

Covering Paths for Planar Point Sets

Adrian Dumitrescu* Dániel Gerbner† Balázs Keszegh‡ Csaba D. Tóth§

February 28, 2013

Abstract

Given n points in the plane, a *covering path* is a polygonal path that visits all the points. If no three points are collinear, every covering path requires at least $n/2$ segments, and $n - 1$ straight line segments obviously suffice even if the covering path is required to be noncrossing. We show that every set of n points in the plane admits a (possibly self-crossing) covering path consisting of $n/2 + O(n/\log n)$ straight line segments. If the path is required to be noncrossing, we prove that $(1 - \varepsilon)n$ straight line segments suffice for a small constant $\varepsilon > 0$, and we exhibit n -element point sets that require at least $5n/9 - O(1)$ segments in every such path. Further, the analogous question for noncrossing *covering trees* is considered and similar bounds are obtained. Finally, it is shown that computing a noncrossing 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 the plane. 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 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 in S are collinear, every covering path consists of at least $\lceil n/2 \rceil$ segments. A *minimum-link* covering path for S is one with the smallest number of segments (links). A point set is said to be in *general position* if no three points are collinear.

We study the following two questions concerning covering paths posed by Morić [20, 21] as a generalization of the well-known puzzle of linking 9 dots in a 3×3 grid with a polygonal path having only 4 segments [16]. Another problem which leads to these questions is separating red from blue points [13].

*Computer Science, University of Wisconsin–Milwaukee, USA. Email: dumitres@uwm.edu. Research was supported by the NSF grant DMS-1001667.

†Alfréd Rényi Institute of Mathematics, Hungarian Academy of Sciences, Budapest, Hungary. Email: gerbner.daniel@renyi.mta.hu. Research was supported by OTKA, grant NK 78439 and by OTKA under EU-ROGIGA project GraDR 10-EuroGIGA-OP-003.

‡Alfréd Rényi Institute of Mathematics, Hungarian Academy of Sciences, Budapest, Hungary. Email: keszegh.balazs@renyi.mta.hu. Research was supported by OTKA, grant NK 78439 and by OTKA under EU-ROGIGA project GraDR 10-EuroGIGA-OP-003.

§Department of Mathematics, California State University, Northridge, USA and University of Calgary, Canada. Email: cdtoth@ucalgary.ca. Research was supported by the NSERC grant RGPIN 35586, the NSF grant CCF-0830734, and the Fields Institute for Research in Mathematical Sciences, Toronto, Canada.

1. What is the minimum number, $f(n)$, such that every set of n points in the plane can be covered by a (possibly self-intersecting) polygonal path with $f(n)$ segments?
2. What is the minimum number, $g(n)$, such that every set of n points in the plane can be covered by a *noncrossing* polygonal path with $g(n)$ segments?

If no three points are collinear, then each segment of a covering path contains at most two points, thus $\lceil n/2 \rceil$ is a trivial lower bound for both $f(n)$ and $g(n)$. Morić conjectured that the answer to the first problem is $n(1/2 + o(1))$ while the answer to the second is $n(1 - o(1))$. We confirm his first conjecture (Theorem 1) but refute the second (Theorem 2)¹. A consideration of these questions in retrospect appears in [9].

Theorem 1. *Every set of n points in the plane admits a (possibly self-crossing) covering path consisting of $n/2 + O(n/\log n)$ line segments. Consequently, $\lceil n/2 \rceil \leq f(n) \leq n/2 + O(n/\log n)$. A covering path with $n/2 + O(n/\log n)$ segments can be computed in $O(n^{1+\varepsilon})$ time, for every $\varepsilon > 0$.*

As expected, the noncrossing property is much harder to deal with. Every set of n points in the plane trivially admits a noncrossing 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, if no three points are collinear. We provide the first nontrivial upper and lower bounds for $g(n)$, in particular disproving the conjectured relation $g(n) = n(1 - o(1))$.

Theorem 2. *Every set of n points in the plane admits a noncrossing covering path with at most $\lceil (1 - 1/601080391)n \rceil - 1$ segments. Consequently, $g(n) \leq \lceil (1 - 1/601080391)n \rceil - 1$. A noncrossing covering path with at most this many segments can be computed in $O(n \log n)$ time.*

Theorem 3. *There exist n -element point sets that require at least $(5n - 4)/9$ segments in any noncrossing covering path. Consequently, $g(n) \geq (5n - 4)/9$.*

In the proof of Theorem 2, we construct a noncrossing covering path that can easily be extended to a noncrossing covering cycle by adding one Steiner point (and two segments).

Corollary 1. *Every set of $n \geq 2$ points in the plane admits a noncrossing covering cycle with at most $\lceil (1 - 1/601080391)n \rceil + 1$ segments. A noncrossing covering cycle of at most this many segments can be computed in $O(n \log n)$ time.*

Covering trees. For covering a finite point set in the plane, certain types of geometric graphs other than paths may also be practical. A noncrossing path or tree, for example, are equally useful for separating a red and blue set of points [13], which is one of the motivating problems. A *covering tree* for a planar point set S is a tree drawn in the plane with straight-line edges such that every point in S lies at a vertex or on an edge of the tree. The lower and upper bounds $\lceil n/2 \rceil \leq f(n) \leq n/2 + O(n/\log n)$ of Theorem 1 trivially carry over for the number of edges of covering trees (with possible edge crossings).

Let $t(n)$ be the minimum integer such that every set of n points in the plane admits a *noncrossing covering tree* with $t(n)$ straight-line edges. Since every path is a tree, we have $t(n) \leq g(n) \leq \lceil (1 - 1/601080391)n \rceil$ from Theorem 2. However, a noncrossing covering tree is significantly easier to obtain than a noncrossing covering path. By simplifying the proof of Theorem 2, we derive a stronger upper bound for covering trees.

¹The first item was observed by the current authors during the Canadian Conference CCCG 2010 and was also communicated to the authors of [9].

Theorem 4. *Every set of n points in the plane admits a noncrossing covering tree with at most $\lfloor 5n/6 \rfloor$ edges. Consequently, $t(n) \leq \lfloor 5n/6 \rfloor$. A covering tree with at most $\lfloor 5n/6 \rfloor$ edges can be computed in $O(n \log n)$ time.*

By modifying the lower bound analysis in the proof of Theorem 3, we show that the same point set used there yields a slightly weaker lower bound for noncrossing covering trees.

Theorem 5. *There exist n -element point sets in the plane that require at least $(9n - 4)/17$ edges in any noncrossing covering tree. Consequently, $t(n) \geq (9n - 4)/17$.*

Instead of minimizing the number of edges in a covering tree, one can try to minimize the number of line segments, where each segment is either a single edge or a chain of several collinear edges of the tree. Let $s(n)$ be the minimum integer such that every set of n points in the plane admits a noncrossing covering tree with $s(n)$ line segments. By definition, we trivially have $s(n) \leq t(n)$. In addition, we determine an exact formula for $s(n)$:

Proposition 1. *We have*

$$s(n) = \begin{cases} n - 1 & \text{if } n = 2, 3, 4 \\ \lceil n/2 \rceil & \text{if } n \geq 5. \end{cases}$$

Bicolored variants. 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 noncrossing* if each of π_R and π_B is noncrossing, and moreover, π_R and π_B do not cross (intersect) each other. A natural extension of the monochromatic noncrossing covering path problem is: What is the minimum number $j(n)$ such that every bicolored set of n points in the plane can be covered by two monochromatic mutually noncrossing polygonal paths with $j(n)$ segments in total? Using the construction in the proof of Theorem 3 we obtain the following corollary.

Corollary 2. *Given a bicolored set of n points, there are two mutually noncrossing 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 noncrossing covering paths. Consequently, $5n/9 - O(1) \leq j(n) \leq 3n/2 + O(1)$.*

Similarly, two covering trees τ_R and τ_B , one for the red and one for the blue points, are *mutually noncrossing* if each is noncrossing and τ_R and τ_B do not cross each other. The analogous question is in this case: What is the minimum number $k(n)$ such that every bicolored set of n points in the plane can be covered by two monochromatic mutually noncrossing polygonal trees with $k(n)$ edges in total? The construction in the proof of Theorem 5 yields the following corollary.

Corollary 3. *Given a bicolored set of n points, there are two mutually noncrossing covering trees with a total of at most n edges. Such a pair of trees can be computed in $O(n \log n)$ time. On the other hand, there exist bicolored sets that require at least $9n/17 - O(1)$ edges in any pair of mutually noncrossing covering trees. Consequently, $9n/17 - O(1) \leq k(n) \leq n$.*

Computational complexity. We establish an $\Omega(n \log n)$ lower bound for computing a noncrossing covering path for a set of n points in the plane.

Theorem 6. *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.*

On the other hand, a noncrossing covering tree for n points can be easily computed in $O(n)$ time; see also Section 8.

Related previous results. Given a set of n points in the plane, 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 [23] 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, 7, 8, 15]. 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 [14]. Other variants of Euclidean TSP can be found in a 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 (strictly) convex position and all points of X lie above (below, respectively) the line connecting the leftmost and the rightmost point of X . Similarly, caps and cups can be defined for arbitrary points (with allowed collinearities), with X being in *weakly* convex position. By slightly abusing notation, we use the same terminology when referring to them, and distinguish them based on the underlying point sets.

According to a classical result of Erdős and Szekeres [10], every 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 -cap or a k -cup. 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 neither k -cups nor l -caps. While Erdős and Szekeres originally proved the above results for points in general position, their arguments go through verbatim for arbitrary point sets (with allowed collinearities), and the same quantitative bounds hold for the resulting caps or cups (now in weakly convex position).

Following the terminology coined by Welzl [24], a set S of n points in the plane 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 (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 every 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. The total number of segments in this path is $n/2 + O(n/\log n)$, as required. Chvátal and Kłincsek [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 Non-Crossing Covering Paths: Upper Bound

It is easy to see that for every set of n points in the plane, there is a noncrossing covering path with at most $n - 1$ segments. For example, an x -monotone spanning path for n points has $n - 1$ segments, no two crossing edges, and no Steiner points either. In this section, we prove Theorem 2 and show that $(1 - \varepsilon)n$ segments suffice for some small constant $\varepsilon > 0$. In the proof of Theorem 2, however, we still use the trivial upper bound $n - 1$ for several subsets of points with the additional constraint that the two endpoints of the path are two given points on the boundary of a convex region containing the points in its interior (Lemma 1).

Lemma 1. *Let X be a set of n points in the interior of a convex region C , and let a, b be two points on the boundary ∂C of C . Then $X \cup \{a, b\}$ admits a noncrossing covering path with $|X| + 1$ segments such that its two endpoints are a and b , and its relative interior lies in the interior of C . Such a covering path can be constructed in $O(n \log n)$ time.*

We include the easy proof for completeness (a similar lemma was also an essential tool in [13]).

Proof. Let ℓ_a and ℓ_b be tangent lines at C incident to a and b , respectively. If ℓ_a and ℓ_b are not parallel, then let $O = \ell_a \cap \ell_b$; otherwise let O be a point at infinity corresponding to the direction of the two parallel lines ℓ_a and ℓ_b . Sort the points in X in the order in which they are encountered by a rotating sweep line from ℓ_a to ℓ_b around O (with ties broken arbitrarily). Let γ be the polygonal path that starts at a , visits the points in X in the above sweep order, and ends at b . The edges of γ are pairwise noncrossing, since they lie in interior-disjoint wedges centered at O . By construction, γ lies in the convex hull of $X \cup \{a, b\} \subseteq \text{int}(C) \cup \{a, b\}$, hence the relative interior of γ lies in the interior of C , as required. \square

Before proving Theorem 2, we show how to reduce the trivial bound $n - 1$ on the size of noncrossing covering paths by an arbitrarily large constant, provided that the number of points n is sufficiently large.

Lemma 2. *Let S be a set of n points in the plane that contains a cap or cup of even size k . Then S admits a noncrossing covering path γ with at most $n + 1 - \lfloor k/6 \rfloor$ segments. Furthermore, if S lies in the interior of a vertical strip H bounded by two vertical lines, h_1 and h_2 , then we may require that the two endpoints of γ lie on h_1 and h_2 , respectively, and the relative interior of γ lie in the interior of H .*

Proof. Let $H, S \subset H$, be a vertical strip bounded by two vertical lines, h_1 and h_2 , from the left and right, respectively. Assume that S contains a *cap* of size k (the case of a *cup* is analogous). We construct a noncrossing covering path γ with at most $n + 1 - \lfloor k/6 \rfloor$ segments. We may assume that k is a multiple of 6 (by decreasing k , if necessary, without changing the value of $\lfloor k/6 \rfloor$).

Let $P = \{p_1, p_2, \dots, p_k\} \subset S$ be a cap of size of k , labeled in increasing order of x -coordinates. Note that P admits a covering path $\gamma_0 = (q_1, q_2, \dots, q_{k/2+1})$ of $k/2$ segments, where each segment of γ_0 contains two consecutive points of the cap. We may assume (by extending or shortening γ_0 if necessary) that the endpoints of γ_0 are on the boundary of the vertical strip H , that is, $q_1 \in h_1$ and $q_{k/2+1} \in h_2$. Let $q_0 \in h_1$ and $q_{k/2+2} \in h_2$ be arbitrary points vertically below the endpoints of γ_0 . With this notation, the polygonal path $(q_0q_1) \cup \gamma_0 \cup (q_{k/2+1}q_{k/2+2})$ is a convex arc. Let $s_{k/2+1} \in h_2$ be an arbitrary Steiner point above the right endpoint of γ_0 . The two endpoints of our final covering path for S will be $q_0 \in h_1$ and $s_{k/2+1} \in h_2$.

In the remainder of the proof, we first construct a noncrossing covering path for S with $n + 1$ segments from q_0 to $s_{k/2+1}$ and then modify this construction to “save” $\lfloor k/6 \rfloor$ segments.




Figure 1: A set with $n = 36$ points. A covering path γ_0 for a cap of $k = 18$ points (bold). The region above γ_0 is decomposed into 9 convex regions by vertical rays r_i , $i = 2, \dots, 9$, each passing through a Steiner point s_i . We obtain a covering path with $n + 1 = 37$ segments for all points by concatenating three covering paths.

Preliminary approach. The path γ_0 decomposes the vertical strip H into two regions: a convex region below and a nonconvex region above, with $k/2 - 1$ reflex vertices at $q_2, \dots, q_{k/2}$. Decompose the nonconvex region above γ_0 into $k/2$ convex regions by upward vertical rays r_i emitted by q_i , for $i = 2, \dots, q_{k/2}$; and place an arbitrary Steiner point s_i on the ray r_i .

We construct a noncrossing covering path for S as a concatenation of the following three paths: The first path is a covering path for the points in S lying strictly below γ_0 , from the Steiner point $q_0 \in h_1$ to $p_k \in \gamma_0$, obtained by Lemma 1. The second path is the part of γ_0 from p_k to p_1 . The third path consists of $k/2$ covering subpaths for the points in S lying above γ_0 ; the first of these subpaths runs from p_1 to s_2 , and the others run from s_i to s_{i+1} , for $i = 1, \dots, k/2$, as obtained by Lemma 1.

The resulting path visits all points of S , since the second part visits all points along γ_0 , and the convex regions jointly contain all points below or above γ_0 . The $k/2 + 2$ parts of the covering path

are pairwise noncrossing, since they lie either on γ_0 or in pairwise interior-disjoint regions whose interiors are also disjoint from γ_0 . The second part of the covering path (the part along γ_0) covers k points with $k/2$ segments, and the remaining $k/2 + 1$ parts each require one more segment than the number of points covered. Hence, the total number of segments is $n + 1 - (k/2) + (k/2 + 1) = n + 1$.

Modified construction. We now modify the above construction and “save” $\lfloor k/6 \rfloor$ segments. The savings come from the following two ideas. (1) It is not necessary to decompose the entire region above γ_0 into convex pieces. If a region above γ_0 contains all points in S above γ_0 , and has fewer than $k/2$ reflex vertices, then we can decompose this region into fewer than $k/2$ convex pieces, using fewer than $k/2$ Steiner points and thus reducing the size of the resulting covering path. (2) If a ray emitted by a reflex vertex passes through a point of S and decomposes the reflex angle into two convex angles, then a Steiner point can be replaced by a point of S , which saves one segment in the resulting covering path.




Figure 2: Local modifications for γ_0 . Case 1: $S \cap Q_9 \neq \emptyset$, we choose a ray r_9 emitted by q_9 that passes through a point in $S \cap Q_9$. Case 2: $S \cap Q_3 = \emptyset$, we construct $\text{conv}(X_3)$, and replace the arc $(q_2, q_3, q_4) \subset \gamma_0$ by $(q_2, q'_2, q'_4, q_4) \subset \gamma_1$. Similarly, $S \cap Q_6 = \emptyset$, we construct $\text{conv}(X_6)$, and replace the arc $(q_5, q_6, q_7) \subset \gamma_0$ by $(q_5, q'_5, q'_7, q_7) \subset \gamma_1$. Note that the triangles $\Delta(c_3, q'_2, q'_4)$ and $\Delta(c_6, q'_5, q'_7)$ contain no points from S .

We show that one of the two ideas is always applicable locally. Specifically, we modify γ_0 by replacing some of the arcs $(q_{i-1}, q_i, q_{i+1}) \subset \gamma_0$, where i is a multiple of 3, by different arcs and obtain a new convex polygonal arc γ_1 . The path γ_1 retains the property that every segment contains two points from S , however it may consist of fewer segments than γ_0 . We keep the modifications “local” in the sense that a modification in the neighborhood of a vertex q_i , where $i > 0$ is a multiple of 3, is carried out independently of modifications at all other vertices q_j , where $j > 0$ is a multiple of 3. Even though γ_0 may be modified in the neighborhoods of vertices q_i and q_{i+3} , where $i > 0$ is a multiple of 3, the intermediate edge $q_{i+1}q_{i+2}$ of γ_0 will be preserved: it will be either an edge of γ_1 or contained in a longer edge of γ_1 .

For every i , where i is a positive multiple of 3, let Q_i be the region above *both* lines $q_{i-1}q_i$ and q_iq_{i+1} , and between the vertical rays r_{i-1} and r_{i+1} . We distinguish two cases (refer to Fig. 2):

Case 1: $S \cap Q_i \neq \emptyset$. In this case, the polygonal arc (q_{i-1}, q_i, q_{i+1}) of γ_0 is not modified but the ray r_i is defined differently. In addition, the Steiner points $s_{i-1} \in r_{i-1}$ and $s_{i+1} \in r_{i+1}$ are defined differently. Pick an arbitrary point $s_i \in S \cap Q_i$, and let r_i be the ray emitted by q_i and passing through s_i . Let r_{i-1} and r_{i+1} be vertical rays emitted by q_{i-1} and q_{i+1} , respectively, like before. Decompose the region above (q_{i-1}, q_i, q_{i+1}) by the two vertical rays r_{i-1} and r_{i+1} , and then by the (possibly nonvertical) ray r_i . Choose Steiner points $s_{i-1} \in r_{i-1}$ and $s_{i+1} \in r_{i+1}$, such that they each lie on the common boundary of two consecutive regions of the decomposition.

Case 2: $S \cap Q_i = \emptyset$. In this case we modify the polygonal arc (q_{i-1}, q_i, q_{i+1}) of γ_0 . Let c_i be the intersection point of lines $q_{i-2}q_{i-1}$ and $q_{i+1}q_{i+2}$ (possibly $q_{i-2} = q_0$ or $q_{i+2} = q_{k/2+2}$). Let S_i denote the set of points of S in the interior of the triangle $\Delta(q_{i-1}c_iq_{i+1})$; and let $X_i = S_i \cup \{q_{i-1}, q_{i+1}\}$. Since Q_i is empty and each of the segments $q_{i-1}q_i$ and q_iq_{i+1} contains two points of S , the convex hull of X_i has at least four vertices, i.e., $\text{conv}(X_i)$ is not a triangle. Let $a_i b_i$ be an arbitrary edge of $\text{conv}(X_i)$, that is incident to neither q_{i-1} nor q_{i+1} . The line $a_i b_i$ intersects the sides $q_{i-1}c_i$ and $q_{i+1}c_i$ of the triangle $\Delta(q_{i-1}c_iq_{i+1})$. Denote the intersection points by $q'_{i-1} \in q_{i-1}c_i$ and $q'_{i+1} \in q_{i+1}c_i$. Replace (q_{i-1}, q_i, q_{i+1}) by $(q_{i-1}, q'_{i-1}, q'_{i+1}, q_{i+1})$ to obtain γ_1 .

The segment $q'_{i-1}q'_{i+1}$ contains points $a_i, b_i \in S$. Notice that $q_{i-1}q'_{i-1}$ is collinear with $q_{i-2}q_{i-1}$; and similarly $q'_{i+1}q_{i+1}$ is collinear with $q_{i+1}q_{i+2}$. Therefore, q_{i-1} and q_{i+1} are not vertices of γ_1 . Let r_i be a vertical ray emitted by c_i , and pick an arbitrary Steiner point $s_i \in r_i$. Decompose the region above $(q_{i-1}, q'_{i-1}, q'_{i+1}, q_{i+1}) \subset \gamma_1$ by the polygonal arc (q_{i-1}, c_i, q_{i+1}) and the vertical ray r_i . Since the triangle $\Delta(c_i, q'_{i-1}, q'_{i+1})$ contains no point from S , the two convex regions adjacent to r_i (only one region in the extremal case $i = k/2$) contain all the points of S lying above $(q_{i-1}, q'_{i-1}, q'_{i+1}, q_{i+1})$.




Figure 3: The set of $n = 36$ points from Fig. 1. A covering path γ_1 for a cap of 14 points (bold). A dashed noncrossing path below γ_1 , and a dashed noncrossing path above γ_1 cover all remaining points in S . The resulting covering path has $n - 2 = 34$ segments.

After the local modifications, we proceed analogously to our initial approach. Construct a noncrossing covering path for S as a concatenation of the following three paths: A trivial path for the points strictly below γ_1 from the Steiner point $q_0 \in h_1$ to the rightmost vertex on γ_1 (Lemma 1); followed by part of γ_1 from the rightmost to the leftmost point on γ_1 ; the third path visits all points in the convex regions above γ_1 , passing through the points $s_i \in r_i$ between consecutive regions.

The resulting path visits all points of S lying below, on, and above γ_1 (in this order). We use $\lfloor k/6 \rfloor$ fewer segments than in our initial construction, since each local modification saves one segment: In Case 1, we use a point $s_i \in S$ instead of a Steiner point. In Case 2, we decrease the number of segments along γ_1 by one, and decrease the number of relevant convex regions above γ_1 by two. This concludes the proof of Lemma 2. \square

Proof of Theorem 2. Let S be a set of $n + 1$ points in the plane, no three of which are collinear. Assume, by rotating the point set if necessary, that no two points have the same x -coordinate.

Lay out a raster of vertical lines in the plane such that there are exactly $m = \binom{32}{16} + 1 = 601,080,391$ points between consecutive lines; no points on the lines or to the left of the leftmost line; and less than m points to the right of the rightmost line. By the result of Erdős and Szekeres [10], there is a cap or cup of 18 points between any two raster lines. By Lemma 2 (for $n = m$ and $k = 18$), the m points between consecutive raster lines admit a noncrossing covering path with $m + 1 - 3 = m - 2$ segments such that the two endpoints of the path are Steiner points on the two raster lines, and the relative interior of the path lies strictly between the raster lines. The $n_0 < m$ points to the right of the rightmost raster line can be covered by an x -monotone path with n_0 segments starting at the (unique) Steiner point on that line assigned by the previous group of points.

The noncrossing covering paths between consecutive raster lines can be joined into a single noncrossing covering path for S by adding one vertical segment on each raster line except for the first and the last one, as depicted in Fig. 4. The total number of segments is $n - \lfloor n/m \rfloor - 1 \leq \lceil (1 - 1/601080391)n \rceil - 1$, as claimed. \square

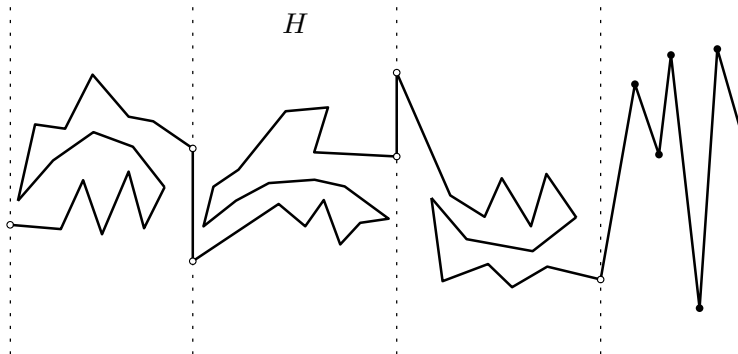


Figure 4: Joining noncrossing covering paths in the last step.

Proof of Corollary 1. In the proof of Theorem 2, we constructed a noncrossing covering path γ for S such that the two endpoints of γ are leftmost and the rightmost vertices of γ . Hence γ can be augmented to a noncrossing covering *cycle* by adding a new vertex of sufficiently large y -coordinate, and thus proving Corollary 1. \square

4 Noncrossing Covering Paths: Lower Bound

Proof of Theorem 3 (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 $(5n - 4)/9$ segments. The lower bound is based on a charging scheme: we distinguish *perfect* and *imperfect* segments in γ , containing 2 and

fewer 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 $(5n - 4)/9$ follows. We continue with the details.

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




Figure 5: 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 $2 \arcsin(\delta/\varepsilon)$.

Lemma 3. *Let $\varphi \in (0, \frac{\pi}{2})$ be an angle. For every $\varepsilon > 0$, there exists $\delta \in (0, \varepsilon)$ 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. 5. 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 $2 \arcsin(\delta/\varepsilon)$. For every $\varepsilon > 0$, $\exists \delta \in (0, \varepsilon)$ such that $2 \arcsin(\delta/\varepsilon) < \varphi$. \square

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 a suitable $\delta > 0$ and $i = 1, \dots, k$. The pairs $\{a_i, b_i\}$ are called *twins*. The value of $\delta > 0$ is specified in the next paragraph. See Fig. 6 for a sketch of the construction.

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 a line intersects two disks in $\{D_1(\varepsilon), \dots, D_k(\varepsilon)\}$, then it meets the line passing through the centers of the two disks at an angle less than $\varphi/2$. It follows that any 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 3, 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 ready to define $\delta > 0$; let $\delta = \min\{\delta_i : i = 0, 1, \dots, k - 1\}$.




Figure 6: A sketch of our construction S with $k = 8$ twins. (The figure is not to scale.) It is also indicated how 5 consecutive segments of a noncrossing path can cover 4 consecutive twins.

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 such that its 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 such that its 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 $D_i(\varepsilon)$.

Write $B = \{b_1, \dots, b_k\}$. We also ensure in each iteration that in the set $S = A \cup B$, (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 covering path, we may ignore the perturbation.

Let γ be a noncrossing covering path for S . By perturbing the vertices of γ if necessary, 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 , respectively, the number of segments in γ that contain 0, 1, and 2 points from S . We establish the following inequality.

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

Before the proof of Lemma 4, we show that it immediately implies Theorem 3.

Proof of Theorem 3. Counting the number of points incident to the segments, we have $n = s_1 + 2s_2$. The number of segments in γ is $s_0 + s_1 + s_2$. This number can be bounded from below by using

Lemma 4 as follows.

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

as claimed □

For the proof of Lemma 4, we introduce a charging scheme: each perfect segment is charged to either an endpoint of an imperfect segment, or one of the two endpoints of γ such that every 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 a 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 noncrossing path covering S . 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 5. *Let pq be a perfect segment in γ that contains one point from each of the twins $\{a_i, b_i\}$ and $\{a_j, b_j\}$, where $i < j$. 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 . □

Lemma 6. *Let pq be a perfect segment of γ that contains one point from each of the twins $\{a_i, b_i\}$ and $\{a_j, b_j\}$, where $i < j$. 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 p 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 .




Figure 7: The four cases in Lemma 6 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 . The disk $D_i(\varepsilon)$ is shaded.

Proof. If p is incident to an imperfect segment of γ , or p is an endpoint of γ , then Case 1 occurs. Assume therefore 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 shall 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, every 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 $\{a_j, b_j\}$, $i < j$, 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 3, 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. \square

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

Lemma 7. *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 . \square

Lemma 8. *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 7, qr does not contain a twin. By Lemma 5, qr cannot contain one point from each of two twins, either, since then its endpoints would lie below β . It follows that qr is an imperfect segment of γ , as required. \square

Proof of Lemma 4. Let Γ' be the set of maximal x -monotone chains of perfect segments in γ . Consider a chain $\gamma' \in \Gamma'$. By Lemma 8, only the rightmost segment of γ' may contain a twin. It is possible that the leftmost segment of γ' contains one point from each of two twins, and the left endpoint of γ' is incident to another perfect segment, which is the left endpoints of another x -monotone chain in Γ' .

Let pq be a perfect segment of γ , and part of an x -monotone chain $\gamma' \in \Gamma'$. We charge pq to a point $\sigma(pq)$ that is either an endpoint of some imperfect segment or an endpoint of γ .

The point $\sigma(pq)$ is defined as follows. 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 8. Assume now that pq does not contain 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 6.

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 6. 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' , the point $\sigma(s)$ is already defined, and we let $\sigma(pq) = \sigma(s)$. This completes the definition of $\sigma(pq)$.

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

Remark. We do not know whether the lower bound $(5n - 4)/9$ 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. 6, where 5 consecutive segments (4 perfect and one imperfect) cover 4 consecutive twins.

5 Noncrossing Covering Trees

An upper bound $t(n) \leq \lceil (1 - 1/601080391)n \rceil$ for noncrossing covering trees follows from Theorem 2. However, the argument can be greatly simplified while also improving the bound.

Proof of Theorem 4. Any set of 7 points with distinct x -coordinates contains a cap or cup of 4 points, say, a, b, c, d , from left to right. The 4 points a, b, c, d , of a cap or cup admit a covering path with 2 segments, i.e., a 2-edge star centered at the intersection point, say v , of the lines through ab and cd , respectively. Augment this 2-edge star covering 4 points to a 5-edge star centered at v and covering all 7 points. The star is contained in the vertical strip bounded by vertical lines incident to the leftmost and the rightmost point, respectively.

We may assume, by rotating the point set if necessary, that no two points have the same x -coordinate. Let p_1, p_2, \dots, p_n be the points in S listed in left to right order. Decompose S into groups of 7 by drawing vertical lines incident to p_{7+6i} , $i = 0, 1, \dots$. Any two consecutive groups in this decomposition share a point (the last point in group i is also the first point in group $i+1$). Thus the stars covering the groups (using 5 edges per group) are already connected in a tree covering all points, that yields the claimed bound. \square

We prove a lower bound for $t(n)$ by analyzing noncrossing covering trees of the point set $S = \{a_1, \dots, a_k, b_1, \dots, b_k\}$, $n = 2k$, defined in Section 4 above. Let τ be a noncrossing covering tree for S . By perturbing the vertices of τ if necessary, we may assume that every point in S lies in the relative interior of a segment of τ . Let s_0, s_1 and s_2 , respectively, denote the number of segments in τ that contain 0, 1, and 2 points from S ; hence $n = s_1 + 2s_2$. We establish the following weaker version of Lemma 4.

Lemma 9. $s_2 \leq 8(s_0 + s_1) + 4$

Before the proof of Lemma 9, we show that it directly implies Theorem 5.

Proof of Theorem 5. The total number of segments in τ is

$$\begin{aligned} s_0 + s_1 + s_2 &= \frac{8(s_0 + s_1) + 4 + 9s_0 + 9s_1 - 4}{17} + s_2 \\ &\geq \frac{s_2 + 9s_0 + 9s_1 - 4}{17} + s_2 \\ &\geq \frac{9(s_1 + 2s_2) - 4}{17} = \frac{9n - 4}{17}, \end{aligned}$$

where we used Lemma 9 and the fact that the total number of points is $n = s_1 + 2s_2$. \square

For the proof of Lemma 9, we set up a charging scheme, similar to the proof of Lemma 4. Lemma 8 continues to hold for τ in place of a noncrossing covering path γ if we replace the “endpoints of γ ” by the “leaves of τ .” While the path γ has exactly two endpoints, the tree τ may have arbitrarily many leaves. Therefore, the charging scheme has to be modified so that no perfect segment is charged to the leaves of τ .

Proof of Lemma 9. Since S is in general position, no three perfect segments have a common endpoint. Therefore, the perfect segments of τ form disjoint paths. Let Γ be the set of maximal chains of perfect segments; and let Γ_x denote the set of maximal x -monotone chains of perfect segments in τ .

Choose an arbitrary vertex r_0 in τ as a *root*, and direct all edges of τ towards r_0 . Every chain in Γ is incident to either vertex r_0 , or to a unique outgoing imperfect edge. Since every chain in Γ has exactly two endpoints, at most two vertices of a chain can have degree 1 in the tree τ .

In the proof of Lemma 4, we charged every perfect segment pq of a covering path γ to a point $\sigma(pq)$, which was an endpoint of an imperfect segment or the endpoint of γ . The function σ relied on the properties established in Lemmas 5–8. These Lemmas also hold for the covering tree τ , if we

replace the endpoints of γ by the leaves (i.e., vertices of degree 1) in τ . With this interpretation, every perfect segment pq is assigned to a point $\sigma(pq)$, which is either an endpoint of an imperfect segment of τ or an endpoint of a chain in Γ . We are now ready to define our charging scheme for τ . Let pq be a perfect segment of τ .

- (i) If $\sigma(pq)$ is an endpoint of an imperfect segment of τ , then charge pq to $\sigma(pq)$.
- (ii) Otherwise $\sigma(pq)$ is an end point of a chain $\gamma_{pq} \in \Gamma$. In this case, if γ_{pq} is incident to r_0 , then charge pq to the root r_0 of τ , else charge pq to the outgoing imperfect edge of τ incident to γ_{pq} .

Similarly to the proof of Lemma 4, each endpoint of an imperfect segment is charged at most twice. Since every imperfect segment has two endpoints, rule (i) is responsible for a total charge of at most $4(s_0 + s_1)$. Each of the two endpoints of a chain $\gamma \in \Gamma$ is charged at most twice by σ , and so the root r_0 and every (directed) imperfect segment is charged at most four times. Thus rule (ii) is responsible for a total charge of at most $4(s_0 + s_1 + 1)$. Altogether the total charge assigned by rules (i) and (ii) is $s_2 \leq 8(s_0 + s_1) + 4$, as required. \square

Proof of Proposition 1. The case of small n ($n \leq 4$) is easy to handle, so assume that $n \geq 5$. Given S , compute $\text{conv}(S)$ and let s_1 be a segment extension of an arbitrary edge of $\text{conv}(S)$; s_1 is long enough so that it intersects all non-parallel lines induced by pairs of points in $S' := S \setminus S \cap s_1$. For simplicity of exposition assume that s_1 is a vertical segment with all other points in S' lying left of s_1 . If $\text{conv}(S')$ is a vertical segment, since $|S| \geq 5$, it is easy to find a 3-segment covering tree for S . If $\text{conv}(S')$ is not a vertical segment, select a non-vertical hull edge of $\text{conv}(S')$ and extend it to the right until it hits s_1 and to the left until it hits all other non-parallel lines induced by pairs of points. Let s_2 be this segment extension.

Continue in a similar way on the set of remaining points, $S'' := S' \setminus S' \cap s_2$, by choosing an arbitrary edge of $\text{conv}(S'')$ and extending it until it hits s_1 or s_2 . If a single point is left at the end, pick a segment incident to it and extend it until it hits the tree made from the previously chosen segments. Otherwise continue by extending an arbitrary hull edge of the remaining points until it hits the tree made from the previously chosen segments. Clearly the resulting tree covers all points and has at most $\lceil n/2 \rceil$ segments.

For the lower bound, it is clear that n points in general position require at least $\lceil n/2 \rceil$ segments in any covering tree. \square

6 Bicolored Variants

Proof of Corollary 2. 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 noncrossing. 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

noncrossing 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 3, 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 noncrossing covering path, the resulting n -element point set requires at least $5n/9 - O(1)$ segments in any pair of mutually noncrossing covering paths. \square

Proof of Corollary 3. For the lower bound we use two copies, red and blue, of the point-set from the proof of Theorem 5 (which is the same as the point-set from the proof of Theorem 3). It remains to show the upper bound. Assume without loss of generality that no two points have the same x -coordinate. Cover the blue points by a blue star with a center high above, and the red points by a red star with a center way below. Obviously, each star is non-crossing, and the distinct x -coordinates of the points suffice to guarantee that the two stars are mutually non-crossing for suitable center positions. The two centers can be easily computed after sorting the points in the above order. \square

7 Computational Complexity

Proof of Theorem 6. 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 [22]. 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 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

8 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. It remains an open problem to close or narrow the gap between the lower and upper bounds for $g(n)$, $(5n-4)/9 \leq g(n) \leq \lceil (1-1/601080391)n \rceil - 1$; and for $t(n)$, $(9n-4)/17 \leq t(n) \leq \lfloor 5n/6 \rfloor$.
3. Let $p(n)$ denote the maximum integer such that every set of n points in the plane has a perfect subset of size $p(n)$. As noticed by Welzl [9, 24], $p(n) = \Omega(\log n)$ immediately follows from the theorem of Erdős and Szekeres [10]. Any improvement in this lower bound would lead to a better upper bound on $f(n)$ in Theorem 1, and thus to a smaller gap relative to the trivial lower bound $f(n) \geq n/2$. 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, 15]. It is also NP-complete for axis-parallel paths in \mathbb{R}^{10} , as shown in [14]. 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 noncrossing paths?
7. Given n points (n even), is it possible to compute a noncrossing 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 noncrossing covering paths and noncrossing spanning paths, and this is asymptotically optimal by Theorem 6. Observe finally that a noncrossing 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] G. Aloupis, A history of linear-time convex hull algorithms for simple polygons, <http://cgm.cs.mcgill.ca/~athens/cs601/>.
- [2] E. M. Arkin, J. S. B. Mitchell and C. D. Piatko, Minimum-link watchman tours, *Inf. Proc. Lett.* **86** (2003), 203–207.
- [3] S. Bereg, P. Bose, A. Dumitrescu, F. Hurtado and P. Valtr, Traversing a set of points with a minimum number of turns, *Discrete Comput. Geom.* **41(4)** (2009), 513–532.
- [4] H. Brönnimann and T. M. Chan, Space-efficient algorithms for computing the convex hull of a simple polygonal line in linear time, *Comput. Geom.* **34** (2006), 75–82.
- [5] B. Chazelle, Triangulating a simple polygon in linear time, *Discrete Comput. Geom.* **6** (1991), 485–524.

- [6] V. Chvátal and G. Klincsek, Finding largest convex subsets, *Congressus Numerantium* **29** (1980), 453–460.
- [7] M. J. Collins and M. E. Moret, Improved lower bounds for the link length of rectilinear spanning paths in grids, *Inf. Proc. Lett.* **68** (1998), 317–319.
- [8] M. J. Collins, Covering a set of points with a minimum number of turns, *Internat. J. Comput. Geom. Appl.* **14(1-2)** (2004), 105–114.
- [9] E.D. Demaine and J. O’Rourke, Open problems from CCCG 2010, in *Proc. 23rd Canadian Conf. Comput. Geom.*, 2011, Toronto, ON, pp. 153–156.
- [10] P. Erdős and G. Szekeres, A combinatorial problem in geometry, *Compositio Mathematica* **2** (1935), 463–470.
- [11] P. Erdős and G. Szekeres, On some extremum problems in elementary geometry, *Annales Univ. Sci. Budapest* **3-4** (1960–1961), 53–62.
- [12] V. Estivill-Castro, A. Heednacram, and F. Suraweera, NP-completeness and FPT results for rectilinear covering problems, *Journal of Universal Computer Science* **15** (2010), 622–652.
- [13] R. Fulek, B. Keszegh, F. Morić and I. Uljarević, On polygons excluding point sets, *Graphs and Combinatorics* (2012), in print.
- [14] M. Jiang, On covering points with minimum turns, *Frontiers in Algorithmics and Algorithmic Aspects in Information and Management - Joint International Conference (FAW-AAIM)*, vol. 7285 of LNCS, Springer, 2012, pp. 58–69.
- [15] E. Kranakis, D. Krizanc and L. Meertens, Link length of rectilinear Hamiltonian tours in grids, *Ars Combinatoria* **38** (1994), 177–192.
- [16] S. Loyd, *Cyclopedia of Puzzles*, The Lamb Publishing Company, 1914.
- [17] J. Matoušek, *Lectures on Discrete Geometry*, Springer, New York, 2002.
- [18] A. Melkman, On-line construction of the convex hull of a simple polygon, *Inf. Proc. Lett.* **25** (1) (1987), 11–12.
- [19] J. S. B. Mitchell, Geometric shortest paths and network optimization, in *Handbook of Computational Geometry (J-R. Sack and J. Urrutia, eds.)*, Chap. 15, Elsevier, 2000, pp. 633–701.
- [20] F. Morić, Separating and covering points in the plane, communication at the Open Problem Session of *22nd Canadian Conf. Comput. Geom.*, Winnipeg, MB, 2010.
- [21] F. Morić, Covering points by a polygonal line, in *Gremo’s Workshop on Open Problems*, Problem Booklet, Wergenstein, Switzerland, 2011.
- [22] F. P. Preparata and M. I. Shamos, *Computational Geometry: An Introduction*, Springer, New York, 1985.
- [23] C. Stein and D. P. Wagner, Approximation algorithms for the minimum bends traveling salesman problem, in *ICALP*, vol. 2081 of LNCS, 2001, pp. 406–421.
- [24] E. Welzl, communication at the *9th Gremo’s Workshop on Open Problems*, Wergenstein, Switzerland, 2011.