

Covering Paths for Planar Point Sets^{*}

Adrian Dumitrescu¹ and Csaba D. Tóth²

¹ Computer Science, University of Wisconsin–Milwaukee
dumitres@uwm.edu

² Mathematics and Statistics, University of Calgary
cdtoth@ucalgary.ca

Abstract. Given a set of points, a *covering path* is a directed polygonal path that visits all the points. We show that for any n points in the plane, there exists a (possibly self-crossing) covering path consisting of $n/2 + O(n/\log n)$ straight line segments. If no three points are collinear, any covering path (self-crossing or non-crossing) needs at least $n/2$ segments. If the path is required to be non-crossing, $n - 1$ straight line segments obviously suffice and we exhibit n -element point sets which require at least $5n/9 - O(1)$ segments in any such path. Further, we show that computing a non-crossing covering path for n points in the plane requires $\Omega(n \log n)$ time in the worst case.

1 Introduction

In this paper we study polygonal paths visiting a finite set of points in general position (we say that a point set is in *general position* if no three points are collinear). A *spanning path* is a directed Hamiltonian path drawn with straight line edges. Each edge in the path connects two of the points, so a spanning path can only turn at one of the given points. Every spanning path of a set of n points in general position consists of $n - 1$ segments. A *covering path* is a directed polygonal path in the plane that visits all the points. A covering path can make a turn at any point, i.e., either at one of the given points or at a (chosen) Steiner point. Obviously, a spanning path for a point set S is also a covering path for S . If no three points are collinear, every covering path consists of at least $n/2$ segments. Given a set of points, a *minimum-link* covering path is one with the smallest number of segments (links).

We study the following two questions concerning covering paths posed by Morić [20] and (independently) B. Keszegh, both motivated by a problem about separating a red and blue set of points [13]. A recent consideration of these questions in retrospect appears in [9].

1. What is the minimum number, $f(n)$, such that any set of n points in the plane, no 3 collinear, can be covered by a (possibly self-intersecting) polygonal path with $f(n)$ segments?

^{*} Dumitrescu acknowledges support from the NSF grant DMS-1001667. Tóth acknowledges support from the NSERC grant RGPIN 35586 and the NSF grant CCF-0830734. Part of the research was conducted at the Fields Institute, Toronto, ON.

2. What is the minimum number, $g(n)$, such that any set of n points in the plane, no 3 collinear, can be covered by a *non-crossing* polygonal path with $g(n)$ segments?

Since no three points are collinear, each segment of a covering path contains at most two points, thus $n/2$ is a trivial lower bound for both $f(n)$ and $g(n)$. Morić conjectured that the answer to the first problem is $n(1 + o(1))/2$ while the answer to the second is $n(1 + o(1))$. He only mentioned trivial bounds, $n/2$ as a lower bound, and respectively $n - 1$ as an upper bound for both problems. We start by confirming his first conjecture¹.

Theorem 1. *For any set of n points in the plane there exists a (possibly self-crossing) path consisting of $n/2 + O(n/\log n)$ line segments that visits all the points. Consequently, $n/2 \leq f(n) \leq n/2 + O(n/\log n)$. A covering path with $n/2 + O(n/\log n)$ segments can be computed in $O(n^{1+\varepsilon})$ time, for any $\varepsilon > 0$.*

As expected, the non-crossing property is much harder to deal with. For any set of n points in the plane, trivially there exists a non-crossing path consisting of $n - 1$ straight line segments that visits all the points, e.g., by sorting the points along some direction, and then connecting them in this order. On the other hand, again trivially, any such covering path requires at least $\lceil n/2 \rceil$ segments, since no three points are collinear. We provide the first non-trivial lower bound for $g(n)$.

Theorem 2. *There exist n -element point sets in general position that require at least $\frac{5}{9}(n - 4)$ segments in any non-crossing covering path. Consequently, $\frac{5}{9}(n - 4) \leq g(n) \leq n - 1$.*

Bicolored Variant. Let S be a bicolored set of n points, with $S = B \cup R$, where B and R are the set of blue and red points, respectively. Two covering paths, π_R and π_B , one for the red and one for the blue points, are *mutually non-crossing* if each of π_R and π_B is non-crossing, and moreover, π_R and π_B do not cross (intersect) each other. A natural extension of the monochromatic non-crossing covering path problem is: What is the minimum number $k(n)$, such that any bicolored set of n points in the plane, can be covered by two monochromatic mutually non-crossing polygonal paths with $k(n)$ segments in total? Using the construction in the proof of Theorem 2 we obtain the following corollary.

Corollary 1. *Given a bicolored set of n points, one can find two mutually non-crossing covering paths with a total of at most $3n/2 + O(1)$ segments. Such a pair of paths can be computed in $O(n \log n)$ time. On the other hand, there exist bicolored sets that require at least $5n/9 - O(1)$ segments in any pair of mutually non-crossing covering paths. Consequently, $5n/9 - O(1) \leq k(n) \leq 3n/2 + O(1)$.*

Computational Complexity. We establish an $\Omega(n \log n)$ lower bound for computing a non-crossing covering path with $O(n)$ vertices for a set of n points in the plane.

¹ This was observed by the current authors during the conference CCCG 2010; it was also communicated to the authors of [9].

Theorem 3. *The sorting problem for n numbers is linear-time reducible to the problem of computing a noncrossing covering path for n points in the plane. Therefore, computing a noncrossing covering path for a set of n points in the plane requires $\Omega(n \log n)$ time in the worst case in the algebraic decision tree model of computation.*

Related previous results. Given a set of points, the MINIMUM-LINK COVERING PATH problem asks for a covering path with the smallest number of segments (links). Arkin et al. [2] proved that (the decision version of) this problem is NP-complete. Stein and Wagner [22] gave a $O(\log z)$ -approximation where z is the maximum number of collinear points.

Various upper and lower bounds on the minimum number of links needed in an axis-aligned path traversing an n -element point set in \mathbb{R}^d have been obtained in [3,8,7,16]. Approximation algorithms with constant ratio (depending on the dimension d) for this problem are developed in [3], while some NP-hardness results have been claimed in [12], and further revised in [15]. Other variants of Euclidean TSP can be found in the survey article by Mitchell [19].

2 Covering Paths with Possible Self-crossings

A set X of k points in general position in the plane, no two on a vertical line, is a k -cap (k -cup, respectively) if X is in convex position and all points of X lie above (below, respectively) the line connecting the leftmost and the rightmost point of X . According to a classical result of Erdős and Szekeres [10], any set of at least $\binom{2k-4}{k-2} + 1$ points in general position in the plane, no two on a vertical line, contains a k -cup or a k -cap. In particular, every such set contains k points in convex position; see also [11,17]. They also showed that this bound is the best possible, i.e., there exist sets of $\binom{2k-4}{k-2}$ points containing no k -cup or k -cap. More generally, there exist sets of $\binom{k+l-4}{k-2}$ points containing no k -cup or l -cap.

Following the terminology coined by Welzl, a set S of n points in the plane, no 3 of which are collinear, is called *perfect* if it can be covered by a (possibly self-crossing) polygonal path consisting of at most $\lceil n/2 \rceil$ segments. It is easy to see that a cup or a cap is perfect: indeed, a suitable covering path can be obtained by extending the odd numbered edges of the x -monotone polygonal chain connecting the points (note that since no two points lie on a vertical line, any consecutive pair of these edges properly intersect).

Proof of Theorem 1. Let S be a set of n points in the plane, no three of which are collinear. Choose an orthogonal coordinate system such that no two points have the same x -coordinate. By the result of Erdős and Szekeres [10], every m -element subset of S contains a k -cup or a k -cap for some $k = \Omega(\log m)$. Since any such subset is perfect, it can be covered by a path of $\lceil k/2 \rceil$ segments.

To construct a covering path, we partition S into caps and cups of size $\Omega(\log n)$ each, and a set of less than $n/\log n$ “leftover” points. Set $T = S$. While $|T| \geq n/\log n$, repeatedly find a maximum-size cup or cap in T and

delete those elements from T . Note that $\log(n/\log n) = \Omega(\log n)$, and we have found a k -cup or k -cap for some $k = \Omega(\log n)$ in each step. Therefore, we have found $O(n/\log n)$ pairwise disjoint caps and cups in S , and we are left with a set T of less than $n/\log n$ points.

For each k -cup (or k -cap), construct a covering sub-path with $\lceil k/2 \rceil$ segments. Link these paths arbitrarily into one path, that is, append them one after another in any order. Finally append to this path an arbitrary spanning path of the remaining less than $n/\log n$ points in T , with one point per turn.

A covering path for S is obtained in this way. It is easy to see that the total number of segments in this path is $n/2 + O(n/\log n)$, as required. Chvátal and Klincsek [6] showed that a maximum-size cap (and cup) in a set of n points in the plane, no 3 of which are collinear, can be found in $O(n^3)$ time. With $O(n/\log n)$ calls to their algorithm, a covering path with $n/2 + O(n/\log n)$ segments can be constructed in $O(n^4/\log n)$ time in the RAM model of computation. Now if the problem can be solved in time $O(n^4/\log n)$, it can also be solved in time $O(n^{1+\varepsilon})$ for any $\varepsilon > 0$: arbitrarily partition the points into $n^{1-\varepsilon/3}$ subsets of $n^{\varepsilon/3}$ points each, solve each subset separately, move the leftover points to the next subset, then link the paths together with one extra segment per path. \square

3 Noncrossing Covering Paths

Proof of Theorem 2 (outline). For every $k \in \mathbb{N}$, we construct a set S of $n = 2k$ points in the plane in general position, where all points are very close to the parabola $x \rightarrow x^2$. We then show that every noncrossing covering path γ consists of at least $\frac{5}{9}(n - 4)$ segments. The lower bound is based on a charging scheme: we distinguish *perfect* and *imperfect* segments in γ , containing 2 and less than 2 points of S , respectively. We charge every perfect segment to a “nearby” endpoint of an imperfect segment or an endpoint of γ , such that each of these endpoints is charged at most twice. This implies that at most about $\frac{4}{5}$ of the segments are perfect, and the lower bound of $\frac{5}{9}(n - 4)$ follows. We continue with the details.

A technical lemma. We start with a simple lemma, showing that certain segments in a noncrossing covering path are almost parallel. We say that a line segment s *traverses* a circular disk B if s intersects the boundary of B twice.

Lemma 1. *Let $\varphi \in (0, \frac{\pi}{2})$ be an angle. For every $\varepsilon > 0$, there is a $\delta \in (0, \varepsilon)$ such that if two noncrossing line segments ab and cd both traverse two concentric disks of radii ε and δ , then the supporting lines of the segments ab and cd meet at an angle at most φ .*

Proof. Let ab and cd be two noncrossing line segments that both traverse two concentric disks of radii $\varepsilon > \delta > 0$. Refer to Fig. 1. By translating the segments, if necessary, we may assume that both are tangent to the disk of radius δ . Clip the segments in the disk of radius ε to obtain two noncrossing chords. The angle between two noncrossing chords is maximal if they have a common endpoint. In this case, they meet at an angle $2 \sin^{-1}(\delta/\varepsilon)$. For every $\varepsilon > 0$, there is a $\delta \in (0, \varepsilon)$ such that $2 \sin^{-1}(\delta/\varepsilon) < \varphi$. \square




Fig. 1. Left: Two noncrossing segments, ab and cd , traverse two concentric disks of radii $\varepsilon > \delta > 0$. Middle: The segments are translated to be tangent to the disk of radius δ . Right: The angle between two adjacent chords is $2 \sin^{-1}(\delta/\varepsilon)$.

Construction. For every $k \in \mathbb{N}$, we define a set $S = \{a_1, \dots, a_k, b_1, \dots, b_k\}$ of $n = 2k$ points. Initially, let $A = \{a_1, \dots, a_k\}$ be a set of k points on the first-quadrant part of the parabola $\alpha : x \rightarrow x^2$ such that no two lines determined by A are parallel. (To achieve strong general position, we shall slightly perturb the points in S in the last step of the construction.) Label the points in A by a_1, \dots, a_k in increasing order of x -coordinates. Each point b_i will be in a small δ -neighborhood of a_i , for $i = 1, \dots, k$. The pairs $\{a_i, b_i\}$ are called *twins*. The value of $\delta > 0$ is specified in the next paragraph.

For every $r > 0$, let $D_i(r)$ denote the disk of radius r centered at $a_i \in A$. Since the points in A are in strictly convex position, points in A determine $\binom{k}{2}$ distinct lines. Let $(2\varphi) \in (0, \frac{\pi}{2})$ be the minimum angle between two lines determined by A (recall that no two such lines are parallel). Let $\varepsilon > 0$ be sufficiently small such if two lines intersect two different pairs of disks from $\{D_1(\varepsilon), \dots, D_k(\varepsilon)\}$, then they meet at an angle at least φ . Observe that every line intersects at most two disks $D_1(\varepsilon), \dots, D_k(\varepsilon)$ (i.e., the ε -neighborhoods of at most two points in A). By Lemma 1, there exists $\delta_0 > 0$ such that if two noncrossing segments traverse both $D_i(\varepsilon)$ and $D_i(\delta_0)$, then their supporting lines meet at an angle less than φ . For $i = 1, \dots, k-1$, let $\delta_i > 0$ be the maximum distance between the supporting line of $a_i a_{i+1}$ and points on the arc of the parabola α between a_i and a_{i+1} . We are now ready to define $\delta > 0$. Let $\delta = \min\{\delta_i : i = 0, 1, \dots, k-1\}$.

We now choose points $b_i \in D_i(\delta)$, for $i = 1, \dots, k$, in *reverse* order. Let ℓ_k be a line that passes through a_k whose slope is larger than the tangent of the parabola α at a_k . Let b_k be a point in $\ell_k \cap D_k(\delta)$ above the parabola α . Having defined line ℓ_j and point b_j for all $j > i$, we choose ℓ_i and $b_i \in \ell_i \cap D_i(\delta)$ as follows:

- let ℓ_i be a line passing through a_i whose slope is larger than that of ℓ_{i+1} ;
- let $b_i \in \ell_i \cap D_i(\delta)$ be above the parabola α ; and
- let b_i be so close to a_i that for every $j, i < j \leq k$, the supporting lines of segments $a_i a_j$ and $b_i b_j$ meet in the disk $D_i(\varepsilon)$.

We also ensure in each iteration that in the set $A \cup \{b_1, \dots, b_k\}$, (1) no three points are collinear; (2) no two lines determined by the points are parallel; and (3) no three lines determined by disjoint pairs of points are concurrent.

Note that S is not in strong general position: for instance, all points in A lie on a parabola. (By strong general position it is meant here there is no nontrivial algebraic relation between the coordinates of the points.) In the last step of our construction, we slightly perturb the points in S . However, for the analysis of a minimum link covering path, we may ignore the perturbation.

Let γ be a minimum-link noncrossing covering path for S . Since no three points in S are collinear, we may assume that every point in S lies in the relative interior of a segment of γ . Denote by s_0, s_1 and s_2 , resp., the number of segments in γ that contain 0, 1, and 2 points from S . We establish the following inequality.

Lemma 2. $s_2 \leq 4(s_0 + s_1 + 1)$.

Before the proof of Lemma 2, we show that it immediately implies Theorem 2. Counting the number of points incident to the segments, we have $n = s_1 + 2s_2$. The total number of segments in γ is

$$\begin{aligned} s_0 + s_1 + s_2 &= \frac{4(s_0 + s_1 + 1) + 5s_0 + 5s_1 - 4}{9} + s_2 \\ &\geq \frac{s_2 + 5s_0 + 5s_1 - 4}{9} + s_2 \\ &\geq \frac{5(s_1 + 2s_2) - 4}{9} = \frac{5n - 4}{9} \end{aligned}$$

For proving Lemma 2, we introduce a charging scheme, where each perfect segment is charged to either an endpoint of an imperfect segment, or to one of the two endpoints of γ such that each such endpoint is charged at most twice. The charges will be defined for maximal x -monotone chains of perfect segments. A subpath $\gamma' \subseteq \gamma$ is called x -monotone, if the intersection of γ' with any vertical line is connected (i.e., the empty set, a point, or a vertical segment).

Recall that all points in $A = \{a_1, \dots, a_k\}$ lie on the parabola $\alpha : x \rightarrow x^2$. Let β be the graph of a strictly convex function that passes through the points b_1, \dots, b_k , and lies strictly above α and below the curve $x \rightarrow x^2 + \delta$.

Properties of a minimum noncrossing covering path. We start by characterizing the perfect segments in γ . Note that if pq is a perfect segment in γ , then pq contains either a twin, or one point from each of two twins. First we make a few observations about perfect segments containing points from two twins.

Lemma 3. *Let pq be a perfect segment in γ that contains one point from each of the twins $\{a_i, b_i\}$ and $\{a_j, b_j\}$, $i < j$. Then pq intersects both $D_i(\delta)$ and $D_j(\delta)$, and its endpoints lie below the curve β .*

Proof. The distance between any two twin points is less than δ , so pq intersects the δ -neighborhood of a_i and a_j (even if pq contains b_i or b_j). The line pq intersects the parabolas $\alpha : x \rightarrow x^2$ and $x \rightarrow x^2 + \delta$ twice each. It also intersects β exactly twice: at least twice, since β is between the two parabolas; and at most twice since the region above β is strictly convex. All points in $\{a_i, b_i\}$ and $\{a_j, b_j\}$ are on or below β ; but pq is above β at some point between its intersections with $\{a_i, b_i\}$ and $\{a_j, b_j\}$, since $\delta \leq \delta_i$. Hence the endpoints of pq are below β . \square

Lemma 4. *Let pq be a perfect segment of γ that contains one point from each of the twins $\{a_i, b_i\}$ and $\{a_j, b_j\}$, with $i < j$. Assume that p is the left endpoint of pq . Let s be the segment of γ containing the other point of the twin $\{a_i, b_i\}$. Then one of the following four cases occurs.*

- Case 1: p is incident to an imperfect segment of γ , or it is an endpoint of γ ;
- Case 2: s is imperfect;
- Case 3: s is perfect, one of its endpoints v lies in $D_i(\varepsilon)$, and v is either incident to some imperfect segment or it is an endpoint of γ ;
- Case 4: s is perfect and p is the common left endpoint of segments pq and s .




Fig. 2. The four cases in Lemma 4 for a perfect segment pq that contains one point from each of the twins $\{a_i, b_i\}$ and $\{a_j, b_j\}$. The points $a_i, i = 1, \dots, k$, lie on the dotted parabola α .

Proof. If p is incident to an imperfect segment of γ , or it is an endpoint of γ , then Case 1 occurs. Assume that p is incident to two perfect segments of γ , pq and pr . If $pr = s$, then p is the common left endpoint of two perfect segments, pq and s , and Case 4 occurs. If s is imperfect, then Case 2 occurs.

Assume now that $pr \neq s$ and s is perfect. We need to show that Case 3 occurs. We claim that the segment pq traverses $D_i(\varepsilon)$. It is enough to show that p and q lie outside of $D_i(\varepsilon)$. Note that pr does not contain any point from the twin $\{a_i, b_i\}$ (these points are covered by segments pq and s). Since pr is a perfect segment, it contains two points from $S \setminus \{a_i, b_i\}$. By construction, any

line determined by $S \setminus \{a_i, b_i\}$ is disjoint from $D_i(\varepsilon)$, hence pr (including p) is outside of $D_i(\varepsilon)$. Since pq contains a point from $\{a_j, b_j\}$, point q is also outside of $D_i(\varepsilon)$. Hence pq traverses $D_i(\varepsilon)$.

We also claim that s cannot traverse $D_i(\varepsilon)$. Suppose, to the contrary, that s traverses $D_i(\varepsilon)$. By Lemma 1, the supporting lines of pq and s meet at an angle less than φ . By the choice of ε , the supporting line of s can intersect the ε -neighborhoods of a_i and a_j only. However, by the choice of b_i , if s contains one point from each of $\{a_i, b_i\}$ and $\{a_j, b_j\}$, then the supporting lines of s and pq intersect in $D_i(\varepsilon)$. This contradicts the fact that the segments of γ do not cross, and proves the claim.

Since s does not traverse $D_i(\varepsilon)$, it has an endpoint v in $D_i(\varepsilon)$. If v is the endpoint of γ , then Case 3 occurs. If v is incident to some other segment of γ , this segment cannot be perfect since every line intersects the ε -neighborhoods of at most two points in A . Hence v is incident to an imperfect segment, and Case 3 occurs. □

We continue with two simple observations about perfect segments containing twins.

Lemma 5. *The supporting lines of any two twins intersect below α .*

Proof. By construction, the supporting line of every twin has positive slope; and $a_i b_i$ has larger slope than $a_j b_j$ if $1 \leq i < j \leq k$. Furthermore, the line $a_i b_i$ has larger slope than the tangent line of the parabola $x \rightarrow x^2$ at a_i , hence a_i lies above the supporting line of $a_j b_j$ for $1 \leq i < j \leq k$. It follows that the supporting lines of segments $a_i b_i$ and $a_j b_j$ intersect below α . □

Lemma 6. *Let pq be a perfect segment of γ that contains a twin $\{a_i, b_i\}$, and let q be the upper (i.e., right) endpoint of pq . Then either q is incident to an imperfect segment of γ or q is an endpoint of γ .*

Proof. Observe that q lies above β . If q is an endpoint of γ , then our proof is complete. Suppose that q is incident to segments pq and qr of γ . By Lemma 5, qr does not contain a twin. By Lemma 3, qr also cannot contain one point from each of two twins. It follows that qr is an imperfect segment of γ , as required. □

Proof (of Lemma 2). Let Γ' be the set of maximal x -monotone chains of perfect segments in γ . Consider a chain $\gamma' \in \Gamma'$. Only the rightmost segment of γ' may contain a twin by Lemma 6. It is possible that the leftmost segment of γ' does not contain a twin, and its left endpoint is incident to some other perfect segment, which is the leftmost segment of some other chain in Γ' .

We charge each segment of γ' to an endpoint of some imperfect segment or to an endpoint of γ as follows. Let pq be a (perfect) segment of γ' . If pq contains a twin, then charge pq to the top vertex of pq , which is the endpoint of an imperfect segment or an endpoint of γ by Lemma 6. Assume now that pq does

not contains a twin, its left endpoint is p , and it contains a point from each of the twins $\{a_i, b_i\}$ and $\{a_j, b_j\}$, with $i < j$. We consider the four cases presented in Lemma 4.

In Case 1, charge pq to p , which is the endpoint of an imperfect segment or an endpoint of γ . In Case 2, charge pq to the left endpoint of the imperfect segment s containing a point of the twin $\{a_i, b_i\}$. In Case 3, charge pq to either an endpoint an imperfect segment or an endpoint of γ located in $D_i(\varepsilon)$. So far, every endpoint of an imperfect segment and every endpoint of γ is charged at most once. Now, consider Case 4 of Lemma 4. In this case, pq is the leftmost segment of γ' . If γ' contains exactly one perfect segment, namely pq , then charge pq to its right endpoint, which is the endpoint of some imperfect segment or the endpoint of γ . If γ' contains at least two perfect segments, then pick an arbitrary perfect segment s , $s \neq pq$, from γ' . Since s is not the leftmost segment of γ' , it has already been charged to some endpoint. Charge pq to the same endpoint as s . Each endpoint of γ and each endpoint of every imperfect segment is now charged at most twice. Since γ and every imperfect segment has two endpoints, we have $s_2 \leq 4(s_0 + s_1) + 4$, as required. \square




Fig. 3. A sketch of our construction S with $k = 8$ twins. (The figure is not to scale). It indicates how 5 consecutive segments of a noncrossing path can cover 4 consecutive twins.

Remark. We do not know whether the lower bound $5(n - 4)/9$ for the number of segments in a minimum noncrossing covering path is tight for the n -element point set S we have constructed. The set S certainly has a covering path with $5n/8 + O(1)$ segments. Such a path is indicated in Fig. 3, where 5 consecutive segments (4 perfect and one imperfect) cover 4 consecutive twins.

3.1 A Two-Colored Version: Proof of Corollary 1

For the upper bound, we proceed as follows. Assume without loss of generality that no two points have the same x -coordinate (after a suitable rotation of the point set, if needed). We have $|B| + |R| = n$, and assume w.l.o.g. that $|B| \leq n/2 \leq |R|$. Cover the red points by an x -monotone spanning path π_R , which is clearly non-crossing. Let $B = B_1 \cup B_2$ be the partition of the blue points induced by π_R into points above and below the red path (remaining points are partitioned arbitrarily). Cover the points in B_1 (above π_R) by an x -monotone covering path: for each consecutive pair of points in the x -order, extend two almost vertical rays that meet far above π_R without crossing π_R . Proceed similarly for covering the points in B_2 (below π_R). Connect the two resulting blue covering paths for B_1 and B_2 by using at most $O(1)$ additional segments.

The number of segments in the red path is $|R| - 1$. The number of segments in the blue path is $2|B| + O(1)$. Consequently, since $|B| \leq n/2$, the two covering paths comprise at most $3n/2 + O(1)$ segments. After sorting the red and blue points along a suitable direction, a pair of mutually non-crossing covering paths as above can be obtained in $O(n)$ time. So the entire procedure takes $O(n \log n)$ time.

For the lower bound, use a red and a blue copy of the point set constructed in the proof of Theorem 2, each with $n/2$ points, so that no three points are collinear. Since covering each copy requires at least $(5n/9 - O(1))/2$ segments in any non-crossing covering path, the resulting n -element point set requires at least $5n/9 - O(1)$ segments in any pair of mutually non-crossing covering paths. \square

4 Computational Complexity

Proof of Theorem 3. We make a reduction from the sorting problem in the algebraic decision tree model of computation. Given n distinct numbers, x_1, \dots, x_n , we map them in $O(n)$ time to n points on the parabola $y = x^2$: $x_i \rightarrow (x_i, x_i^2)$; similar reductions can be found in [21]. Let S denote this n -element point set. Since no 3 points are collinear, any covering path for S has at least $\lceil n/2 \rceil + 1$ vertices. We show below that, given a noncrossing covering path of S with $m = \Omega(n)$ vertices, the points in S can be sorted in left to right order in $O(m)$ time; equivalently, given a noncrossing covering path with m vertices, the $n = O(m)$ input numbers can be sorted in $O(m)$ time. Consequently, the $\Omega(n \log n)$ lower bound is then implied. Thus it suffices to prove the following.

Given a noncrossing covering path γ of S with m vertices, the points in S can be sorted in left to right order in $O(m)$ time.

The boundary of the convex hull of γ is a closed polygonal curve, denoted $\partial\text{conv}(\gamma)$. Melkman's algorithm [18] computes $\partial\text{conv}(\gamma)$ in $O(m)$ time. (See [1] for a thorough review of convex hull algorithms for simple polygons, and [4] for space-efficient variants). Triangulate all faces of the plane graph $\gamma \cup \partial\text{conv}(\gamma)$ within $O(m)$ time [5], and let T denote the triangulation. The parabola $y = x^2$

intersects the boundary of each triangle at most 6 times (at most twice per edge). The intersection points can be sorted in each triangle in $O(1)$ time. So we can trace the parabola $y = x^2$ from triangle to triangle through the entire triangulation, in $O(1)$ time per triangle, thus in $O(m)$ time overall. Since all points of S are on the parabola, one can report the sorted order of the points within the same time. \square

5 Conclusion

We conclude with a few (new or previously posed) questions and some remarks.

1. It seems unlikely that every point set with no three collinear points admits a covering path with $n/2 + O(1)$ segments. Can a lower bound of the form $f(n) = n/2 + \omega(1)$ be established?
2. Recently, Gerbner and Keszegh [14] have shown that $g(n) \leq cn$ for some positive constant $c < 1$, thereby disproving the conjectured relation $g(n) = n(1+o(1))$. However, a sizable gap to the lower bound in Theorem 2 remains.
3. Let $p(n)$ denote the maximum integer such that any set of n points in the plane has a perfect subset of size $p(n)$. As noticed by Welzl [9,23], $p(n) = \Omega(\log n)$ immediately follows from the theorem of Erdős and Szekeres [10]. Any improvement in this lower bound would lead to a better upper bound in Theorem 1, and thus to a smaller gap versus the $n/2$ lower bound. It is a challenging question whether Welzl's lower bound $p(n) = \Omega(\log n)$ can be improved; see also [9].
4. It is known that the minimum-link covering path problem is NP-complete for planar paths whose segments are unrestricted in orientation [2,16]. It is also NP-complete for axis-parallel paths in \mathbb{R}^{10} , as shown in [15]. Is the minimum-link covering path problem still NP-complete for axis-aligned paths in \mathbb{R}^d for $2 \leq d \leq 9$? It is known [3] that a minimum-link axis-aligned covering path in the plane can be approximated with ratio 2. Can the approximation ratio of 2 be reduced?
5. Is the minimum-link covering path problem still NP-complete for points in general position and arbitrary oriented paths?
6. Is the minimum-link covering path problem still NP-complete for points in general position and arbitrary oriented non-crossing paths?
7. Given n points (n even), is it possible to compute a non-crossing perfect matching in $O(n)$ time? Observe that such a matching can be computed in $O(n \log n)$ time by sorting the points along some direction. The same upper bound $O(n \log n)$ holds for non-crossing covering paths and non-crossing spanning paths, and this is asymptotically optimal by Theorem 3. Observe finally that a non-crossing spanning tree can be computed in $O(n)$ time: indeed, just take a star rooted at an arbitrary point in the set.

Acknowledgment. The authors are grateful to an anonymous reviewer for the running time improvement in Theorem 1, and to the members of the MIT-Tufts Computational Geometry Research Group for stimulating discussions.

References

1. Aloupis, G.: A history of linear-time convex hull algorithms for simple polygons, <http://cgm.cs.mcgill.ca/~athens/cs601/>
2. Arkin, E.M., Mitchell, J.S.B., Piatko, C.D.: Minimum-link watchman tours. *Inf. Proc. Lett.* 86, 203–207 (2003)
3. Bereg, S., Bose, P., Dumitrescu, A., Hurtado, F., Valtr, P.: Traversing a set of points with a minimum number of turns. *Discrete Comput. Geom.* 41(4), 513–532 (2009)
4. Brönnimann, H., Chan, T.M.: Space-efficient algorithms for computing the convex hull of a simple polygonal line in linear time. *Comput. Geom.* 34, 75–82 (2006)
5. Chazelle, B.: Triangulating a simple polygon in linear time. *Discrete Comput. Geom.* 6, 485–524 (1991)
6. Chvátal, V., Klineček, G.: Finding largest convex subsets. *Congressus Numerantium* 29, 453–460 (1980)
7. Collins, M.J.: Covering a set of points with a minimum number of turns. *Internat. J. Comput. Geom. Appl.* 14(1-2), 105–114 (2004)
8. Collins, M.J., Moret, M.E.: Improved lower bounds for the link length of rectilinear spanning paths in grids. *Inf. Proc. Lett.* 68, 317–319 (1998)
9. Demaine, E.D., O’Rourke, J.: Open problems from CCCG 2010. In: *Proc. 23rd Canadian Conf. Comput. Geom.*, Toronto, ON, pp. 153–156 (2011)
10. Erdős, P., Szekeres, G.: A combinatorial problem in geometry. *Compositio Mathematica* 2, 463–470 (1935)
11. Erdős, P., Szekeres, G.: On some extremum problems in elementary geometry. *Annales Univ. Sci. Budapest* 3-4, 53–62 (1960-1961)
12. Estivill-Castro, V., Heednacram, A., Suraweera, F.: NP-completeness and FPT results for rectilinear covering problems. *Journal of Universal Computer Science* 15, 622–652 (2010)
13. Fulek, R., Keszegh, B., Morić, F., Uljarević, I.: On polygons excluding point sets. In: *Proc. 22nd Canadian Conf. Comput. Geom.*, Winnipeg, MB, pp. 273–276 (2010)
14. Gerbner, D., Keszegh, B.: Personal Communication (July 2012)
15. Jiang, M.: On Covering Points with Minimum Turns. In: Snoeyink, J., Lu, P., Su, K., Wang, L. (eds.) *FAW-AAIM 2012*. LNCS, vol. 7285, pp. 58–69. Springer, Heidelberg (2012)
16. Kranakis, E., Krizanc, D., Meertens, L.: Link length of rectilinear hamiltonian tours in grids. *Ars Combinatoria* 38, 177–192 (1994)
17. Matoušek, J.: *Lectures on Discrete Geometry*. Springer, New York (2002)
18. Melkman, A.: On-line construction of the convex hull of a simple polygon. *Inf. Proc. Lett.* 25(1), 11–12 (1987)
19. Mitchell, J.S.B.: Geometric shortest paths and network optimization. In: Sack, J.R., Urrutia, J. (eds.) *Handbook of Computational Geometry*, ch. 15, pp. 633–701. Elsevier (2000)
20. Morić, F.: Separating and covering points in the plane. Communication at the Open Problem Session of the *22nd Canadian Conf. Comput. Geom.*, Winnipeg, MB (2010)
21. Preparata, F.P., Shamos, M.I.: *Computational Geometry: An Introduction*. Springer, New York (1985)
22. Stein, C., Wagner, D.P.: Approximation Algorithms for the Minimum Bends Traveling Salesman Problem. In: Aardal, K., Gerards, B. (eds.) *IPCO 2001*. LNCS, vol. 2081, pp. 406–421. Springer, Heidelberg (2001)
23. Welzl, E.: Separating and covering points in the plane. In: Communication at the *9th Gremo’s Workshop on Open Problems*, Switzerland (February 2011)