

Communicated by D. Avis; submitted 20 November 1991; resubmitted 30 March 1996; accepted 26 April 1996

Abstract

For a given set $A \subseteq (-\pi; +\pi]$ of angles, the problem "Angle-Restricted Tour" (ART) is to decide whether a set P of n points in the Euclidean plane allows a closed directed tour consisting of straight line segments, such that all angles between consecutive line segments are from the set A.

We present a variety of algorithmic and combinatorial results on this problem. In particular, we show that any finite set of at least five points allows a "pseudoconvex" tour (i.e., a tour where all angles are nonnegative), and we derive a fast algorithm for constructing such a tour. Moreover, we give a complete classification (from the computational complexity point of view) for the special cases where the tour has to be part of the orthogonal grid. © 1997 Elsevier Science B.V.

Keywords: Angle-Restrictions; Traveling Salesman Problem; Hamiltonian cycle; Convexity; Complexity; Geometry; NP-complete

1. Introduction

The question of traveling a given finite set of locations with minimum overall distance (generally known as the "Traveling Salesman Problem" or TSP) is one of the classical problems in combinatorial optimization. In recent years, there has been some increased interest in studying Hamiltonian cycles that optimize objective functions related to angles between consecutive edges in the tour. Practical motivations come from finding "smooth" tours for nonholonomic robots [3,4,19]. Also, there has been interest in motion planning with "curvature constraints" (see [2]), motivated by nonholonomic motion planning for steering-constrained robots with limited range of change of direction [5,6,15,18,20]. Other practical motivations stem from the planning of curvature-constrained paths for cars [15] and high-speed aircraft [8].

* Corresponding author.

1 Gerhard Woeginger acknowledges support by the Spezialforschungsbereich F 003 "Optimierung und Kontrolle", Projektbereich Diskrete Optimierung.

PII S0925-7721(96)00012-0

Very recently, Aggarwal et al. [1] showed that the so-called "Angular-Metric TSP" is NP-complete. In this problem, one has to find a Hamiltonian path for a given set of points in the Euclidean plane, such that the sum of the direction changes at each vertex along the tour is minimized. This result has important consequences for the problem of matroid parity, as it resolves the long-standing open question about the hardness of the weighted linear matroid parity problem in the affirmative.

In this paper, we discuss aspects of the related problem "Angle-Restricted Tour" (ART) that was first considered in the first author's Ph.D. thesis [13]: for a given set A of feasible angles, decide whether a finite set of points in the Euclidean plane allows a (directed) Hamiltonian cycle, such that all angles between consecutive edges are from the set A. Similarly, we define angle-restricted paths. We can think of this as the question of deciding whether a machine with restricted mobility is able to make a roundtrip through a given set of points. For this motivation, Culberson and Rawlins [9] have discussed the problem of finding a simple polygon with a given sequence of angles. Restricted orientations have also been examined because of their relevance for computer graphics and VLSI design—see Rawlins and Wood [23–26], Schuierer [28] and Widmayer et al. [29].

We will mainly deal with $A^{\text{obtu}} = \{\alpha \mid \alpha < -\pi/2 \text{ or } \pi/2 < \alpha\}$, the set of obtuse angles, with $A^{\text{acut}} = \{\alpha \mid -\pi/2 < \alpha < \pi/2\}$, the set of acute angles, with $A^{\text{orth}} = \{-\pi/2, +\pi/2, \pi\}$, the set of orthogonal angles, and with $A^{\text{pos}} = \{\alpha \mid 0 \leqslant \alpha \leqslant \pi\}$, the set of positive angles. Tours that only use angles from A^{acut}, A^{obtu}, A^{orth} or A^{pos} are called *acute*, *obtuse*, *orthogonal* or *pseudoconvex* tours, respectively.

The main results derived in this paper are

- For A^{pos}: every point set P allows a pseudoconvex tour, if $|P| = 3$ or $|P| \geqslant 5$; point sets that are composed of a triangle and a single point in its interior do not allow a pseudoconvex tour.
- For A^{orth}: it can be decided in polynomial time whether some point set P has a $\{-\pi/2, +\pi/2\}$-tour, whereas it is NP-complete to decide whether P allows a $\{-\pi/2, +\pi/2, \pi\}$-tour or a $\{+\pi/2, \pi\}$-tour or a $\{-\pi/2, \pi\}$-tour.
- For A^{acut} and A^{obtu}: there are arbitrarily large point sets that do *not* allow acute tours and there are arbitrarily large point sets that do *not* allow obtuse tours.

The rest of the paper is organized as follows. Section 2 states some definitions and simple observations. Section 3 gives the results for pseudoconvex tours. Section 4 presents the algorithmic results on orthogonal tours. Section 5 considers obtuse and acute tours and obtuse and acute paths. Section 6 finally lists some open problems.

2. Preliminaries

Let A be a set of feasible angles, $A \subseteq (-\pi; +\pi]$. For a set P of $n \geqslant 3$ points in the Euclidean plane, consider a directed tour $\langle p_1, p_2, \ldots, p_n, p_1 \rangle$. We call a tour *weakly feasible with respect to A*, or an *A-tour* for short, if all the angles $\angle(p_i p_{i+1} p_{i+2})$, $0 \leqslant i \leqslant n-1$, are elements of the angle set A. We call a feasible A-tour *strongly feasible*, if none of the segments $\overline{p_i p_{i+1}}$, $0 \leqslant i \leqslant n-1$, contains another point of P in its interior. This leads to the following problem:

ANGLE-RESTRICTED TOUR (ART)
Given a set $A \subseteq (-\pi; +\pi]$ of angles. The problem "Angle-Restricted Tour" (ART) is to decide whether

a set P of n points in the Euclidean plane allows an A-tour, i.e., a closed directed tour consisting of straight line segments, such that all angles between consecutive line segments are from the set A.

A (weakly) feasible spanning path with respect to some angle set A is defined like a (weakly) feasible spanning tour with the only difference that a path is not closed.

A trivial necessary condition on some angle set A to allow point sets with A-tours is the following. There exist angles $\alpha_1, \ldots, \alpha_k \in A$, k nonnegative integers $c_1, \ldots, c_k$ that are not all zero and an arbitrary integer c_{k+1}, such that

$$c_1\alpha_1 + c_2\alpha_2 + \cdots + c_k\alpha_k = c_{k+1} \cdot 2\pi$$

holds. Otherwise, a tour could never close. Clearly, this condition is not necessary for the existence of A-paths.

One interesting question is whether there is any simple connection between the size of A and the complexity of ART. It is not hard to see that ART is an easy problem when $|A| = 1$. Assume $A = \{\alpha_1\}$. Then the existence of a strongly feasible A-tour for some point set P, $|P| = n$, can be checked in the following way: the starting segment of an A-tour completely determines all the following segments in the tour. Hence, we simply check all $n - 1$ possible starting segments emanating from some fixed point in P.

The problem becomes considerably harder if $|A| = 2$. As it will be shown in Section 4, there exist angle sets with two elements for which determining the existence of A-tours is NP-complete and there exist other two-element angle sets for which determining the existence of A-tours can be done in polynomial time. These angle sets are both subsets of $A^{\text{orth}} = \{-\pi/2, +\pi/2, \pi\}$, a three-element set for which determining the existence of A-tours will also be shown to be NP-complete.

3. Pseudoconvex tours

In this section, we show that every finite set with three, five or more points has a closed tour with all angles being nonnegative, while there are configurations with four points that do not allow such a tour.

3.1. Convexity and pseudoconvexity

If we think of a tour as being a closed polygon consisting of directed line segments between consecutive vertices, we can characterize a (nondegenerate) convex tour by the following two conditions.

(C1) It is *simple*, i.e., noncrossing.

(C2) For any three consecutive vertices p_{i-1}, p_i and p_{i+1}, the angle $\angle(p_{i-1}, p_i, p_{i+1})$ lies in the nonnegative interval $[0, \pi]$.

All angles are considered to be in the interval $(-\pi, \pi]$. We follow the usual convention that for three points x, y and z, $\angle(x, y, z)$ describes the angle by which the ray $\overrightarrow{yx}$ has to be rotated in counterclockwise fashion around y to place it over $\overrightarrow{yz}$. It is not hard to see that any convex tour of a given finite set of points is a unique optimal solution for a TSP instance—note that it is the only noncrossing, i.e., 2-optimal tour. Because of this, convexity can be used as a simple geometric optimality criterion for backward error analysis—see [13,14].

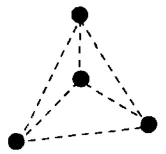

Fig. 1. The bad shape.

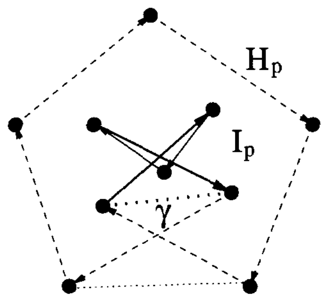

Fig. 2. Extending a pseudoconvex tour.

We can think of condition (C2) as a local condition that requires us to make a "right-hand turn" when following the line segments from p_{i-1} to p_i and then from p_i to p_{i+1}. If we relax the conditions on convexity by dropping the global condition (C1), we get so-called *pseudoconvex tours*.

Observation 3.1. *Not every finite set of points has a pseudoconvex tour—see the four points in Fig. 1. We call this arrangement the bad shape.*

Surprisingly enough, this is essentially the only counterexample.

Theorem 3.2. *Let P be a set of $n \geqslant 5$ points in the Euclidean plane. Then P has a pseudoconvex tour which can be found in* $\mathrm{O}(n \log n)$ *time.*

The objective of this section is to prove Theorem 3.2.

3.2. Setting up the proof

The basic idea for proving Theorem 3.2 can be seen in Fig. 2. Consider the convex hull of the point set P. Assuming that we have a pseudoconvex tour of the points I_P inside the hull, we will try to find an edge of this inner tour that we can replace by a sequence of all the points on the hull H_P. It turns out that a particular arrangement of directed edges in the inner tour guarantees the existence of such a replaceable edge (a so-called "γ-segment"). The proof proceeds by induction over the number of convex layers in the so-called *onion-decomposition* of P. Most of the work will be spent in identifying γ-segments.

Before we start the induction, we introduce some basic definitions, state and prove a key lemma, and finally take care of some basic cases that arise when the set of inner points is too small to allow a pseudoconvex tour. For the ease of description, we assume that the points are in general position. We will discuss the situation for nongeneral position at the very end of Section 3.

Definition 3.3. For any point set P, let H_P be the set of points on the boundary of the convex hull. We call H_P the set of *hull points*. We write $I_P := P \setminus H_P$ for the set of its *interior points*.

For two distinct points x and y, $\overline{xy}$ denotes the line that they determine, and $\overrightarrow{xy}$ is the directed edge from x to y. A point z lies *to the right of* $\overrightarrow{xy}$, if the angle $\angle(x, y, z)$ is positive (i.e., contained in the interval $(0, \pi)$), and z lies *to the left of* $\overrightarrow{xy}$, if the angle $\angle(x, y, z)$ is negative.

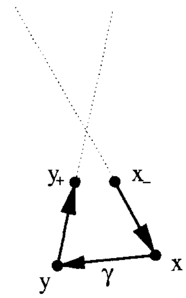

Fig. 3. A γ-segment.

Let I be a point set with a pseudoconvex tour $\mathcal{T}$. Assume there is a consecutive sequence x_-, x, y, y_+ of points in $\mathcal{T}$, such that the lines $\overline{x_-x}$ and $\overline{yy_+}$ do not intersect to the left of $\overrightarrow{xy}$. (See Fig. 3.) In this case we say that $\mathcal{T}$ contains a *γ-segment* $\overrightarrow{xy}$.

The main significance of γ-segments arises from the following Lemma 3.4, which shows that with the help of a γ-segment, we can integrate the next outer layer of points into a pseudoconvex tour of the inner layers. Lemma 3.4 will be the main tool in the induction step.

Lemma 3.4. *Let I be a finite set of points with a pseudoconvex tour $\mathcal{T}$ and a γ-segment $\overrightarrow{xy}$. Let H be a finite set that is disjoint from I such that H constitutes the convex hull of $H \cup I$. Then there exists a pseudoconvex tour for $H \cup I$ in which the elements of H follow each other consecutively, and in clockwise order around the hull.*

Proof. Assume that in $\mathcal{T}$, x_- is the predecessor of x, and y_+ is the successor of y. Consider the halfplane ε_1 lying to the right of $\overrightarrow{x_-x}$ and the halfplane ε_2 to the right of $\overrightarrow{yy_+}$. Since the bounding lines intersect the convex hull of P, both halfplanes ε_1 and ε_2 must each contain at least one point of H.

If the set $\varepsilon_1 \cup \varepsilon_2$ contains at least two points of H, then we are done—see Fig. 4. In this case, we can find two points $h_0 \in \varepsilon_1$ and $h_1 \in \varepsilon_2$ that are neighbors in the clockwise ordering $h_0, h_1, \ldots, h_{m-1}$ of H. By construction, both the angles $\angle(x_-xh_1)$ and $\angle(h_0yy_+)$ are positive, so we get the pseudoconvex tour $\langle x, h_1, \ldots, h_{m-1}, h_0, y, y_+, \mathcal{T}', x_-, x\rangle$, where $\mathcal{T}'$ denotes the part of $\mathcal{T}$ leading from y_+ to x_-. (Whenever we use $\mathcal{T}'$ in one of the figures, it will denote an appropriate part of a pseudoconvex tour.)

If the set $\varepsilon_1 \cup \varepsilon_2$ contains only a single point h_1 of H, then $h_1 \in \varepsilon_1 \cap \varepsilon_2$. (See Fig. 5.) Consider the segment $\overrightarrow{x_-x}$ instead of $\overrightarrow{xy}$. Let $x_=$ be the predecessor of x_- in $\mathcal{T}$; the halfplanes ε_3 and ε_4 lie to the right of $\overrightarrow{x_=x_-}$ and $\overrightarrow{xy}$, respectively. By definition, $\overline{x_-x}$ and $\overline{yy_+}$ do not intersect to the left of $\overrightarrow{xy}$. This implies that the set of all points to the left of $\overrightarrow{yy_+}$ and to the left of $\overrightarrow{x_-x}$ lies to the right of $\overrightarrow{xy}$. This means that ε_4 contains all the points in $H \setminus \{h_1\}$. Like in the previous case, we can find two consecutive points h_i and h_{i+1}, such that $h_i \in \varepsilon_4$ and $h_{i+1} \in \varepsilon_3$.

By construction, all four angles $\angle(x_=x_-h_{i+1})$, $\angle(x_-h_{i+1}h_{i+2})$, $\angle(h_{i-1}h_ix)$, $\angle(h_0xy)$ are positive, so in this case the sequence $\langle x_=, x_-, h_{i+1}, \ldots, h_{i-1}, h_i, x, y, y_+, \mathcal{T}', x_=, x_-\rangle$ forms a pseudoconvex tour. $\Box$

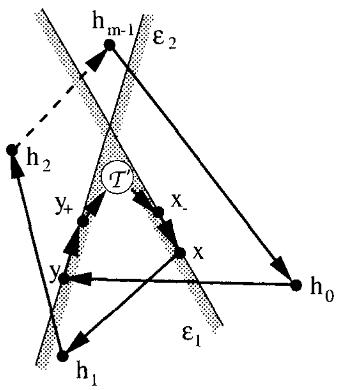

Fig. 4. Using a γ-segment when $\varepsilon_1 \cup \varepsilon_2$ contains more than one hull point.

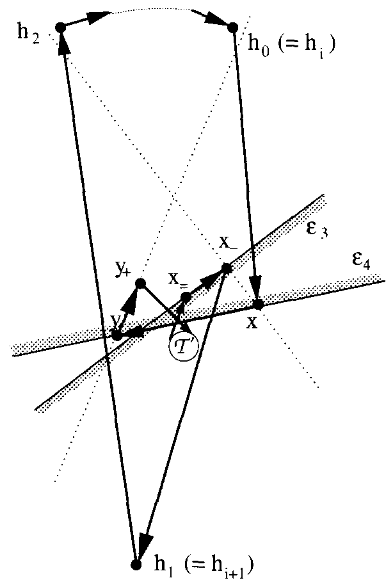

Fig. 5. Using a γ-segment when $\varepsilon_1 \cup \varepsilon_2$ contains only one hull point.

The following Lemma 3.5 will cover all the starting cases for the induction. The reader should be warned that the number of cases that need to be considered will make the proof quite tedious.

Lemma 3.5. *Let P be a set of points in the Euclidean plane that is in general position and not a bad shape. Let I_P be the set of its interior points. If $|I_P| < 5$ and if I_P does not allow a pseudoconvex tour, then P has a pseudoconvex tour containing a γ-segment.*

Proof. Let H_P with $|H_P| = m$ be the hull points of P. Let the points of H_P be ordered as $h_0, \ldots, h_{m-1}$ when going clockwise around the hull. We distinguish four cases and three subcases according to the form of I_P.

Case 0. $|I_P| = 0$.

For $m = 3, 4$, the tour $\langle h_0, \ldots, h_{m-1}, h_0\rangle$ is convex, thus pseudoconvex. Since there must be two consecutive angles whose sum does not exceed π, the tour contains a γ-segment. For $m \geqslant 5$, we can rearrange the order of h_{m-1}, h_0, h_1, h_2, h_3 in the convex tour $\langle h_0, \ldots, h_{m-1}, h_0\rangle$ in order to get a pseudoconvex tour with a γ-segment, as shown in Fig. 6. Clearly, the tour $\langle h_{m-1}, h_1, h_3, h_0, h_2, h_4, \ldots, h_{m-1}\rangle$ is pseudoconvex and $\overrightarrow{h_3h_0}$ is a γ-segment.

Case 1. $|I_P| = 1$.

See Fig. 7. Since P is not a bad shape, we have $|H_P| \geqslant 4$. Let x denote the only point in I_P. There exists a triangle $\triangle h_j h_k h_l$ of points $h_j, h_k, h_l \in H_P$ that contains x in its interior. Without loss of generality we assume that $0 < j < k < l$ and that x is contained in the triangle $\triangle h_0 h_j h_k$. Then $\langle h_0, x, h_l, \ldots, h_{m-1}, h_1, \ldots, h_{l-1}, h_0\rangle$ is a pseudoconvex tour: $\angle(h_{l-1}, h_0 x)$, $\angle(h_0 x h_1)$, $\angle(x h_l h_{l+1})$ are positive by virtue of $x \in \triangle h_j h_k h_l \cap \triangle h_0 h_j h_k$; all other angles are positive by running around the convex hull in the right orientation. Furthermore, $\overrightarrow{x h_l}$ is a γ-segment in the tour.

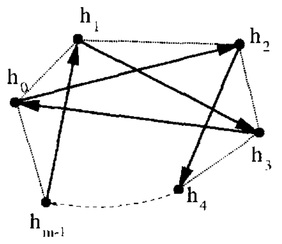

Fig. 6. How to solve Case 0.

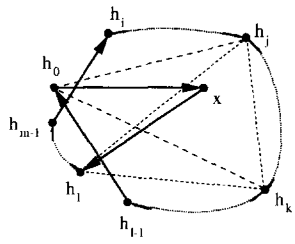

Fig. 7. How to solve Case 1.

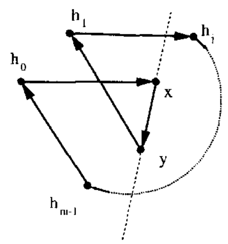

Fig. 8. How to solve Case 2.

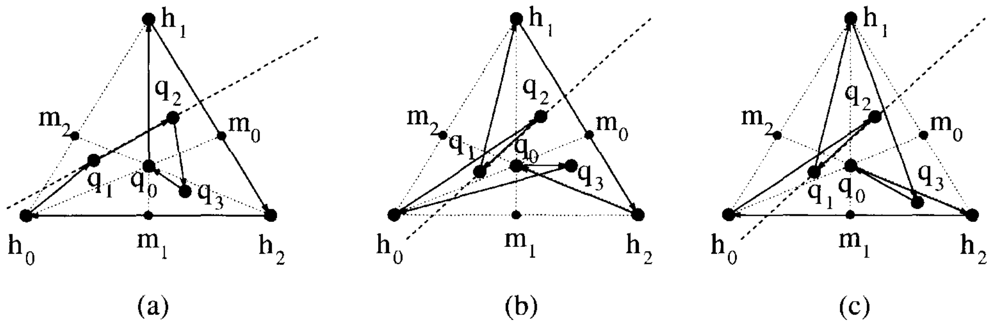

Fig. 9. How to solve Subcase B3.

Case 2. $|I_P| = 2$.

See Fig. 8. Let x and y denote the two points in I. The line through x and y partitions the set H_P into two non-empty subsets. Without loss of generality assume that h_0 and h_1 lie to the right of $\overrightarrow{xy}$. Then $\langle h_0, x, y, h_1, h_2, \ldots, h_{m-1}, h_0\rangle$ is a pseudoconvex tour and $\overrightarrow{xy}$ is a γ-segment in it.

Case B. I_P is a bad shape.

Let $I_P := \{q_0, q_1, q_2, q_3\}$, where q_0 is the point contained in the convex hull of the other three points which have the clockwise order $q_1q_2q_3$. According to the number m of hull points, we branch into three subcases $m = 3$, $m = 4$ and $m \geqslant 5$.

Subcase B3. $m = 3$.

See Fig. 9. The lines $\overline{h_0q_0}$, $\overline{h_1q_0}$ and $\overline{h_2q_0}$ partition the triangle $\triangle h_0h_1h_2$ into the six triangles

$$\triangle_1 := \triangle h_0q_0m_2, \qquad \triangle_2 := \triangle h_1m_2q_0, \qquad \triangle_3 := \triangle h_1q_0m_0,$$

$$\triangle_4 := \triangle h_2m_0q_0, \qquad \triangle_5 := \triangle h_2q_0m_1, \qquad \triangle_6 := \triangle h_0m_1q_0,$$

where m_0, m_1, m_2 are the intersection points of $\overline{h_0q_0}$ and $\overline{h_1h_2}$, $\overline{h_1q_0}$ and $\overline{h_0h_2}$, $\overline{h_2q_0}$ and $\overline{h_0h_1}$, respectively.

Since q_0 is contained in the convex hull of q_1, q_2 and q_3, there must be one of these three points such that it lies in triangle $\triangle_i$ and none of the two other points lie in $\triangle_{i-1}$, $\triangle_i$ or $\triangle_{i+1}$. By symmetry, we may assume that q_1 lies in $\triangle_1$ and neither q_2 nor q_3 lie in $\triangle_6$, $\triangle_1$ or $\triangle_2$. Furthermore, we conclude that ($q_2 \in \triangle_3$ or $q_2 \in \triangle_4$) and ($q_3 \in \triangle_4$ or $q_3 \in \triangle_5$).

Now consider the line $\overline{q_1q_2}$, which has to intersect $\overrightarrow{h_1h_2}$ and either $\overrightarrow{h_0h_1}$ or $\overrightarrow{h_0h_2}$. If $\overline{q_1q_2}$ intersects $\overrightarrow{h_0h_1}$, then the angle $\angle(h_0q_1q_2)$ is positive and one easily checks that all the angles in the tour $\langle h_0, q_1, q_2, q_3, q_0, h_1, h_2, h_0\rangle$ are positive—see Fig. 9(a): $\angle(q_1q_2q_3)$, since q_1, q_2, q_3 follow each other clockwise on the convex hull of I_P; $\angle(q_2q_3q_0)$, since q_2, q_3 follow each other clockwise on the convex hull of I_P and q_0 lies inside the convex hull of I_P; $\angle(q_3q_0h_1)$, since q_3 lies in $\triangle_3 \cup \triangle_4 \cup \triangle_5$; $\angle(q_0h_1h_2)$, since h_1, h_2 follow each other clockwise on the convex hull of H_P and q_0 lies inside the convex hull of H_P; $\angle(h_1h_2h_0)$, since h_1, h_2, h_0 follow each other clockwise on the convex hull of H_P; $\angle(h_2h_0q_1)$, since h_2, h_0 follow each other clockwise on the convex hull of H_P and q_1 lies inside the convex hull of H_P. Clearly, $\overrightarrow{h_2h_0}$ is a γ-segment.

If $\overline{q_1q_2}$ intersects $\overrightarrow{h_0h_2}$, then the angles $\angle(h_0q_2q_1)$ and $\angle(q_2q_1h_1)$ are positive and q_2 lies in $\triangle_3$. As stated above, either $q_3 \in \triangle_4$ or $q_3 \in \triangle_5$.

For $q_3 \in \triangle_4$, the angle $\angle(h_2q_0q_3)$ is positive and one easily checks that all the angles in the tour $\langle h_0, q_2, q_1, h_1, h_2, q_0, q_3, h_0\rangle$ are positive—see Fig. 9(b): $\angle(q_1h_1h_2)$, since h_1, h_2 follow each other clockwise on the convex hull of H_P and q_1 lies inside the convex hull of H_P; $\angle(h_1h_2q_0)$, since h_1, h_2 follow each other clockwise on the convex hull of H_P and q_0 lies inside the convex hull of H_P; $\angle(q_0q_3h_0)$, since q_3 lies in $\triangle_4$; $\angle(q_3h_0q_1)$, since q_2 lies in $\triangle_3$. Clearly, $\overrightarrow{q_2q_1}$ is a γ-segment.

For $q_3 \in \triangle_5$, the angle $\angle(h_2q_3q_0)$ is positive and one easily checks that all the angles in the tour $\langle h_0, q_2, q_1, h_1, q_3, q_0, h_2, h_0\rangle$ are positive—see Fig. 9(c): $\angle(q_1h_1q_3)$, $\angle(h_1q_3q_0)$, $\angle(q_3q_0h_2)$, since q_3 lies in $\triangle_5$; $\angle(q_0h_2h_0)$, since h_2, h_0 follow each other clockwise on the convex hull of H_P and q_0 lies inside the convex hull of H_P; $\angle(h_2h_0q_2)$, since h_2, h_0 follow each other clockwise on the convex hull of H_P and q_2 lies inside the convex hull of H_P. Again, $\overrightarrow{q_2q_1}$ is a γ-segment.

Subcase B4. $m = 4$.

Let s be the intersection point of the diagonals $\overline{h_0h_2}$ and $\overline{h_1h_3}$. Without loss of generality, let q_0 be contained in the triangle $\triangle$ with vertices s, h_3, h_2. Furthermore, we may assume that we have numbered q_1, q_2 and q_3 in a clockwise fashion such that q_1 lies to the left of $\overrightarrow{h_0q_0}$ and q_2 to the right of $\overrightarrow{h_0q_0}$. Both angles $\angle(h_0h_2h_3)$ and $\angle(h_2h_3h_1)$ are positive because of the clockwise order of the h_i.

If q_3 lies to the right of $\overrightarrow{h_1q_0}$, the angle $\angle(h_1q_0q_3)$ is positive. It is not hard to check that all other angles in the tour $\langle h_0, h_2, h_3, h_1, q_0, q_3, q_1, q_2, h_0\rangle$ are positive—see Fig. 10(a): $\angle(h_3h_1q_0)$, since q_0 lies in $\triangle$; $\angle(q_0q_3q_1)$, since q_3, q_1 follow each other clockwise on the convex hull of I_P and q_0 lies inside the convex hull of I_P; $\angle(q_3q_1q_2)$, by the way we numbered the q_i; $\angle(q_1q_2h_0)$, since $\angle(q_2h_0q_1)$ is positive by numbering of the q_i; $\angle(q_2h_0h_2)$, since q_0 lies in $\triangle$ and q_2 to the right of the line $\overrightarrow{h_0q_0}$.

If q_3 lies to the left of $\overrightarrow{h_1q_0}$, it follows that q_3 must lie to the left of $\overrightarrow{h_0q_0}$, otherwise $\langle q_1, q_2, q_3, q_1\rangle$ cannot run clockwise around q_0. By the same argument, q_2 must lie to the right of $\overrightarrow{h_1q_0}$, and q_1 to the left of $\overrightarrow{h_1q_0}$. We easily check that all other angles in the tour $\langle h_0, h_2, h_3, h_1, q_1, q_2, q_3, q_0, h_0\rangle$ are

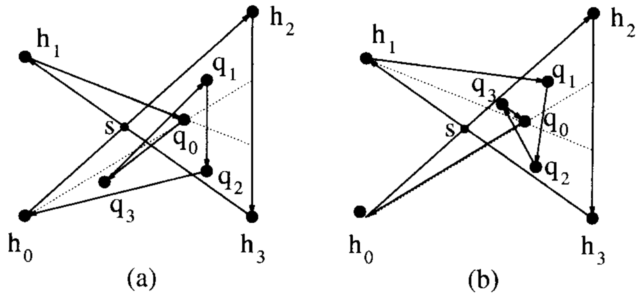

Fig. 10. How to solve Subcase B4.

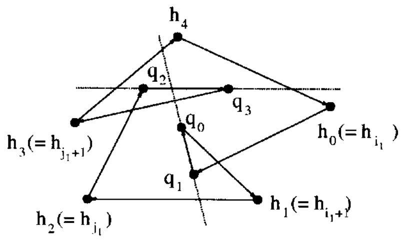

Fig. 11. How to solve Subcase B5+.

positive—see Fig. 10(b): $\angle(h_3h_1q_1)$, since q_0 lies in $\triangle$ and q_1 to the right of $\overrightarrow{h_1q_0}$; $\angle(h_1q_1q_2)$, since $\angle(q_2h_1q_1)$ is positive as q_1 and q_2 lie to the left and to the right of $\overrightarrow{h_1q_0}$; $\angle(q_1q_2q_3)$, by the way we numbered the q_i; $\angle(q_2q_3q_0)$, since q_2, q_3 follow each other clockwise on the convex hull of I_P and q_0 lies inside the convex hull of I_P; $\angle(q_3q_0h_0)$, since q_3 lies to the left of $\overrightarrow{h_0q_0}$; $\angle(q_0h_0h_2)$, since q_0 lies in $\triangle$.

In both cases, $\overrightarrow{h_2h_3}$ is a γ-segment.

Subcase B5+. $m \geqslant 5$.

See Fig. 11. Consider the lines $l_1 := \overline{q_0q_1}$ and $l_2 := \overline{q_2q_3}$. The idea is to find two appropriate pairs of points $\langle h_{i_1}, h_{i_1+1}\rangle$ and $\langle h_{j_1}, h_{j_1+1}\rangle$ for which we can insert $\langle q_0, q_1\rangle$ and $\langle q_2, q_3\rangle$ into the convex tour $h_{i_1}, \ldots, h_{j_1}, \ldots, h_{i_1}$. To get a pseudoconvex tour, we have to make sure that these two pairs are disjoint. For easier notation, we write e_i for a pair $\{h_i, h_{i+1}\}$ of consecutive points of the hull H_P.

There are at least three different pairs e_{i_1}, e_{i_2} and e_{i_3} that are not separated by l_1; similarly, there are at least three pairs e_{j_1}, e_{j_2} and e_{j_3} that are not separated by l_2. Assume that $e_{i_2} \cap e_{j_k} \neq \emptyset$ for $k = 1, 2, 3$, i.e., the edge e_{i_2} intersects all three edges e_{j_1}, e_{j_2}, e_{j_3}. It is not hard to see that this can only happen if e_{j_1}, e_{j_2}, e_{j_3} are adjacent edges and if (for an appropriate numbering), $e_{i_2} = e_{j_2}$. It follows that there must be a j_k for which $e_{i_1} \cap e_{j_k} = \emptyset$; let this be the case for j_1.

Without loss of generality, let the angles $\angle(h_{i_1}q_0q_1)$ and $\angle(h_{j_1}q_2q_3)$ be positive. (Otherwise exchange the order of q_0 and q_1 or q_2 and q_3, respectively, in the following tour to make it pseudoconvex.) Then we can change the convex tour of H_P into $\langle h_{i_1}, q_0, q_1, h_{i_1+1}, \ldots, h_{j_1}, q_2, q_3, h_{j+1}, \ldots, h_{i_1}\rangle$. By choice of i_1 and j_1, the angles $\angle(q_0q_1h_{i_1+1})$ and $\angle(q_2q_3h_{j_1+1})$ are both positive and the tour is pseudoconvex. By construction, $\overrightarrow{q_0q_1}$ is a γ-segment.

This finally concludes the proof of Lemma 3.5. □

3.3. Proving Theorem 3.2

With the help of Lemma 3.4 and Lemma 3.5, we can prove the main Theorem 3.2. As stated before, we will use induction on the number of convex layers. Lemma 3.5 is used to start the induction, while Lemma 3.4 is the main tool for establishing the induction step.

The main idea is to include the points of one layer at a time in their clockwise order into a tour of the inner points. (See Fig. 12.) Since we use γ-segments for this purpose, we may have to rearrange the resulting tour in order to get a γ-segment that we can use for including the *next* layer. Since there are numerous cases to consider, we will use indirect argumentation to establish the validity of the induction step.

Let H_P with $|H_P| = m$ be the extremal points of P. Let the points of H_P be ordered as $h_0, \ldots, h_{m-1}$ when going clockwise around the hull H_P and let $I_P = P \setminus H_P$ denote the interior points of P. Let H_I with $|H_I| = k$ be the hull points of I_P and I_I the interior points of I_P.

If I_P satisfies the assumptions of Lemma 3.5 (i.e., $|I_I| < 5$ and I_I does not admit a pseudoconvex tour), we are done: I_P allows a pseudoconvex tour with a γ-segment, and so by Lemma 3.4, this tour can be extended to a pseudoconvex tour of P. So we only have to consider point sets P for which this is not the case. Assume that $\overline{P}$ is the smallest of these sets that does not have *a pseudoconvex tour where the points of $H_{\overline{P}}$ follow each other in the order in which they appear around the hull.*

We observe that the set $I_{\overline{P}}$ with its interior points $I_{I_{\overline{P}}}$ cannot satisfy the assumptions of Lemma 3.5: one of the constructions in Lemma 3.5 would create a pseudoconvex tour of $I_{\overline{P}}$ with a γ-segment, and therefore (by Lemma 3.4) we would get a pseudoconvex tour for $\overline{P}$ in which the elements of $H_{\overline{P}}$ follow each other in order. Using the minimality assumption on $\overline{P}$, we conclude that $I_{\overline{P}}$ allows a pseudoconvex tour $\mathcal{T}$, where the elements of $H_{I_{\overline{P}}}$ follow each other in the ordering $q_1, \ldots, q_k$.

In the following, we will either locate a γ-segment in $\mathcal{T}$, find a way to rearrange $\mathcal{T}$ in order to create a γ-segment, or show that $\mathcal{T}$ has a structure that allows us to extend $\mathcal{T}$ into a tour of $\overline{P}$ without using γ-segments.

We distinguish three cases on the number k of points in $H_{I_{\overline{P}}}$:

Case 1. $k \leqslant 4$.

Since the sum of angles in the polygon with vertices $q_1, \ldots, q_k$ is π or 2π, it is not hard to see that one of the segments $\overrightarrow{q_1q_2}$, $\overrightarrow{q_2q_3}$, $\overrightarrow{q_{k-1}q_k}$ is a γ-segment.

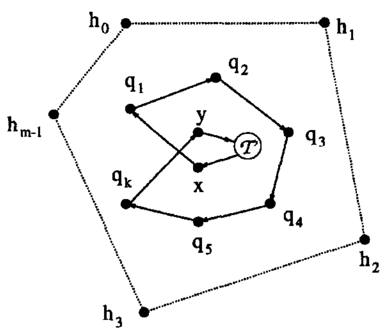

Fig. 12. The setup for the induction step.

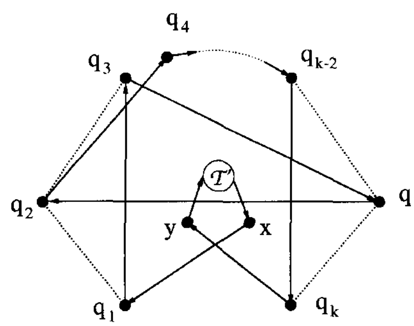

Fig. 13. Rearranging $\mathcal{T}$ in order to get a γ-segment.

Case 2. $k \geqslant 6$.

Let x be the predecessor of q_1 in $\mathcal{T}$ and y be the successor of q_k in $\mathcal{T}$. If x lies to the left of $\overrightarrow{q_1q_3}$, the segment $\overrightarrow{q_1q_2}$ is a γ-segment, contradicting our assumption on $\overline{P}$. Similarly, if y lies to the left of $\overrightarrow{q_{k-2}q_k}$, the segment $\overrightarrow{q_{k-1}q_k}$ is a γ-segment. If neither is the case, we can rearrange the sequence of the q_i in $\mathcal{T}$ to get

$$\langle x, q_1, q_3, q_{k-1}, q_2, q_4, \dots, q_{k-2}, q_k, y, \mathcal{T}, x\rangle,$$

thereby creating another pseudoconvex tour with $\overrightarrow{q_3q_5}$ being a γ-segment. (See Fig. 13.)

Case 3. $k = 5$.

Let x be the predecessor of q_1 in $\mathcal{T}$ and y be the successor of q_5. Let x_- be the predecessor of x and y_+ the successor of y. Now look at Fig. 14 and read the following five properties (P1)–(P5) from it.

(P1) x lies to the right of $\overrightarrow{q_1q_3}$.

In case x would lie to the left of $\overrightarrow{q_1q_3}$, then the segment $\overrightarrow{q_1q_2}$ would form a γ-segment in $\mathcal{T}$.

(P2) $\overline{q_1q_2}$ and $\overline{q_3q_4}$ intersect in a point s_{23} that lies to the left of $\overrightarrow{q_2q_3}$. q_2 lies between q_1 and s_{23}, and q_3 between s_{23} and q_4.

Otherwise, if the lines $\overline{q_1q_2}$ and $\overline{q_3q_4}$ would not intersect to the left of $\overrightarrow{q_2q_3}$, then $\overrightarrow{q_2q_3}$ would yield a γ-segment in $\mathcal{T}$.

(P3) $\overline{q_2q_3}$ and $\overline{q_4q_5}$ intersect in a point s_{34} that lies to the left of $\overrightarrow{q_3q_4}$. q_3 lies between q_2 and s_{34}, and q_4 between s_{23} and q_5.

Otherwise, $\overrightarrow{q_3q_4}$ is a γ-segment in $\mathcal{T}$.

(P4) $\overline{q_3q_4}$ and $\overline{q_5y}$ intersect in a point s_{45} that lies to the left of $\overrightarrow{q_4q_5}$. q_4 lies between q_3 and s_{45}, and q_5 between s_{45} and y.

Otherwise, $\overrightarrow{q_4q_5}$ is a γ-segment in $\mathcal{T}$.

(P5) The line $\overline{yy_+}$ does not separate q_4 and q_5.

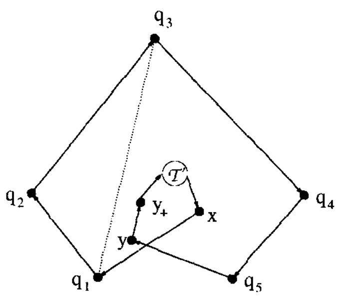

Fig. 14. Potential γ-segments in $\mathcal{T}$.

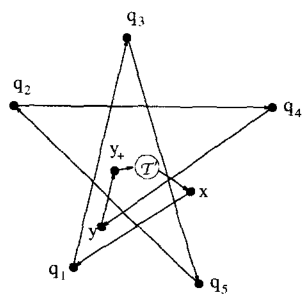

Fig. 15. A pseudoconvex tour for y to the right of $\overrightarrow{q_2q_4}$.

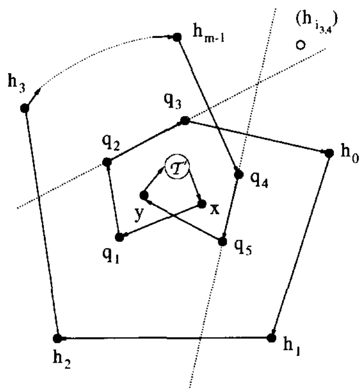

Fig. 16. There must be a point $h_{i_{34}}$ to the left of $\overrightarrow{q_2q_3}$ and to the right of $\overrightarrow{q_4q_5}$.

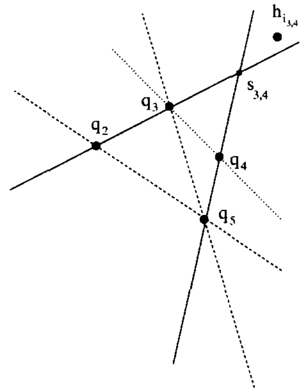

Fig. 17. $h_{i_{34}}$ lies to the left of $\overrightarrow{q_2q_4}$ and to the left of $\overrightarrow{q_3q_5}$.

Otherwise, $\overrightarrow{q_5y}$ is a γ-segment in $\mathcal{T}$. From (P1) and (P5), it now follows that the tour

$$\langle x_-, x, q_1, q_3, q_5, q_2, q_4, y, y_+, \ldots, x_- \rangle$$

is pseudoconvex with $\overrightarrow{q_3q_5}$ being a γ-segment, in case that y lies to the right of $\overrightarrow{q_2q_4}$. (See Fig. 15.) Therefore, we assume that

(P6) y lies to the left of $\overrightarrow{q_2q_4}$. Let t_1 be the intersection of $\overline{q_2q_4}$ and $\overline{q_5y}$. Clearly, t_1 lies between q_2 and q_4.

If none of the points $h_i \in H$ lies to the left of $\overrightarrow{q_2q_3}$ and to the left of $\overrightarrow{q_4q_5}$, the edge $\overrightarrow{q_3q_4}$ can be used to extend $\mathcal{T}$ to a pseudoconvex tour of $\overline{P}$ by inserting the points h_i in the way they appear clockwise around H: assuming that h_0 is the first of the h_i that lies to the left of $\overrightarrow{q_4q_5}$, h_0 lies to the right of $\overrightarrow{q_2q_3}$ and h_{m-1} lies to the right of $\overrightarrow{q_4q_5}$.

This means that the tour $\langle x, q_1, q_2, q_3, h_0, \ldots, h_{m-1}, q_4, q_5, y, y_+, \ldots, x_- \rangle$ is pseudoconvex, thus contradicting our assumption on $\overline{P}$. (See Fig. 16.) We conclude that

(P7) there is a point $h_{i_{34}}$ to the left of $\overrightarrow{q_2q_3}$ and to the left of $\overrightarrow{q_4q_5}$; without loss of generality, we assume that $h_{i_{34}+1}$ lies to the right of $\overrightarrow{q_2q_3}$.

Considering the position of $h_{i_{34}}$ relative to the points q_2, q_3, q_4, q_5 and s_{34} (see Fig. 17), we see that $h_{i_{34}}$ lies to the left of $\overrightarrow{q_2q_4}$ and to the left of $\overrightarrow{q_3q_5}$.

Considering the position of $h_{i_{45}}$ relative to q_3, q_4, q_5, t_1 and s_{45}, we see that we have a similar situation as for $h_{i_{34}}$ relative to q_2, q_3, s_{34}, q_4 and q_5—see Fig. 18. Thus,

(P8) $h_{i_{45}}$ lies to the right of $\overrightarrow{t_1q_4}$, i.e., to the right of $\overrightarrow{q_2q_4}$. This implies $h_{i_{34}} \neq h_{i_{45}}$.

Next consider the intersection of $\overline{q_5y}$ with the boundary of the triangle $\triangle = (q_4s_{23}q_2)$, as shown in Fig. 19. Since $\{s_{45}\} = \overline{q_3q_4} \cap \overline{q_5y}$ lies to the right of $\overrightarrow{q_4q_5}$, the point s_{45} cannot belong to the triangle; so $\overline{q_5y}$ must intersect $\overline{q_2s_3}$ in a point t_2 between q_2 and s_3. This means that

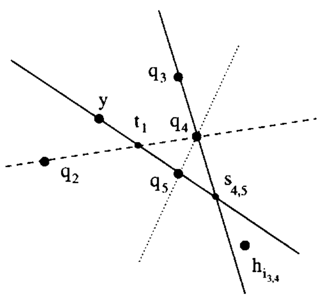

Fig. 18. $h_{i_{45}}$ lies to the right of $\overrightarrow{q_2q_4}$.

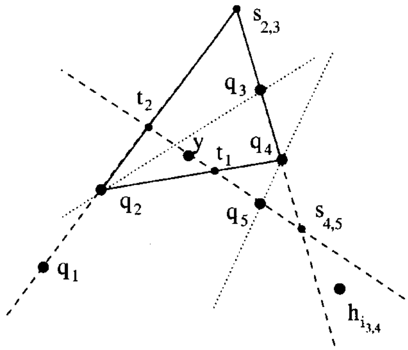

Fig. 19. $h_{i_{45}}$ lies to the right of $\overrightarrow{q_1q_2}$.

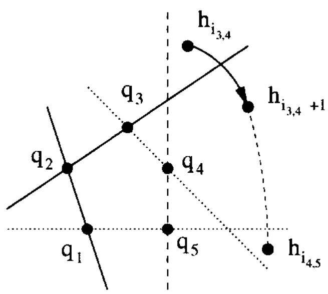

Fig. 20. $h_{i_{34}+1}$ lies to the right of $\overrightarrow{q_1q_2}$.

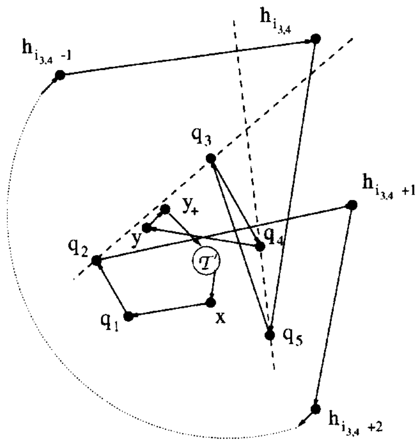

Fig. 21. The final pseudoconvex tour.

(P9) $h_{i_{45}}$ lies to the right of $\overrightarrow{t_2s_3}$, i.e., to the right of $\overrightarrow{q_1q_2}$.

Finally, $h_{i_{34}+1}$ must lie to the right of $\overrightarrow{q_1q_2}$ (see Fig. 20): by construction, $h_{i_{34}+1}$ must lie between $h_{i_{34}}$ and $h_{i_{45}}$ when going clockwise around the hull H; if we do cross $\overline{q_1q_2}$ on this way, this cannot happen after we cross $\overline{q_2q_3}$. This means that

(P10) $h_{i_{34}+1}$ lies on the same side of $\overrightarrow{q_1q_2}$ as $h_{i_{45}}$, i.e., on the right.

But now all the angles in the tour

$$\langle x, q_1, q_2, h_{i_{34}+1}, \ldots, h_{i_{34}}, q_5, q_3, q_4, y, y_+ \ldots, x_-, x\rangle$$

are positive—see Fig. 21.

The angle $\angle(q_1q_2h_{i_{34}+1})$, as stated; $\angle(q_2h_{i_{34}+1}h_{i_{34}+2})$, since $h_{i_{34}+1}$, $h_{i_{34}+2}$ follow each other clockwise on the convex hull of $\overline{P}$ and q_2 lies inside the convex hull of $\overline{P}$; $\angle(h_{i_{34}-1}h_{i_{34}}q_5)$, since $h_{i_{34}-1}$, $h_{i_{34}}$ follow each other clockwise on the convex hull of $\overline{P}$ and q_5 lies inside the convex hull of $\overline{P}$; $\angle(h_{i_{34}}q_5q_3)$, as stated; $\angle(q_5q_3q_4)$, since $\angle(q_3q_4q_5)$ is positive; $\angle(q_3q_4y)$, since q_3, q_4 follow each other

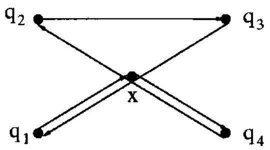

Fig. 22. A pseudoconvex tour that encounters a vertex more than once.

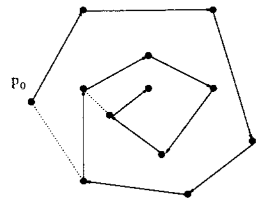

Fig. 23. Any point set has a pseudoconvex spanning path.

clockwise on the convex hull of $I_{\overline{P}}$ and y lies inside the convex hull of $I_{\overline{P}}$; $\angle(q_4yy_+)$, since $\angle(q_5yy_+)$ is positive and $\overline{yy_+}$ does not separate q_4 and q_5.

We conclude that $\overline{P}$ *does* indeed have a pseudoconvex tour with the elements of the hull $H_{\overline{P}}$ following each other in order. This completes the proof of the combinatorial part of Theorem 3.2.

The time complexity $O(n \log n)$ as claimed in the algorithmic part of Theorem 3.2 follows by constructing the so-called onion decomposition of a planar point set into convex layers in $O(n \log n)$ time (see [7,12]). Once we have this decomposition, we can construct the pseudoconvex tour in amortized time $O(n)$: starting from the inside, at each of the $O(n)$ stages, we include another layer L_i of size l_i of the decomposition into the tour. At each stage, we include a set of roughly l_i edges along the new convex hull, plus $O(1)$ many others by rerouting edges as described above. Thus, updating γ-elements (or the constellations used at the end of the preceding proof) can be done in time $\sum_i (O(l_i) + O(1)) = O(n)$, resulting in the above overall complexity. Thus the proof of Theorem 3.2 is complete. □

If the points in P are not in general position, there may be no pseudoconvex tour that encounters each of the vertices exactly once. Easy examples for this situation arise if P is a set of collinear points or if P consists of the four vertices of a rectangle together with a fifth point at the intersection of the diagonals.

We may choose to permit that a vertex may be "run over" by other edges—see Fig. 22, where we have the *weakly feasible* pseudoconvex tour $\langle x, q_4, q_2, q_3, q_1, x\rangle$. If we content ourselves with weakly feasible tours, it is not hard to check that the above steps of the proof remain valid even if the points are not in general position (since all regions used in our arguments include their boundaries, if we allow the use of the angles 0 and π.) We finally note the following.

Corollary 3.6. *Let P be a set in the plane in general position with $|P| \geqslant 5$. Then P has a strictly feasible pseudoconvex tour such all angles lie in the open interval $(0, \pi)$.*

3.4. Pseudoconvex paths

Any nondegenerate convex tour is a pseudoconvex tour that is simple. While the existence of a convex tour is a very special property of a point set, we have seen that any set of at least 5 points has a pseudoconvex tour. If we relax the question for a *Hamiltonian cycle* with nonnegative angles to the question for a *Hamiltonian path* with nonnegative angles, the situation becomes considerably easier as the simple proof for the following Theorem 3.7 shows.

Theorem 3.7. *Any set P of n points allows a non-intersecting pseudoconvex path.*

Proof. See Fig. 23. Choose any point p_0 on the convex hull of P to be the starting point and remove p_0 from P. Find some line through p_0 that does not intersect the convex hull of P and turn it clockwise until it hits the first point in P. This point is the next point in our path, we remove it from P and repeat the procedure. The resulting path will always be intersection-free. (It is not hard to see that we follow the convex layers of P in an inward spiral.) □

4. Orthogonal tours

In this section we investigate the computational complexity of detecting strongly feasible *orthogonal* tours. Let $A_1 = \{-\pi/2, +\pi/2, \pi\}$, $A_2 = \{\pi/2, \pi\}$ and $A_3 = \{-\pi/2, +\pi/2\}$. We will prove that detecting strongly feasible A_1- and A_2-tours is NP-complete, whereas finding A_3-tours can be done in polynomial time.

We start with the NP-completeness result on A_1-tours. A related result was derived by Rappaport [21,22] who gave an NP-completeness proof for the case of A_1-tours that *must not intersect themselves.*

Theorem 4.1. *Given a set P of n points in the plane, deciding whether P allows a strongly feasible A_1-tour is NP-complete.*

Proof. We show that the NP-complete problem HAMILTONIAN CYCLE IN GRID GRAPHS (Itai et al. [17]) can be reduced to detecting A_1-tours. A grid graph $G = (V, E)$ consists of a set V of n integer grid points. There is an edge between two vertices if and only if the corresponding points are at distance 1.

In our reduction, we first partition the points in V into horizontal classes. Two points belong to the same horizontal class if they have identical y-coordinates and if they are connected by a path that uses only edges of the graph with this y-coordinate. The points in the ith horizontal class are shifted by the vector $(0, 1/4^i)$. (The reason for doing this is to give all horizontal classes distinct y-coordinates.) In an analogous way, we define vertical classes and shift them by vectors $(1/4^i, 0)$. Clearly, all points adjacent in the grid graph maintain their horizontal (vertical) connections. All points not adjacent in the grid graph either keep the distinctness of their vertical (horizontal) coordinates, or these coordinates are made distinct by the shifts.

Next, we consider the leftmost point p_0 with smallest y-coordinate. Let p_1 be its horizontal neighbor. As shown in Fig. 24, we add the four points q_1, q_2, q_3, q_4 to the point set, such that q_1 has only neighbors p_0 and q_2, q_2 has only neighbors q_1 and q_3, q_3 has only neighbors q_2 and q_4, q_4 has only neighbors q_3 and p_1, and such that q_2 becomes the new lowermost and leftmost point. Clearly, there can be no orthogonal tour of the point set where q_2 is not adjacent to both q_1 and q_3, so this arrangement forces every strongly feasible A_1-tour to be axes-parallel. Now it is easy to see that a Hamiltonian cycle in the original grid graph exactly corresponds to a strongly feasible A_1-tour in the shifted and extended point set. □

Note that an alternative solution to adding the four points q_1, q_2, q_3, q_4 is provided by Theorem 4.7: if the grid graph is connected, the only orientation for which we can connect *all* vertices in *any* manner is axes-parallel.

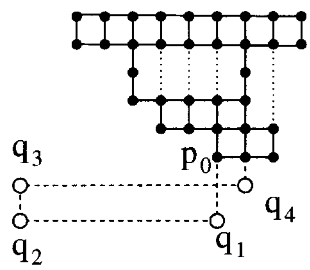

Fig. 24. Shifting vertical classes removes the dotted "crossovers".

Next, we describe the NP-completeness proof for detecting A_2-tours.

Theorem 4.2. *Given a set P of n points in the plane, deciding whether P allows a strongly feasible A_2-tour is NP-complete.*

Proof. We will show that the problem of detecting an *axes-parallel* A_2-tour is NP-complete. The claim then follows by adding some extra points as in the proof of the above theorem. We show the NP-completeness by reducing Hamiltonian cycle in cubic directed graphs to it. (See Garey and Johnson [16].) So let $D = (V, E)$ be some cubic, directed graph with $|V| = n$. We may assume that all vertices have either indegree or outdegree two, which partitions V into *in-vertices* and *out-vertices*. This partition induces a bipartition of the graph and the edges are either *mandatory* (being the only edge leaving or entering for both its end points) or *optional* (being one of two edges leaving or entering for both end points.) Furthermore, the optional edges form a set of disjoint (undirected) cycles in D. Let $C_1, \ldots, C_z$ denote the set of these cycles.

We will construct a point set P_D that allows an axes-parallel A_2-tour if and only if G has a Hamiltonian cycle.

In a first step, we choose n disjoint parallel boxes in the plane as shown in Fig. 28. Let v_1 be some out-vertex in C_1 and $c_1 = |C_1|$. Let $v_2, \ldots, v_{c_1}$ be the other vertices of C_1 in an order in which they appear when running through C_1. Then assign the first c_1 boxes to $v_1, \ldots, v_{c_1}$. In the same manner, assign the other boxes to the vertices in $C_2, \ldots, C_z$.

Next, the optional edges are represented by horizontal edges connecting the appropriate boxes, such that all edges get different vertical coordinates. We start with C_1 as follows. Clearly, the boxes for v_1 and v_{c_1} will have their two edges lying on the same side ("type A"), while any other box corresponding to C_1 will have one edge on each side ("type B"). For a box of type B, we will refer to the appropriate adjacent optional edges as the "right" and as the "left" optional edge. After placing the optional edges adjacent to v_1 at any vertical level, we run through the boxes for the vertices $v_2, \ldots, v_{c_1 - 1}$ of type B to place the other optional edges of C_1. The right edge for v_i is placed *below* the left edge whenever i is even, i.e., if v_i is an in-vertex. If i is odd, i.e., v_i an out-vertex, we place the right optional edge *above* the left optional edge.

We repeat this procedure for all C_i, $i \geqslant 2$. Afterwards, we represent each mandatory edge by a rectilinear path consisting of two vertical and one horizontal line segment, such that the appropriate boxes are connected. Again, we make sure that all vertical coordinates are distinct.

Next, consider some fixed out-vertex v with out-degree two, and let Bo_v be its corresponding box. Depending on whether v is of type A or B, we distinguish the following cases.

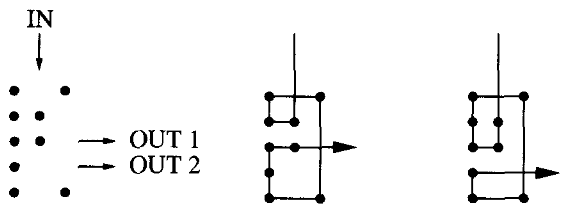

Fig. 25. Box type A: OUT1 and OUT2 leave on the same side.

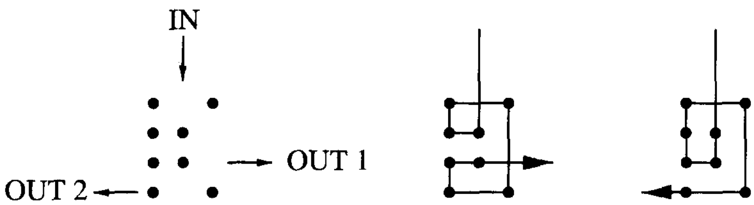

Fig. 26. Box type B: OUT1 and OUT2 leave on different sides.

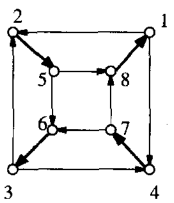

Fig. 27. A cubic digraph D.

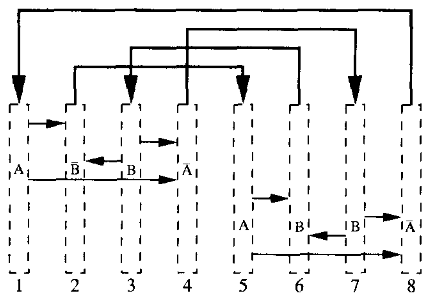

Fig. 28. The representation for D.

Case A. If both outgoing edges leave the box on the same side, we place a stretched copy of the set A as depicted in Fig. 25 into the box Bo_v. This is done in such a way that the heights of OUT1 and OUT2 coincide with the heights of the line segments corresponding to the two outgoing edges. If (u, v) is the mandatory edge adjacent to v, reroute it into IN by placing one appropriate point at each right turn of (u, v) and three appropriate points at each left turn of (u, v).
Case B. If both outgoing edges leave the box on different sides, we use the point set B depicted in Fig. 26 in a similar manner as point set A in Case A.

Finally, if the vertex v is an in-vertex, we simply use reflected versions of the two point sets in Case A (Fig. 25) and Case B (Fig. 26). Thus inputs become outputs and vice versa. We denote the resulting point sets by $\overline{A}$ and $\overline{B}$ in Fig. 28.

To provide the reader with some intuition for our constructions, we state the following observations. For an example, see Fig. 28.

(a) Groups of points in different boxes have different coordinates with the exception of the three in- and outputs. Consequently, a strongly feasible tour will enter a box, run through all points, leave the box, and it will never enter it again (otherwise it would have to use at least two inputs and at least two outputs).

(b) The point set used to reroute a mandatory edge *must* be used in each strongly feasible tour, since there is no possibility for a strongly feasible tour to enter through OUT1 and to leave through OUT2 (or to enter through OUT2 and to leave through OUT1) and to visit all points inside the box at the same time.

(c) If we reflect an A_3-path from a to b, it becomes a strongly feasible A_3-path from b to a. Hence, for in-vertices, the paths in the figures are strongly feasible paths from OUT to IN.

Now assume we are given a Hamiltonian cycle for G. We want to construct a strongly feasible A_2-tour visiting all points in the point set P_G. First, we insert all line segments corresponding to edges in G that belong to the Hamiltonian cycle. Next we consider a box corresponding to some out-vertex. We know at which positions the box is entered and at which positions it is left again. We include the point set that was used to reroute the input from IN to the height of the ingoing edge. Then we use the corresponding sets of segments shown in Figs. 25 and 26 to connect IN to one of the outgoing segments. (In case of an in-vertex we use the reflected versions.) Clearly, we end up with an A_2-tour.

Finally, assume that the constructed point set allows an A_2-tour $\mathcal{T}$ and consider some out-vertex v. The only possibilities for $\mathcal{T}$ to enter and to leave the box corresponding to v is via the line segments OUT1, OUT2 and IN corresponding to the three incident edges (all other points have distinct x- and y-coordinates). It is easily checked that the tour cannot visit all points inside the box in a valid way if it enters through OUT1 or OUT2 or if it leaves through IN. Hence, the tour must enter all boxes at places corresponding to ingoing edges and it must leave all boxes at places corresponding to outgoing edges. Contracting the subpaths of the tour in every single box yields a Hamiltonian cycle in G. Our proof is complete. □

The polynomial time result on A_3-tours is a direct implication of the following proposition due to O'Rourke [27].

Proposition 4.3 (O'Rourke [27]). *For a set P of n points in the Euclidean plane, it can be checked in $O(n \log n)$ time whether P allows a strongly feasible A_3-tour such that all segments in the tour are parallel to the x- or to the y-axis. Furthermore, there can be at most one strongly feasible A_3-tour for any orientation of the axes.*

This yields the following Corollary 4.4. For the sake of the proof of the following Theorem 4.5, we describe the proof idea.

Corollary 4.4. *For a set P of n points in the Euclidean plane, it can be checked in $O(n^2 \log n)$ time whether P admits a strongly feasible A_3-tour.*

Proof. See Fig. 29. Consider some fixed line segment s in some valid A_3-tour $\mathcal{T}$. Then all other line segments in $\mathcal{T}$ are either parallel or normal to s. That means we may assume an underlying coordinate system such that all segments in $\mathcal{T}$ are parallel to one of the axes. Since there are only $n-1$ segments emanating from some fixed point in P, there are only $O(n)$ possible orientations for such a coordinate system.

For a fixed coordinate system, consider some horizontal line that contains at least one point of P. Let $p_1, p_2, \ldots, p_k$ be the points of P on this line sorted from left to right. Assume there exists an A_3-tour $\mathcal{T}$ that visits all points of P. Then the point p_1 has two neighbors on the polygonal line $\mathcal{T}$,

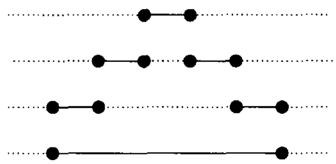

Fig. 29. Finding strongly feasible A_3-tours.

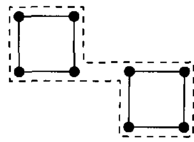

Fig. 30. Merging subtours to find weakly feasible A_3-tours.

one of them vertical, the other one horizontal. The only possible horizontal neighbor is the point p_2. Hence, we may connect p_1 to p_2. Similarly, the only possible horizontal neighbor of p_3 is p_4, the only possible horizontal neighbor of p_5 is p_6 and so on. This yields $k/2$ segments that must be part of $\mathcal{T}$. We repeat this process for every horizontal and every vertical line that contains at least one point of P. If we meet some line that contains an *odd* number of points, we stop immediately, as a tour cannot exist for the given orientation.

We end up with a set of n line segments, such that every point of P is incident to exactly two segments. If the segments form a single closed polygonal line we have constructed a (strongly feasible) A_3-tour, otherwise none exists for the orientation considered. □

Furthermore, we can extend the result to weakly feasible A_3-tours.

Theorem 4.5. *Let P be a set of n points in the Euclidean plane. Then we can decide in $O(n^2 \log n)$ time whether P allows a weakly feasible A_3-tour.*

Proof. If a solution exists for some fixed coordinate system, the method described in Corollary 4.4 either detects a strictly feasibly tour (and we are done) or it will find a partition of P into disjoint sets $P_1, \ldots, P_k$ such that each P_i allows a weakly feasible A_3-tour. Assume there exist two points $p_i \in P_i$ and $p_j \in P_j$ such that p_i and p_j have identical x- or y-coordinate. In this case, there exists an axes-parallel line containing a segment of the tour for P_i and a segment of the tour for P_j, and it is easy to see that the A_3-tours for P_i and P_j can be merged into a weakly feasible A_3-tour for $P_i \cup P_j$. (See Fig. 30.)

Since we may repeat this merging procedure, we see that the problem reduces to determining whether there exists an orthogonal coordinate system such that P is connected with respect to axes-parallel connections. But this is done easily in time $O(n^2 \log n)$ by checking $n - 1$ coordinate systems and sorting the corresponding x- and y-coordinates. □

In [27], O'Rourke stresses the aspect of *uniqueness* of strongly feasible A_3-tours. He motivates the question as a pattern recognition problem: We are trying to find a simple rectilinear polygon for a given set of vertices. For his uniqueness result, he assumes that the orientation of the polygon edges is axes-parallel in some given coordinate system.

If we have to find the underlying rectilinear polygonal shape of a given set of vertices in the plane, we cannot necessarily assume knowledge of an underlying coordinate system. Interestingly enough, we can show that O'Rourke's uniqueness result remains valid even if no orientations are known in advance. This nicely illustrates our point about the relation between relative and fixed restricted orientations.

Definition 4.6. Given a subset V of the Euclidean plane $\mathbb{R}^2$. For any orientation O of a Cartesian coordinate system, we define the *O-orthogonality graph* $G_O(V) = (V, E)$ as follows: two vertices v_1 and v_2 are connected by an edge, if and only if they have the same x- or y-coordinate.

Theorem 4.7. *Given a set P of n vertices in the Euclidean plane that all have rational coordinates. Then there is at most one orientation O for which the graph $G_O(P)$ is connected.*

Proof. Suppose that we have two such orientations O_1 and O_2. Then we can write the direction of the x-axis in O_2 as (a, b) with $a, b \neq 0$ in O_1-coordinates. Since all the coordinates are rational, we may assume that a and b are relatively prime integers. Furthermore we may assume that all O_1-coordinates of the given points are relatively prime integers. This implies that there are two points of the same y-coordinate that have a positive distance in x-direction that is not a multiple of $a^2 + b^2 > 1$, say k.

Since these two points can be connected by an axis-parallel path with orientation (a, b) and all endpoints of segments being grid points, the vector $(k, 0)$ must be an integer combination of the vectors (a, b) and $(-b, a)$, i.e., the system

$$ax - by = k, \qquad bx + ay = 0$$

must have an integer solution (x, y)—which is equivalent to

$$(a^2 + b^2)x = ka, \qquad (a^2 + b^2)y = kb.$$

Since a and b are relatively prime, $a^2 + b^2$ and a are, so a must divide x. This implies that k is a multiple of $(a^2 + b^2)$; a contradiction. □

5. Acute and obtuse angles

In this section, we will discuss tours and paths for acute or obtuse angles. The set A^{acut} may be interpreted as the requirement to make "sharp" turns at every vertex; conversely, A^{obtu} corresponds to the situation where we have to avoid too sudden changes of direction. As already mentioned in the introduction, problems of the latter kind play an important role in the planning of roundtrips with constrained curvature.

Mathematically, the feasible angle sets for pseudoconvex, acute and obtuse tours can all be represented as one (appropriate) half of the unit circle $S^1 = \mathbb{R} \bmod 2\pi$. The following statements, however, demonstrate that each of these three angle sets behaves quite differently with respect to angle-restricted tours.

Example 5.1. Obtuse spanning paths do not exist for all point sets. To see this, consider the three vertices of an equilateral triangle together with $n-3$ points in the interior of the triangle. (See Fig. 31.) Any spanning path must contain one of the corner points somewhere in its middle. But then the corner point together with its two neighbors in the path form an obtuse angle.

Theorem 5.2. *Every finite set of points has an acute path.*

Proof. See Fig. 32. Choose any point p_0 to be the starting point and remove p_0 from P. p_1 is a point in P that is farthest from p_0, and we remove p_1 from P. p_2 is a point in P that is farthest from p_1,

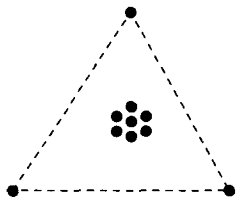

Fig. 31. A point set without obtuse spanning paths.

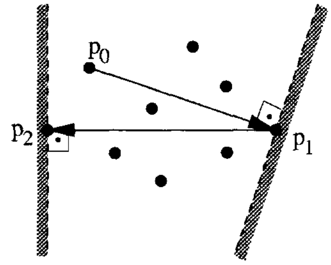

Fig. 32. Every finite set of points has an acute path.

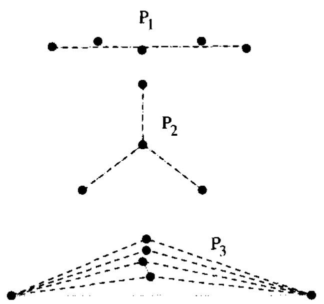

Fig. 33. Sets that do not allow acute tours.

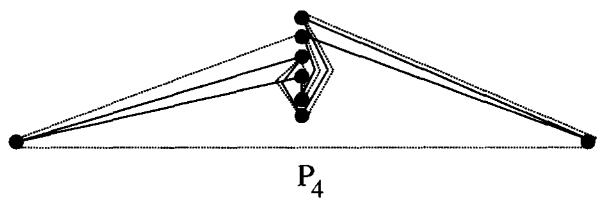

Fig. 34. Some difficulties for acute tours.

and so on. This procedure can never produce an angle $\angle(p_i p_{i+1} p_{i+2})$ that is not acute, as in this case p_{i+2} is further from p_i than p_{i+1}, and we would have chosen p_{i+2} to be the successor of p_i. □

The situation for acute tours is more complicated.

Example 5.3. Let P_1 be a point set containing an odd number of points on the x-axis (maybe perturbated a little bit to get a point set in general position). Suppose P_1 has an acute tour. We call a segment a western (eastern) segment, if the acute tour traverses it in direction of negative (positive) x. The tour must consist of an alternating sequence of eastern and western segments: if the tour traverses some western line segment $\overline{xy}$, it must traverse the following segment $\overline{yz}$ in eastern direction to avoid an obtuse angle at the point y. This forces the tour to use an even number of line segments and, consequently, an even number of points.

An even number of points still does not guarantee an acute tour; see the point sets P_2 and P_3 depicted in Fig. 33.

Example 5.4. The point set P_2 containing the four points $(0,-10)$, $(0,10)$, $(5,0)$, $(10,0)$ does not allow an acute tour. The point set P_3 containing the six points $(-10,0)$, $(10,0)$, $(0,1)$, $(0,2)$, $(\varepsilon,3)$ and $(\varepsilon,4)$ (where $\varepsilon < 10^{-10}$ is some small real) is in general position and does not allow an acute tour.

Surprisingly enough, Example 5.4 does not generalize to larger even numbers of points. In fact, we have the following conjecture.

Conjecture 5.5. Every set P of $2k \geqslant 8$ points admits an acute tour.

After the construction of acute spanning paths in Theorem 5.2, one might be tempted to conjecture some kind of relation between acute paths, acute tours and tours of maximum overall length. However, the following Example 5.6 points out a few difficulties.

Example 5.6. For a positive integer m, the point set P_4 contains the $2m$ points $(m^3, 0)$, $(-m^3, 0)$, $(0, 1)$, $\ldots$, $(0, 2m-2)$ as shown in Fig. 34. One easily checks that for $m \geqslant 4$, the following statements hold.
- P_4 allows an acute tour.
- No acute tour contains the diameter $\overline{(m^3, 0)(-m^3, 0)}$.
- It is impossible to get an acute tour by using the approach for acute spanning paths as described in Theorem 5.2; regardless of the starting point, the diameter would be included in the second step at the latest.
- Any longest tour has length $2m^3 + \sqrt{m^6 + (2m-2)^2} + \sqrt{m^6 + (2m-3)^2} + 4m - 3$ and contains the diameter (indicated by the broken lines in Fig. 34), while a longest acute tour (solid lines) has only length

$$\sqrt{m^6 + (2m-2)^2} + \sqrt{m^6 + (2m-3)^2} + \sqrt{m^6 + (2m-4)^2} + \sqrt{m^6 + (2m-5)^2} + 4m - 4,$$

which is smaller.

Finding spanning paths that minimize the maximum angle or maximize the minimum angle seems to be a difficult algorithmic problem. By the above algorithm for finding acute spanning paths, we can always guarantee a maximum angle of absolute size at most $\pi/2$. A set of $n-1$ collinear points with an nth point far away from the line show that this trivial guarantee is also best possible.

As for lower bounds on the minimum angle, the equilateral triangle together with its centerpoint demonstrates that no angle greater than $\pi/6$ can be guaranteed. We have the following conjecture.

Conjecture 5.7. Every planar set P admits a spanning path with minimum angle at least $\pi/6$.

6. Other questions

We conclude this paper with some more open questions.

Conjecture 6.1. Detecting $\{-2\pi/3, +2\pi/3\}$-tours can be done in polynomial time.

Conjecture 6.2. Detecting $\{-\pi/3, +\pi/3, \pi\}$-tours is NP-complete.

To our knowledge, it has not been established that detecting Hamiltonian cycles in vertex-induced subgraphs of the hexagonal grid is NP-complete. A proof of this would certainly be helpful.

Problem 6.3. What is the computational complexity of detecting acute/obtuse tours?

For this problem as well as for Conjecture 5.7, some of the results on triangulations that minimize the maximal angle or maximize the minimal angle may be useful. Delaunay triangulations may appear as a good choice. Note, however, that Delaunay triangulations or other minimum weight triangulations are in general *not* necessarily Hamiltonian (see Dillencourt [10,11]).

References

[1] A. Aggarwal, D. Coppersmith, S. Khanna, R. Motwani and B. Schieber, Angular-metric TSP and the matroid parity problem, in: Proc. 28th Symp. Theory Comput. (1996).

[2] P.K. Agarwal, P. Raghavan and H. Tamaki, Motion planning for a steering-constrained robot through moderate obstacles, in: Proc. 27th Symp. Theory Comput. (1995) 343–352.

[3] J. Barraquand and J.-C. Latombe, Nonholonomic mobile robots and optimal maneuvering, Rev. Intellig. Artific. 3 (1989) 77–103.

[4] J. Barraquand and J.-C. Latombe, Nonholonomic multibody mobile robots: controllability and motion planning in the presence of obstacles, Algorithmica 10 (1993) 121–155.

[5] J. Boissonat, J. Czyzowicza, O. Devillers, J.-M. Robert and M. Yvinec, Convex tours of bounded curvature, in: Proc. 2nd European Symp. (ESA 94) (Springer, Berlin, 1994) 254–265.

[6] J. Boissonat, A. Cérézo and J. Leblond, Shortest paths of bounded curvature in the plane, J. Intellig. Robotic Systems 11 (1994) 5–20.

[7] B.M. Chazelle, On the convex layers of a planar set, IEEE Trans. Inform. Theory IT-31 (1985) 509–517.

[8] J.C. Clements, Minimum-time turn trajectories to fly-to points, Optimal Control Appl. Methods 11 (1990) 39–50.

[9] J. Culberson and G.J.E. Rawlins, Turtlegons: generating simple polygons from given sequences of angles, in: Proc. ACM Symp. Comput. Geom. (1985) 305–310.

[10] M.B. Dillencourt, Traveling salesman cycles are not always subgraphs of Delaunay triangulations or of minimum weight triangulations, Inform. Process. Lett. 24 (1987) 339–342.

[11] M.B. Dillencourt, A non-Hamiltonian, nondegenerate Delaunay triangulation, Inform. Process. Lett. 25 (1987) 149–151.

[12] H. Edelsbrunner, Algorithms in Combinatorial Geometry (Springer, Berlin, 1987).

[13] S.P. Fekete, Geometry and the Traveling Salesman Problem, Ph.D. Thesis, University of Waterloo, Waterloo, Canada (1992).

[14] S.P. Fekete and W.R. Pulleyblank, Backward error analysis for the Traveling Salesman Problem: generalized convexity, ZPR Technical Report 93-142, Center for Parallel Computing, Universität zu Köln, Germany (1993).

[15] T. Frachard, Smooth trajectory planning for a car in a structured world, in: Proc. IEEE Internat. Conf. Robotics (1989) 318–323.

[16] M.R. Garey and D.S. Johnson, Computers and Intractability: A Guide to the Theory of NP-completeness (Freeman, San Francisco, 1979).

[17] A. Itai, C. Papadimitriou and J.L. Swarcfiter, Hamilton paths in grid graphs, SIAM J. Computing 11 (1982) 676–686.

[18] P. Jacobs and J. Canny, Planning smooth paths for mobile robots, in: Proc. IEEE Internat. Conf. Robotics (1989) 2–7.

[19] J.-C. Latombe, Robot Motion Planning (Kluwer Academic Press, 1991).

[20] J.-P. Laumond, Finding collision-free trajectories for a nonholonomic mobile robot, in: Proc. IEEE Internat. Conf. Robotics (1987) 1120–1123.

[21] D. Rappaport, Computing simple circuits from a set of line segments is NP-complete, SIAM J. Computing 18 (1989) 1128–1139.
[22] D. Rappaport, On the complexity of computing orthogonal polygons from a set of points, McGill University Technical Report No. SOCS-86.9 (1986).
[23] G.J.E. Rawlins, Explorations in restricted orientation geometry, Ph.D. Thesis, University of Waterloo, Waterloo, Canada (1987); available as Research Report CS-87-57.
[24] G.J.E. Rawlins and D. Wood, On the optimal computation of finitely-oriented convex hulls, Inform. and Comput. 72 (1987) 150–166.
[25] G.J.E. Rawlins and D. Wood, Computational geometry with restricted orientations, in: Proc. 13th IFIP Conf. System Modeling and Optimization, Lecture Notes in Computer Science (Springer, Berlin, 1988).
[26] G.J.E. Rawlins and D. Wood, Orthoconvexity and its generalizations, in: G.T. Toussaint, ed., Computational Morphology (North-Holland, Amsterdam, 1988) 137–152.
[27] J. O'Rourke, Uniqueness of orthogonal connect-the-dots, in: G.T. Toussaint, ed., Computational Morphology (North-Holland, Amsterdam, 1988) 97–104.
[28] S. Schuierer, On generalized visibility, Doctoral Thesis, Albert-Ludwigs-Universität Freiburg im Breisgau, Germany (1991).
[29] P. Widmayer, Y.F. Wu and C.K. Wong, On some problems in fixed orientations, SIAM J. Computing 16 (1987) 728–746.