Exploring Increasing-Chord Paths and Trees

Yeganeh Bahoo*   Stephane Durocher*§   Sahar Mehrpour†   Debajyoti Mondal‡§

Abstract

A straight-line drawing Γ of a graph $G = (V, E)$ is a drawing of G in the Euclidean plane, where every vertex in G is mapped to a distinct point, and every edge in G is mapped to a straight line segment between their endpoints. A path P in Γ is called increasing-chord if for every four points (not necessarily vertices) a, b, c, d on P in this order, the Euclidean distance between b, c is at most the Euclidean distance between a, d. A spanning tree T rooted at some vertex r in Γ is called increasing-chord if T contains an increasing-chord path from r to every vertex in T. We prove that given a vertex r in a straight-line drawing Γ, it is NP-complete to decide whether Γ contains an increasing-chord spanning tree rooted at r, which answers a question posed by Mastakas and Symvonis [9]. We also shed light on the problem of finding an increasing-chord path between a pair of vertices in Γ, but the computational complexity question remains open.

1 Introduction

In 1995, Icking et al. [6] introduced the concept of a self-approaching curve. A curve is called *self-approaching* if for any three points a, b and c on the curve in this order, $|bc| \leq |ac|$, where $|xy|$ denotes the Euclidean distance between x and y. A path P in a straight-line drawing Γ is called increasing-chord if for every four points (not necessarily vertices) a, b, c, d on P in this order, the inequality $|bc| \leq |ad|$ holds. Γ is called an *increasing-chord drawing* if there exists an increasing-chord path between every pair of vertices in Γ.

Alamdari et al. [1] examined the problem of recognizing increasing-chord drawings, and the problem of constructing such a drawing on a given set of points. They showed that it is NP-hard to recognize increasing-chord drawings in $\mathbb{R}^3$, and asked whether it is also NP-hard in $\mathbb{R}^2$. They also proved that for every set of n points P in $\mathbb{R}^2$, one can construct an increasing-chord drawing Γ with $O(n)$ vertices and edges, where P is a subset

---
*Department of Computer Science, University of Manitoba, Winnipeg, Canada, {bahoo,durocher}@cs.umanitoba.ca

†School of Computing, University of Utah, Utah (UT), USA, mehrpour@cs.utah.edu

‡Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada, dmondal@uwaterloo.ca

§Work of the author is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC).

of the vertices of Γ. In this case, Γ is called a *Steiner network of P*, and the vertices of Γ that do not belong to P are called Steiner points. Dehkordi et al. [4] proved that if P is a convex point set, then one can construct an increasing-chord network with $O(n \log n)$ edges, and without introducing any Steiner point. Mastakas and Symvonis [8] improved the $O(n \log n)$ upper bound on edges to $O(n)$ with at most one Steiner point. Nöllenburg et al. [11] examined the problem of computing increasing-chord drawings of given graphs. Recently, Bonichon et al. [3] showed that the existence of an angle-monotone path of width $0 \leq \gamma < 180°$ between a pair of vertices (in a straight-line drawing) can be decided in polynomial time, which is very interesting since angle-monotone paths of width $\gamma \leq 90°$ satisfy increasing chord property.

Nöllenburg et al. [10] showed that partitioning a plane graph drawing into a minimum number of increasing-chord components is NP-hard, which extends a result of Tan and Kermarrec [12]. They also proved that the problem remains NP-hard for trees, and gave polynomial-time algorithms in some restricted settings. Recently, Mastakas and Symvonis [9] showed that given a point set S and a point $v \in S$, one can compute a rooted minimum-cost spanning tree in polynomial time, where each point in $S \setminus \{v\}$ is connected to v by a path that satisfies some monotonicity property. They also proved that the existence of a monotone rooted spanning tree in a given geometric graph can be decided in polynomial time, and asked whether the decision problem remains NP-hard also for increasing-chord or self-approaching properties.

2 Technical Background

Given a straight line segment l, the *slab of l* is an infinite region lying between a pair of parallel straight lines that are perpendicular to l, and pass through the endpoints of l. Let Γ be a straight-line drawing, and let P be a path in Γ. Then the *slabs* of P are the slabs of the line segments of P. We denote by $\Psi(P)$ the arrangement of the slabs of P. Figure 1(a) illustrates a path P, where the slabs of P are shown in shaded regions. Let A be an arrangement of a set of straight lines such that no line in A is vertical. Then the *upper envelope* of A is a polygonal chain $U(A)$ such that each point of $U(A)$ belongs to some straight line of A, and they are visible from the point $(0, +\infty)$. The upper envelope of a set

Figure 1: (a) Illustration for $\Psi(P)$, where the upper envelope is shown in dashed line. (b) An increasing-chord extension of $a, b, \ldots, p$ is shown in bold.

of slabs is the upper envelope of the arrangement of lines corresponding to the slab boundaries, as shown in dashed line in Figure 1(a). Let t be a vertex in Γ and let $Q = (a, b, \ldots, p)$ be an increasing-chord path in Γ. A path $Q' = (a, b, \ldots, p, \ldots, t)$ in Γ is called an *increasing-chord extension* of Q if Q' is also an increasing-chord path, e.g., see Figure 1(b).

Observation 1 (Icking et al. [7]) *A polygonal path P is increasing-chord if and only if for each point v on the path, the line perpendicular to P at v does not properly intersect P except possibly at v.*

A straightforward consequence of Observation 1 is that every polygonal chain which is both x- and y-monotone, is an increasing-chord path. We will use Observation 1 throughout the paper to verify whether a path is increasing-chord. Let v be a point in $\mathbb{R}^2$. By the *quadrants of v* we refer to the four regions determined by the vertical and horizontal lines through v.

3 Increasing-Chord Rooted Spanning Trees

In this section we prove the problem (IC-Tree) of computing a rooted increasing-chord spanning tree of a given straight-line drawing to be NP-hard.

Theorem 1 *Given a vertex r in a straight-line drawing Γ, it is NP-complete to decide whether Γ admits an increasing-chord spanning tree rooted at r.*

We reduce the NP-complete problem 3-SAT [5] to IC-Tree. Let $I = (X, C)$ be an instance of 3-SAT, where X and C are the set of variables and clauses. We construct a straight-line drawing Γ and choose a vertex r in Γ such that Γ contains an increasing-chord spanning tree rooted at r if and only if I admits a satisfying truth assignment. Here we give an outline of the hardness proof and describe the construction of Γ. A detailed reduction is given in the full version [2].

Assume that $\alpha = |X|$, and $\beta = |C|$. Let l_h be the line determined by the X-axis. Γ will contain $O(\beta)$ points above l_h, one point t on l_h, and $O(\alpha)$ points below l_h, as shown in Figures 2(a)–(b). Each clause $c \in C$ with j literals, will correspond to a set of $j + 1$ points above l_h, and we will refer to the point with the highest y-coordinate among these $j + 1$ points as the *peak t_c of c*. Among the points below l_h, there are 4α points that correspond to the variables and their negations, and two other points, i.e., s and r. In the reduction, the point t and the points below l_h altogether help to set the truth assignments of the variables.

We will first create a straight-line drawing H such that every increasing-chord path between r and t_c, where $c \in C$, passes through s and t. Consequently, any increasing-chord tree T rooted at r (not necessarily spanning), which spans the points t_c, must contain an increasing-chord path $P = (r, s, \ldots, t)$. We will use this path to set the truth values of the variables.

The edges of H below l_h will create a set of thin slabs, and the upper envelope of these slabs will determine a convex chain W above l_h. Each line segment on W will correspond to a distinct variable, as shown in Figure 2(b). The points that correspond to the clauses will be positioned below these segments, and hence some of these points will be 'inaccessible' depending on the choice of the path P. These literal-points will ensure that for any clause $c \in C$, there exists an increasing-chord extension of P from t to t_c if and only if c is satisfied by the truth assignment determined by P.

By the above discussion, I admits a satisfying truth assignment if and only if there exists an increasing-chord tree T in H that connects the peaks to r. But H may still contain some vertices that do not belong to this tree. Therefore, we construct the final drawing Γ by adding some new paths to H, which will allow us to reach these remaining vertices from r. We now describe the construction in details.

Construction of H: We first construct an arrangement $\mathcal{A}$ of 2α straight line segments. The endpoints of the ith line segment L_i, where $1 \leq i \leq 2\alpha$, are $(0, i)$ and $(2\alpha - i + 1, 0)$. We now extend each L_i downward by scaling its length by a factor of $(2\alpha+1)$, as shown in Figure 3(a). Later, the variable x_j, where $1 \leq j \leq \alpha$, and its negation will be represented using the lines L_{2j-1} and L_{2j}. Let l_v be a vertical line segment with endpoints $(2\alpha + 1, 2\alpha)$ and $(2\alpha + 1, -5\alpha^2)$. Since the slope of a line in $\mathcal{A}$ is in the interval $[-2\alpha, -1/(2\alpha)]$, each L_i intersects l_v. Since the coordinates of the endpoints of L_i and l_v are of size $O(\alpha^2)$, and all the intersection points can be represented using polynomial space.

By construction, the line segments of $\mathcal{A}$ appear on $U(\mathcal{A})$ in the order of the variables, i.e., the first two segments (from right) of $U(\mathcal{A})$ correspond to x_1 and $\overline{x_1}$, the next two segments correspond to x_2 and $\overline{x_2}$, etc.

Figure 2: A schematic representation of Γ: (a) Points below l_h, (b) Points above l_h. The points that correspond to c_1 and c_2 are connected in paths of black, and gray, respectively. The slabs of the edges of H that determine the upper envelope are shown in gray straight lines. Each variable and its negation correspond to a pair of adjacent line segments on the upper envelope of the slabs. See the full version [2] for a better illustration.
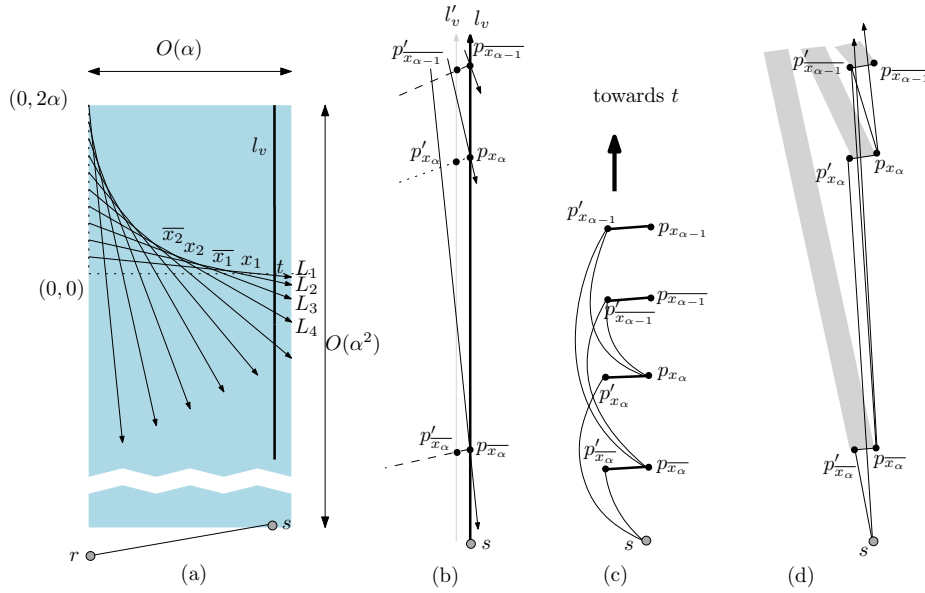

Figure 3: (a) Construction of $\mathcal{A}$. (b)–(c) Construction of the vertices and edges of H_b. (d) Illustration for the straight line segments of H_b, and the slabs corresponding to the needles.

Variable Gadgets: We denote the intersection point of l_h and l_v by t, and the endpoint $(2\alpha + 1, -5\alpha^2)$ of l_v by s. We now create the points that correspond to the variables and their negations. Recall that L_{2j-1} and L_{2j} correspond to the variable x_j and its negation $\overline{x_j}$, respectively. Denote the intersection point of L_{2j-1} and l_v by p_{x_j}, and the intersection point of L_{2j} and l_v by $p_{\overline{x_j}}$, e.g., see Figure 3(b). For each p_{x_j} ($p_{\overline{x_j}}$), we create a new point p'_{x_j} ($p'_{\overline{x_j}}$) such that the straight line segment $p_{x_j}p'_{x_j}$ ($p_{\overline{x_j}}p'_{\overline{x_j}}$) is perpendicular to L_{2j-1} (L_{2j}), as shown using the dotted (dashed) line in Figure 3(b). We may assume that all the points p'_{x_j} and $p'_{\overline{x_j}}$ lie on a vertical line l'_v, where l'_v lies ε distance away to the left of l_v. The value of ε would be determined later. In the following we use the points p_{x_j}, $p_{\overline{x_j}}$, p'_{x_j} and $p'_{\overline{x_j}}$ to create some polygonal paths from s to t.

For each j from 1 to α, we draw the straight line segments $p_{x_j}p'_{x_j}$ and $p_{\overline{x_j}}p'_{\overline{x_j}}$. Then for each k, where $1 < k \le \alpha$, we make p_{x_k} and $p_{\overline{x_k}}$ adjacent to both $p'_{x_{k-1}}$ and $p'_{\overline{x_{k-1}}}$, e.g., see Figure 3(c). We then add the edges from s to p'_{x_α} and $p'_{\overline{x_\alpha}}$, and finally, from t to p_{x_1} and $p_{\overline{x_1}}$. For each x_j ($\overline{x_j}$), we refer to the segment $p_{x_j}p'_{x_j}$ ($p_{\overline{x_j}}p'_{\overline{x_j}}$) as the *needle* of x_j ($\overline{x_j}$). Figure 3(c) illustrates the needles in bold. Let the resulting drawing be H_b.

Recall that l'_v is ε distance away to the left of l_v. We choose ε sufficiently small such that for each needle, its slab does not intersect any other needle in H_b, e.g., see Figure 3(d). The upper envelope of the slabs of all the straight line segments of H_b coincides with $U(\mathcal{A})$. Since the distance between any pair of points that we created on l_v is at least $1/\alpha$ units, it suffices to choose $\varepsilon = 1/\alpha^3$. Note that the points p'_{x_j} and $p'_{\overline{x_j}}$ can be represented in polynomial space using the endpoints of

l'_v and the endpoints of the segments L_{2j-1} and L_{2j}. The proof of the following lemma is omitted (see [2]).

Lemma 2 *Every increasing-chord path P that starts at s and ends at t must pass through exactly one point among p_{x_j} and $p_{\overline{x_j}}$, where $1 \le j \le \alpha$, and vice versa.*

We now place a point r on the y-axis sufficiently below H_b, e.g., at position $(0, -\alpha^5)$, such that the slab of the straight line segment rs does not intersect H_b (except at s), and similarly, the slabs of the line segments of H_b do not intersect rs. Furthermore, the slab of rs does not intersect any segment L_j, and vice versa. We then add the point r and the segment rs to H_b. Let P be an increasing-chord path from r to t. The upper envelope of $\Psi(P)$ is determined by the needles in P, which selects some segments from the convex chain W, e.g., see Figure 2(b). For each x_j, P passes through exactly one point among p_{x_j} and $p_{\overline{x_j}}$. Therefore, for each variable x_j, either the slab of x_j, or the slab of $\overline{x_j}$ appears on $U(P)$. Later, if P passes through point p_{x_j} ($p_{\overline{x_j}}$), then we will set x_j to false (true). Since P is an increasing-chord path, by Lemma 2 it cannot pass through both p_{x_j} and $p_{\overline{x_j}}$ simultaneously. Therefore, all the truth values will be set consistently.

Clause Gadgets: We now complete the construction of H by adding clause gadgets to H_b. For each clause c_i, where $1 \le i \le \beta$, we first create the peak point t_{c_i} at position $(0, 2\alpha + i)$. For each variable x_j, let λ_{x_j} be the interval of L_{2j-1} that appears on the upper envelope of $\mathcal{A}$. Similarly, let $\lambda_{\overline{x_j}}$ be the interval of L_{2j} on the upper envelope of $\mathcal{A}$. For each c_i, we construct a point q_{x_j, c_i} ($q_{\overline{x_j}, c_i}$) inside the cell of $\mathcal{A}$ immediately below λ_{x_j} ($\lambda_{\overline{x_j}}$). We will refer to these points as the *literal-points of c_i*. The full version [2] depicts these points in black squares. We assume that for each variable, the corresponding literal-points lie on the same location. One may perturb them to remove vertex overlaps. For each variable $x \in c_i$, we create a path (t, x, t_c). In the reduction, if at least one of the literals of c_i is true, then we can take the corresponding path to connect t_c to t. Let the resulting drawing be H.

Construction of Γ: Let q be a literal-point in H. We now add an increasing-chord path $P' = (r, a, q)$ to H in such a way that P' cannot be extended to any larger increasing-chord path in H. We place the point a at the intersection point of the horizontal line through q and the vertical line through r, the full version [2] contains the details. We refer to the point a as the *anchor* of q. By the construction of H, all the neighbors of q that have a higher y-coordinate than q lie in the top-left quadrant of q. Let q' be the first neighbor in the top-left quadrant of q in counter clockwise order. Since $\angle aqq' < 90°$, P' cannot be extended to any larger increasing-chord path (r, a, q, w) in H, where the y-coordinate of w is higher than q. On the other hand, every literal-point

w in H with y-coordinate smaller than q intersects the slab of ra. Therefore, P' cannot be extended to any larger increasing-chord path.

For every literal-point q in H, we add such an increasing-chord path from t to q. To avoid edge overlaps, one can perturb the anchors such that the new paths remain increasing-chord and non-extensible to any larger increasing-chord paths. This completes the construction of Γ. We refer the reader to the full version [2] for the formal details of the reduction.

4 Increasing-Chord Paths

In this section we attempt to reduce 3-SAT to the problem of finding an increasing-chord path (IC-PATH) between a pair of vertices in a given straight-line drawing. We were unable to bound the coordinates of the drawing to a polynomial number of bits, and hence the computational complexity question of the problem remains open. We hope that the ideas we present here will be useful in future endeavors to settle the question.

Here we briefly describe the idea of the reduction. Given a 3-SAT instance $I = (X, C)$, the corresponding drawing $\mathcal{D}$ for IC-PATH consists of straight-line drawings $\mathcal{D}_{i-1}$, where $1 \le i \le \beta$, e.g., see Figure 4(a). The drawing $\mathcal{D}_{i-1}$ corresponds to the each clause c_i. We will refer to the bottommost (topmost) point of $\mathcal{D}_{i-1}$ as $t_{c_{i-1}}$ (t_{c_i}). We will choose t_{c_0} and t_{c_β} to be the points t and t', respectively, and show that I admits a satisfying truth assignment if and only if there exists an increasing-chord path P from t to t' that passes through every t_{c_i}. For every i, the subpath P_{i-1} of P between $t_{c_{i-1}}$ and t_{c_i} will correspond to a set of truth values for all the variables in X. The most involved part is to show that the truth values determined by P_{i-1} and P_i are consistent. This consistency will be ensured by the construction of $\mathcal{D}$, i.e., the increasing-chord path P_{i-1} from $t_{c_{i-1}}$ to t_{c_i} in $\mathcal{D}_{i-1}$ will determine a set of slabs, which will force a unique increasing-chord path P_i in $\mathcal{D}_i$ between t_{c_i} and $t_{c_{i+1}}$ with the same truth values as determined by P_{i-1}.

Construction of $\mathcal{D}$: The construction of $\mathcal{D}_{i-1}$ depends on an arrangement of lines $\mathcal{A}^{i-1}$. The construction of $\mathcal{A}^0$ is the same as the construction of arrangement $\mathcal{A}$, which we described in Section 3. Figure 4(c) illustrates $\mathcal{A}^0$ in dotted lines. For each variable x_j, where $1 \le j \le \alpha$, there exists an interval $\lambda_{x_j}^0$ of L_{2j-1} on the upper envelope of $\mathcal{A}^0$. Similarly, for each $\overline{x_j}$, there exists an interval $\lambda_{\overline{x_j}}^0$ of L_{2j} on the upper envelope of $\mathcal{A}^0$.

We now describe the construction of $\mathcal{D}_0$. Choose t_{c_0} (t_{c_1}) to be the bottommost (topmost) point of $\lambda_{x_1}^0$ ($\lambda_{\overline{x_\alpha}}^0$). We then slightly shrink the intervals $\lambda_{x_1}^0$ and $\lambda_{\overline{x_\alpha}}^0$ such that t_{c_0} and t_{c_1} no longer belong to these segments. Assume that c_1 contains δ literals, where $\delta \le 3$, and let $\sigma_1, \ldots, \sigma_{2^\delta - 1}$ be the satisfying truth assignments for c_1. We construct a graph G_{c_1} that corresponds to these sat-

Figure 4: (a) Idea for the reduction. (b) The graph corresponding to the truth assignment satisfying $c_1 = (x_1 \vee x_4)$. Only the construction for the truth assignments $\sigma_1 = \{x_1 = true, x_4 = true\}$ and $\sigma_2 = \{x_1 = true, x_4 = false\}$ are shown. (c) A schematic representation for $\mathcal{D}$, where $I = (x_1 \vee x_3) \wedge (\overline{x_2} \vee \overline{x_3})$. An increasing-chord path is shown in bold, and the corresponding truth value assignment is: $x_1 = true$, $x_2 = false$, $x_3 = true$. (d) An s-segment.

isfying truth assignments, e.g., see Figure 4(b) and the full version [2] for formal details. The idea is to ensure that any path between t_{c_0} and t_{c_1} passes through exactly one point in $\{q_{x_j}^{\sigma_k}, q_{\overline{x_j}}^{\sigma_k}\}$, for each truth assignment σ_k, which will set the truth value of x_j. In $\mathcal{D}_0$, the point $q_{x_j}^{\sigma_k}$ ($q_{\overline{x_j}}^{\sigma_k}$) is chosen to be the midpoint of $\lambda_{x_j}^{i-1}$ ($\lambda_{\overline{x_j}}^{i-1}$). Later, we will refer to these points as *q-points*, e.g., see Figure 4(c). We may assume that for each x_j, the points $q_{x_j}^{\sigma_k}$ lie at the same location. One may later perturb them to remove vertex overlaps.

By Observation 1, any y-monotone path P' between t_{c_0} and t_{c_1} must be an increasing-chord path. If P' passes through $q_{x_j}^\sigma$, then we set x_j to true. Otherwise, P' must pass through $q_{\overline{x_j}}^\sigma$, and we set x_j to false. In the following we replace each q-point by a small segment. The slabs of these segments will determine $\mathcal{A}^1$. Consider an upward ray r^1 with positive slope starting at the q-point on λ_{x_1}, e.g., see Figure 4(c). Since all the edges that are currently in $\mathcal{D}_0$ have negative slopes, we can choose a sufficiently large positive slope for r^1 and a point a^1 on r^1 such that all the slabs of $\mathcal{D}_0$ lie below a^1. We now find a point b^1 above a^1 on r^1 with sufficiently large y-coordinate such that the slab of $t_{c_1} b^1$ does not intersect the edges in $\mathcal{D}_0$. Let $l_{\overline{x_1}}^1$ be the line determined by r^1. For each x_j and $\overline{x_j}$ (except for $j = 1$), we now construct the lines $l_{\overline{x_j}}^1$ and $l_{x_j}^1$ that pass through their corresponding q-points and intersect r^1 above b^1. The lines $l_{x_j}^1$ and $l_{\overline{x_j}}^1$ determine the arrangement $\mathcal{A}^1$. Observe that one can construct these lines in the decreasing order of the x-coordinates of their q-points, and ensure that for each $l_{x_j}^1$ ($l_{\overline{x_j}}^1$), there exists an interval $\lambda_{x_j}^1$ ($\lambda_{\overline{x_j}}^1$) on the upper envelope of $\mathcal{A}^1$. Note that the correspondence is inverted, i.e., in $\mathcal{A}^1$, $\lambda_{\overline{x_j}}^1$ corresponds to $\lambda_{x_j}^0$, and $\lambda_{x_j}^1$ corresponds to $\lambda_{\overline{x_j}}^0$.

For each j, we draw a small segment $s_{x_j}^0$ ($s_{\overline{x_j}}^0$) perpendicular to $l_{x_j}^1$ ($l_{\overline{x_j}}^1$) that passes through the q-point and lies to the left of q, e.g., see Figure 4(d). The construction of $\mathcal{D}_i$, where $i > 1$, is more involved. The upper envelope of $\mathcal{A}^{i+1}$ is determined by the upper envelope of the slabs of the s-segments in $\mathcal{D}_{i-1}$. For each i, we construct the q-points and corresponding graph G_{c_i}. The full version [2] includes the formal details.

In the reduction we show that any increasing-chord path P from t to t' contains the points t_{c_i}. We set a variable x_j true or false depending on whether P passes through $s_{x_j}^0$ or $s_{\overline{x_j}}^0$. The construction of $\mathcal{D}$ imposes the constraint that if P passes through $s_{x_j}^{i-1}$ ($s_{\overline{x_j}}^{i-1}$), then it must pass through $s_{x_j}^i$ ($s_{\overline{x_j}}^i$). Hence the truth values in all the clauses are set consistently. By construction of G_{c_i}, any increasing-chord path between $t_{c_{i-1}}$ to t_{c_i} determines a satisfying truth assignment for c_i. On the other hand, if I admits a satisfying truth assignment, then for each clause c_i, we choose the corresponding increasing-chord path P_i between $t_{c_{i-1}}$ and t_{c_i}. The union of all P_i yields the required increasing-chord path P from t to t'. The full version [2] presents the construction in details, and explains the challenges of encoding $\mathcal{D}$ in a polynomial number of bits.

5 Open Problems

The most intriguing problem in this context is to settle the computational complexity of the increasing-chord path (IC-PATH) problem. Another interesting question is whether the problem IC-TREE remains NP-hard under the planarity constraint; a potential attempt to adapt our hardness reduction could be replacing the edge intersections by dummy vertices.

References

[1] S. Alamdari, T. M. Chan, E. Grant, A. Lubiw, and V. Pathak. Self-approaching graphs. In *Proc. of GD*, volume 7704 of *LNCS*, pages 260–271. Springer, 2013.

[2] Y. Bahoo, S. Durocher, S. Mehrpour, and D. Mondal. Exploring increasing-chord paths and trees. *CoRR*, abs/1702.08380, 2017.

[3] N. Bonichon, P. Bose, P. Carmi, I. Kostitsyna, A. Lubiw, and S. Verdonschot. Gabriel triangulations and angle-monotone graphs: Local routing and recognition. In *Proc. of GD*, volume 9801 of *LNCS*, pages 519–531. Springer, 2016.

[4] H. R. Dehkordi, F. Frati, and J. Gudmundsson. Increasing-chord graphs on point sets. *Journal of Graph Algorithms and Applications*, 19(2):761–778, 2015.

[5] M. R. Garey and D. S. Johnson. *Computers and intractability*. Freeman, San Francisco, 1979.

[6] C. Icking and R. Klein. Searching for the kernel of a polygon - A competitive strategy. In *Proc. of SoCG*, pages 258–266. ACM, 1995.

[7] C. Icking, R. Klein, and E. Langetepe. Self-approaching curves. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 125, pages 441–453. Cambridge Univ Press, 1999.

[8] K. Mastakas and A. Symvonis. On the construction of increasing-chord graphs on convex point sets. In *Proc. of IISA*, pages 1–6. IEEE, 2015.

[9] K. Mastakas and A. Symvonis. Rooted uniform monotone minimum spanning trees. In *Proc. of CIAC*, volume 10236 of *LNCS*, pages 405–417. Springer, 2017.

[10] M. Nöllenburg, R. Prutkin, and I. Rutter. Partitioning graph drawings and triangulated simple polygons into greedily routable regions. In *Proc. of ISAAC*, volume 9472 of *LNCS*, pages 637–649. Springer, 2015.

[11] M. Nöllenburg, R. Prutkin, and I. Rutter. On self-approaching and increasing-chord drawings of 3-connected planar graphs. *JoCG*, 7(1):47–69, 2016.

[12] G. Tan and A. Kermarrec. Greedy geographic routing in large-scale sensor networks: A minimum network decomposition approach. *IEEE/ACM Trans. Netw.*, 20(3):864–877, 2012.