Extending Simple Monotone Drawings

Jan Kynčl Jan Soukup

Submitted: Sep 30, 2025; Accepted: May 8, 2026; Published: Jul 3, 2026
© The authors. Released under the CC BY-ND license (International 4.0).

Abstract

We prove the following variant of Levi's Enlargement Lemma: for an arbitrary arrangement $\mathcal{A}$ of x-monotone pseudosegments in the plane and a pair of points a, b with distinct x-coordinates and not on the same pseudosegment, there exists a simple x-monotone curve with endpoints a, b that intersects every curve of $\mathcal{A}$ at most once. As a consequence, every simple monotone drawing of a graph can be extended to a simple monotone drawing of a complete graph. We also show that extending an arrangement of cylindrically monotone pseudosegments is not always possible; in fact, the corresponding decision problem is NP-hard.

Mathematics Subject Classifications: 05C10, 68R10

1 Introduction

Given $k \geqslant 1$, a finite set $\mathcal{A}$ of simple curves in the plane is called an *arrangement of k-strings* if every pair of the curves of $\mathcal{A}$ intersects at most k times, and every intersection point is either a proper crossing or an endpoint of at least one of the curves of the intersecting pair. Multiple k-strings can intersect at a common point, an endpoint of one curve can lie on another, but two curves cannot touch in their inner points. An arrangement of 1-strings is also called an *arrangement of pseudosegments*, and each curve in the arrangement is called a *pseudosegment*. In this paper, we represent simple curves as subsets of the plane (or another surface) that are homeomorphic images of a closed interval. When necessary, we also specify one endpoint of a curve as the starting point to introduce an orientation to the curve.

A simple curve γ in the plane is *x-monotone*, shortly *monotone*, if γ intersects every line parallel to the y-axis at most once.

Let $\mathcal{F}$ be a family of arrangements of pseudosegments in the plane, and let a, b be a pair of points in the plane. An arrangement $\mathcal{A}$ from $\mathcal{F}$ is *(a, b)-extendable in $\mathcal{F}$* if there exists a simple curve α with endpoints a, b so that $\mathcal{A} \cup \{\alpha\} \in \mathcal{F}$. The arrangement $\mathcal{A}$ is *extendable in $\mathcal{F}$* if it is (a, b)-extendable in $\mathcal{F}$ for all possible choices of a and b with

Department of Applied Mathematics, Faculty of Mathematics and Physics, Charles University, Czech Republic (kyncl@kam.mff.cuni.cz, soukup@kam.mff.cuni.cz).

 https://doi.org/10.37236/14683

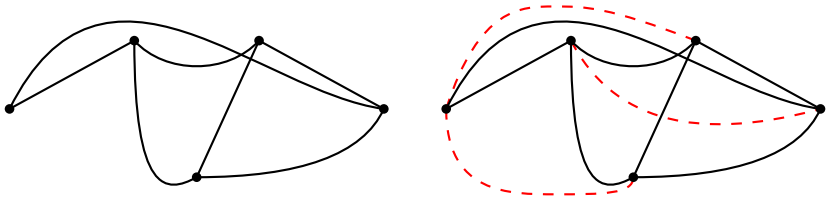


Figure 1: Left: a simple monotone drawing of a graph. Right: an extension of the drawing on the left to a simple monotone drawing of a complete graph. The added edges are dashed.

distinct x-coordinates and not on the same pseudosegment. We omit the family $\mathcal{F}$ from the notation whenever it is clear from context. The family $\mathcal{F}$ is *extendable* if all its elements are extendable in $\mathcal{F}$.

Our main result is the following.

Theorem 1. *The family of all arrangements of monotone pseudosegments in the plane is extendable.*

Moreover, in Subsection 2.2 we show that the proof of Theorem 1 can be turned into an efficient algorithm: given a suitable representation of an arrangement of n monotone pseudosegments with m incidences between endpoints or intersection points and pseudosegments, we can find the new pseudosegment extending the arrangement in time $O(m)$.

A drawing of a graph in the plane is *simple* if every pair of edges has at most one common point, either a common endpoint or a proper crossing. A drawing of a graph is *monotone* if every edge is drawn as a monotone curve and no two vertices share the same x-coordinate. We have the following direct consequence of Theorem 1, illustrated in Figure 1.

Corollary 2. *Every simple monotone drawing of a graph in the plane can be extended to a simple monotone drawing of the complete graph with the same set of vertices.*

1.1 Drawings on a cylinder

A cylinder can be represented as the surface $S^1 \times \mathbb{R}$ embedded in $\mathbb{R}^3$, where S^1 is the unit circle in the xy-plane. A simple curve γ on the cylinder is *cylindrically monotone* if γ intersects every line parallel to the z-axis at most once. A finite system of circular arcs on the circle S^1 is *normal* if no pair of arcs covers the whole circle [14]. An arrangement of cylindrically monotone pseudosegments is *normal* if the projections of its curves to S^1 form a normal system of circular arcs. Drawings of graphs whose edges are cylindrically monotone have also been called *angularly monotone* [8], and in the case when the edges form a normal arrangement they have also been called *strongly c-monotone* [1].

Let $\mathcal{F}$ be a family of arrangements of pseudosegments on the cylinder, and let a, b be a pair of points on the cylinder. An arrangement $\mathcal{A}$ from $\mathcal{F}$ is (a, b)-*extendable in* $\mathcal{F}$ if there

exists a simple curve α with endpoints a, b so that $\mathcal{A} \cup \{\alpha\} \in \mathcal{F}$. The arrangement $\mathcal{A}$ is *extendable in* $\mathcal{F}$ if it is (a, b)-extendable in $\mathcal{F}$ for all possible choices of a and b not on the same vertical line and not on the same pseudosegment. We omit the family $\mathcal{F}$ from the notation whenever it is clear from context. The family $\mathcal{F}$ is *extendable* if all its elements are extendable in $\mathcal{F}$.

Theorem 1 generalizes in a straightforward way to normal arrangements of cylindrically monotone pseudosegments:

Corollary 3. *The family of all normal arrangements of cylindrically monotone pseudosegments is extendable.*

It may be helpful to consider an alternative definition of normal arrangements as those where the curves are drawn the "shorter way" around the cylinder. The following result implies that these two definitions are combinatorially equivalent.

Proposition 4. *Given a normal system $\mathcal{C}$ of circular arcs on the circle S^1, there exists a homeomorphism of S^1 that maps each arc in $\mathcal{C}$ to an arc of length smaller than π.*

We suspect that Proposition 4 might be a known result, but we were not able to find it in the literature.

We show that further generalization of Corollary 3 to arbitrary arrangements of cylindrically monotone pseudosegments is not possible.

Proposition 5. *There exists an arrangement $\mathcal{A}$ of five cylindrically monotone pseudosegments and a pair of points a, b not on the pseudosegments of $\mathcal{A}$ and not on the same vertical line such that $\mathcal{A}$ is not (a, b)-extendable in the family of all arrangements of cylindrically monotone pseudosegments.*

Moreover, the decision problem of extendability of cylindrically monotone arrangements turns out to be NP-hard.

Theorem 6. *Given an arrangement $\mathcal{A}$ of cylindrically monotone pseudosegments and a pair of points a, b, it is NP-hard to decide whether $\mathcal{A}$ is (a, b)-extendable in the family of all arrangements of cylindrically monotone pseudosegments.*

The problem in Theorem 6 is, in fact, NP-complete: the membership in NP is rather straightforward since an arrangement of n cylindrically monotone pseudosegments can be encoded by a sequence consisting of $O(n)$ endpoints and $O(n^2)$ crossings, and by the above/below relations between the pseudosegments and the endpoints or crossings of other pseudosegments.

We prove Theorem 1 in Section 2 and we prove Corollary 3, Proposition 4, Proposition 5 and Theorem 6 in Section 3.

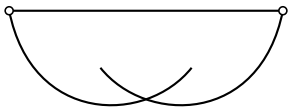


Figure 2: An example of an arrangement of three pseudosegments that cannot be extended to pseudolines forming a pseudoline arrangement since any extension of the pseudosegments into pseudolines would contain a pair of pseudolines with two mutual intersections.

1.2 Related results

A *pseudoline* in the plane is an image of a Euclidean line under a homeomorphism of the plane; in other words, a pseudoline is a homeomorphic image of the set $\mathbb{R}$, unbounded in both directions. An *arrangement of pseudolines* is a finite set of pseudolines such that every pair of them has exactly one crossing, and no other common intersection point. Pseudolines are also often defined in the projective plane, as nonseparating simple closed curves.

Levi's Enlargement Lemma [13] states that for every arrangement of pseudolines and every pair of points a, b not on the same pseudoline, one can draw a new pseudoline through a and b, crossing every curve from the given arrangement exactly once. The lemma has several alternative proofs in the literature [5, 15].

Snoeyink and Hershberger [16] generalized Levi's Lemma to a sweeping theorem for pseudoline arrangements, which allows "rotating" a new pseudoline through a given point, sweeping the whole plane in the process.

By a classical result of Goodman [9], [7, Theorem 5.1.4], every arrangement of pseudolines can be transformed by a homeomorphism of the plane into an arrangement of monotone pseudolines, or a so-called *wiring diagram*. Therefore, monotone arrangements of pseudosegments can be considered as a generalization of pseudoline arrangements. On the other hand, Figure 2 shows an example that not every monotone arrangement of pseudosegments can be seen as a "restriction" of a pseudoline arrangement, and so Theorem 1 does not easily follow from Levi's Lemma. See Arroyo, Bensmail and Richter [3, Figure 2] for more examples. Since a pseudoline (in the projective plane) can be considered as a union of two internally disjoint pseudosegments (one bounded and the other one crossing the line at infinity), Theorem 1 can also be considered as a generalization of "a half" of Levi's Lemma.

By $P_3 + P_3$ we denote the union of two vertex-disjoint paths of length 2. A simple drawing of $P_3 + P_3$ that cannot be extended to a simple drawing of a complete graph was constructed by Eggleton [6, Diagram 15(ii)] and later rediscovered by the first author [10, Figure 9]. Later a few more examples of non-extendable simple drawings were constructed [11, Figures 1, 10]. None of these drawings are homeomorphic to monotone drawings, which follows, for example, from Corollary 2, but some of them can easily be transformed into cylindrically monotone drawings; see Figure 3.

The first author together with Pach, Radoičić and Tóth [11] proved that for every $k \geqslant 1$, for every arrangement $\mathcal{A}$ of k-strings and every pair of points a, b not on the same

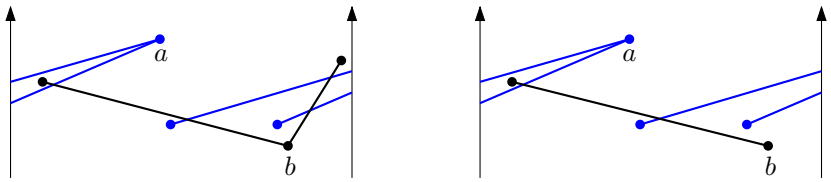


Figure 3: Left: a simple cylindrically monotone drawing of $P_3 + P_3$ where the edge ab cannot be added [6, 11, 10] as a simple curve without crossing some edge of $P_3 + P_3$. Since all edges of $P_3 + P_3$ are incident to either a or b, the added edge ab would have at least two intersections with some edge of $P_3 + P_3$.
Right: a simple cylindrically monotone drawing of $P_3 + P_2$ where the edge ab cannot be added as a cylindrically monotone curve without intersecting some edge of $P_3 + P_2$ twice.

curve of $\mathcal{A}$, there exists a simple curve joining a and b intersecting every curve of $\mathcal{A}$ at most $2k$ times. They also constructed examples showing that the constant $2k$ cannot be improved.

Arroyo, Klute, Parada, Vogtenhuber, Seidel and Wiedera [4] showed that it is NP-hard to decide, given an arrangement $\mathcal{A}$ of pseudosegments and a pair of points a, b, whether a and b can be joined by a simple curve crossing each pseudosegment of $\mathcal{A}$ at most once. Our proof of Theorem 6 is a simple adaptation of this result to cylindrically monotone arrangements. Recently, Aichholzer Orthaber and Vogtenhuber [2] proved a result similar to Theorem 1 where they showed that both the crossing-minimizing and the so-called separable drawings can be extended to simple drawings of the complete graph but not necessarily to the crossing-minimizing or separable drawings, respectively, of the complete graph.

Arroyo, Bensmail and Richter [3] studied a slightly different extendability question: given an arrangement of pseudosegments, can it be extended to an arrangement of pseudolines by extending every given pseudosegment to a pseudoline? They determined the full infinite set of minimal obstructions, and found a polynomial time algorithm for detecting the obstructions and extending the arrangement.

An extended abstract of this paper appeared in the proceedings of IWOCA 2025 [12].

2 Monotone arrangements in the plane

We start with a few definitions and tools for analyzing x-monotone arrangements. Given a pair of points a, b in the plane, we write $a \prec b$ if a has a smaller x-coordinate than b. Clearly, $\prec$ is a strict linear order on the points of any monotone curve.

We can naturally talk about objects lying "below" and "above" monotone curves. Let a, b be points such that $a \prec b$. For any monotone curve γ we denote by $\gamma[a, b]$ and $\gamma(a, b)$ the subset of γ formed by the points x of γ satisfying $a \preceq x \preceq b$ and $a \prec x \prec b$, respectively. Similarly, for an arrangement $\mathcal{A}$ of monotone pseudosegments we denote by $\mathcal{A}[a, b]$ the arrangement of pseudosegments where we replace each $\gamma \in \mathcal{A}$ by $\gamma[a, b]$.

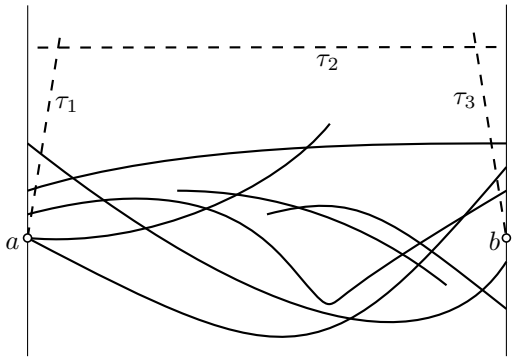


Figure 4: An arrangement of monotone pseudosegments with three added segments τ_1, τ_2, τ_3 connecting points a, b "from above".

By *consecutive intersections* of two monotone curves with finitely many intersections we mean consecutive intersections with respect to their x-coordinates. Let α, β be two monotone curves with finitely many intersections. Let a, b be two consecutive intersections of α, β such that $a \prec b$. Then the only intersections of $\alpha[a, b]$ with $\beta[a, b]$ are the points a and b. In this case we say that the curves α and β form a *bigon*. Furthermore, if $\alpha(a, b)$ lies above $\beta(a, b)$ we say that α and β form an *α-top*, or equivalently, a *β-bottom* bigon. Note that in general α and β can form both an α-bottom and a β-bottom bigon.

The *lower envelope* $\text{low}(\mathcal{U})$ of a set $\mathcal{U}$ of curves is the set of all points p of these curves such that no other point of any curve of $\mathcal{U}$ with the same x-coordinate as p is below p. Note that if $\mathcal{U}$ is an arrangement of monotone pseudosegments, then $\text{low}(\mathcal{U})$ is a finite union of connected parts of pseudosegments.

We continue by proving Theorem 1, followed by a discussion of the time complexity of extending the arrangement.

2.1 Proof of Theorem 1

Let $\mathcal{A}$ be an arrangement of monotone pseudosegments. Let a, b, with $a \prec b$, be points that are not on the same pseudosegment of $\mathcal{A}$. We need to find a monotone curve from a to b that intersects every curve of $\mathcal{A}$ at most once. Since every curve of $\mathcal{A}$ is monotone, we can without loss of generality assume that $\mathcal{A} = \mathcal{A}[a, b]$.

Let $\mathcal{A}'$ be an arrangement of monotone pseudosegments formed by all pseudosegments of $\mathcal{A}$ together with three new segments τ_1, τ_2, τ_3, defined as follows. The segment τ_1 is an almost vertical segment starting in a and ending in some new point to the right of a and above all pseudosegments of $\mathcal{A}$. Similarly, τ_3 is an almost vertical segment ending in b and starting in some new point to the left of b and above all pseudosegments of $\mathcal{A}$. Finally, τ_2 is a horizontal segment crossing τ_1 and τ_3, and lying entirely above all pseudosegments of $\mathcal{A}$; see Figure 4. In this way, $\text{low}(\{\tau_1, \tau_2, \tau_3\})$ is a monotone curve connecting a and b "from above", so that every pseudosegment $\gamma \in \mathcal{A}$ intersects it at most twice. Furthermore no $\gamma \in \mathcal{A}$ forms a γ-top bigon with $\text{low}(\{\tau_1, \tau_2, \tau_3\})$ (it can only form a γ-bottom bigon).

In order to find an extending curve we do the following. We find a nonempty subset $\mathcal{U} \subseteq \mathcal{A}'$ of pseudosegments such that the lower envelope of $\mathcal{U}$ is a monotone curve connecting a to b, intersecting every pseudosegment of $\mathcal{A}' \setminus \mathcal{U}$ at most once. Furthermore, we find $\mathcal{U}$ so that no pseudosegment α touches $\text{low}(\mathcal{U})$ from below in an inner point of α. After finding such $\mathcal{U}$, a new pseudosegment connecting a and b can clearly be drawn slightly below the lower envelope of $\mathcal{U}$ and will indeed intersect every pseudosegment of $\mathcal{A}'$ at most once. Thus, if such $\mathcal{U}$ exists, $\mathcal{A}'$, and consequently $\mathcal{A}$, is (a, b)-extendable.

We find $\mathcal{U}$ inductively. We start with $\mathcal{U}_0 = \{\tau_1, \tau_2, \tau_3\}$ and always look at the lower envelope of $\mathcal{U}_i$. In the ith step we select an arbitrary pseudosegment γ_i of $\mathcal{A}' \setminus \mathcal{U}_{i-1}$ intersecting $\text{low}(\mathcal{U}_{i-1})$ at least twice. If there is no such γ_i then $\mathcal{U} = \mathcal{U}_{i-1}$ and we are done. Otherwise, we set $\mathcal{U}_i = \mathcal{U}_{i-1} \cup \{\gamma_i\}$. The number of pseudosegments is finite, so this process finishes with a set $\mathcal{U}$ such that the lower envelope of $\mathcal{U}$ intersects every pseudosegment of $\mathcal{A}' \setminus \mathcal{U}$ at most once.

Additionally, we prove that the induction preserves the following invariants for every $\mathcal{U}_i$.

(I1) No pseudosegment α of $\mathcal{A}' \setminus \mathcal{U}_i$ forms an α-top bigon with $\text{low}(\mathcal{U}_i)$.

(I2) No pseudosegment α of $\mathcal{A}' \setminus \mathcal{U}_i$ touches $\text{low}(\mathcal{U}_i)$ from below in an inner point of α.

(I3) The lower envelope of $\mathcal{U}_i$ is connected and contains a and b.

In particular, by (I3), the lower envelope of $\mathcal{U}$ is a monotone curve connecting a to b and, by (I2), no pseudosegment α of $\mathcal{A}' \setminus \mathcal{U}$ touches $\text{low}(\mathcal{U})$ from below in an inner point of α. Since $\text{low}(\mathcal{U})$ intersects every pseudosegment of $\mathcal{A}' \setminus \mathcal{U}$ at most once by its construction, $\mathcal{A}$ is (a, b)-extendable by the previous discussion. Thus, it suffices to prove the correctness of these invariants to finish the proof.

The invariants hold for $\mathcal{U}_0$ by the construction of τ_1, τ_2 and τ_3. Suppose all invariants hold for $\mathcal{U}_{i-1}$. In particular, $\text{low}(\mathcal{U}_{i-1})$ is a monotone curve connecting a to b by invariant (I3). We show that all invariants also hold for $\mathcal{U}_i$.

The pseudosegment γ_i intersects $\text{low}(\mathcal{U}_{i-1})$ at least twice. We show that γ_i intersects $\text{low}(\mathcal{U}_{i-1})$ exactly twice. Suppose, for contradiction, that there are three consecutive intersections c, d and e of γ_i with $\text{low}(\mathcal{U}_{i-1})$ such that $c \prec d \prec e$. Then $\gamma_i[c, d]$ with $\text{low}(\mathcal{U}_{i-1})[c, d]$ forms a bigon and so does $\gamma_i[d, e]$ with $\text{low}(\mathcal{U}_{i-1})[d, e]$. By invariant (I1) both of these bigons must be $\text{low}(\mathcal{U}_{i-1})$-top bigons. However, in this case γ_i touches $\text{low}(\mathcal{U}_{i-1})$ from below in the point d. That is not possible by invariant (I2). Thus, γ_i intersects $\text{low}(\mathcal{U}_{i-1})$ exactly twice. Furthermore, by invariant (I1), γ_i and $\text{low}(\mathcal{U}_{i-1})$ form a γ_i-bottom bigon.

Let x and y be the two intersection points of γ_i and $\text{low}(\mathcal{U}_{i-1})$. Refer to Figure 5. Since γ_i and $\text{low}(\mathcal{U}_{i-1})$ form a γ_i-bottom bigon, the only part of the curve γ_i that lies below $\text{low}(\mathcal{U}_{i-1})$ is exactly $\gamma_i(x, y)$. Thus, the lower envelope of $\mathcal{U}_{i-1} \cup \{\gamma_i\}$ is a monotone curve connecting a and b. Therefore, invariant (I3) holds also for $\mathcal{U}_i$.

Now, suppose that invariant (I2) does not hold for $\mathcal{U}_i$, that is, there exists some pseudosegment β of $\mathcal{A}' \setminus \mathcal{U}_i$ that touches $\text{low}(\mathcal{U}_i)$ from below in an inner point h of β. Refer to Figure 6. Since $\mathcal{U}_i = \mathcal{U}_{i-1} \cup \{\gamma_i\}$ and the curves $\text{low}(\mathcal{U}_{i-1})$ and $\text{low}(\mathcal{U}_i)$ are

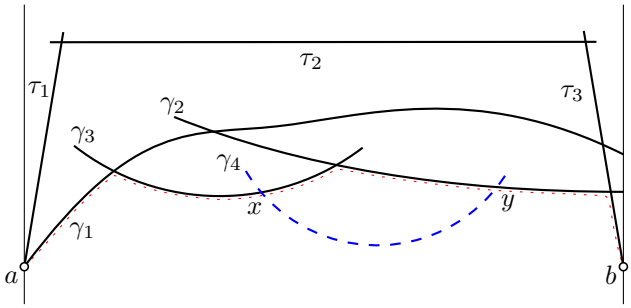


Figure 5: Induction step in the proof of Theorem 1. In the ith step (fourth step in the figure) we add pseudosegment γ_i (dashed) intersecting the lower envelope (dotted) of the previous segments twice. The lower envelope remains a connected curve connecting a with b even after this addition.

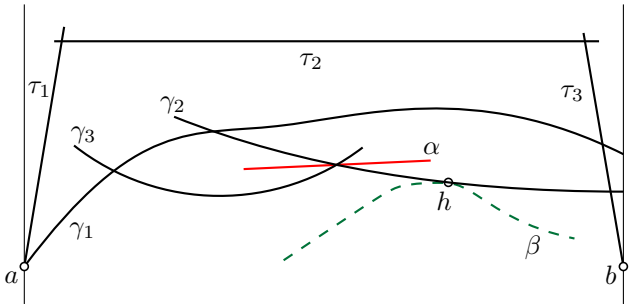


Figure 6: Induction step in the proof of Theorem 1. During the selection of $\mathcal{U}$ some pseudosegments may touch $\text{low}(\mathcal{U})$ from above but never from below. Pseudosegment α touches $\text{low}(\{\tau_1, \tau_2, \tau_3, \gamma_1, \gamma_2, \gamma_3\})$ from above. On the other hand, β cannot be in the same arrangement of pseudosegments since it touches γ_2.

both connected curves containing a and b by invariant (I3), the point h has to be a point of $\text{low}(\mathcal{U}_{i-1})$ or an inner point of γ_i. In the first case, β touches $\text{low}(\mathcal{U}_{i-1})$ in h, an inner point of β, a contradiction with invariant (I2) for $\mathcal{U}_{i-1}$. In the second case, β touches γ_i in h, an inner point of both γ_i and β, a contradiction with the definition of an arrangement of pseudosegments. Hence, invariant (I2) also holds for $\mathcal{U}_i$. Note that the analogous statement for touchings from above does not hold, that is, there may exist some pseudosegment α of $\mathcal{A}' \setminus \mathcal{U}_i$ that both touches $\text{low}(\mathcal{U}_i)$ from above in an inner point of α and touches neither γ_i nor $\text{low}(\mathcal{U}_i)$ in an inner point of α. See Figure 6 for an example of such α.

Finally, suppose that invariant (I1) does not hold, that is, there exists some pseudosegment ρ of $\mathcal{A}' \setminus \mathcal{U}_i$ that together with $\text{low}(\mathcal{U}_i)$ forms a ρ-top bigon. Call s and t the vertices of this bigon and assume $s \prec t$. See Figure 7.

If s and t both lie on $\gamma_i[x, y]$, then ρ and γ_i intersect twice, a contradiction. Otherwise s or t does not lie on $\gamma_i[x, y]$. Without loss of generality assume that t does not lie on

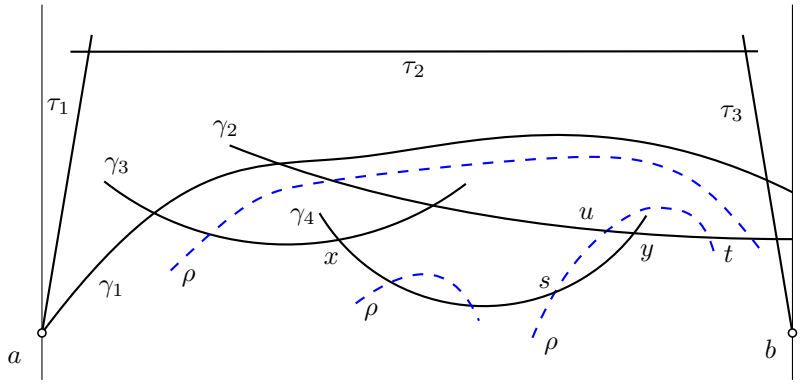


Figure 7: Induction step in the proof of Theorem 1. If there was some pseudosegment ρ that together with the lower envelope ($\mathrm{low}(\{\tau_1, \tau_2, \tau_3, \gamma_1, \gamma_2, \gamma_3, \gamma_4\})$ in the picture) formed a ρ-top bigon, it would either form a ρ-top bigon with the previous lower envelope or intersect twice the segment that was added as the last. In the picture, there are three such possible ρ's.

$\gamma_i[x, y]$ and $y \prec t$. Then s either lies on $\mathrm{low}(\mathcal{U}_{i-1})$ or below it. In both cases $\rho[s, t]$ intersects $\mathrm{low}(\mathcal{U}_{i-1})$ in some point other than t since $\rho[s, t]$ together with $\mathrm{low}(\mathcal{U}_i)$ forms a ρ-top bigon. Denote the rightmost intersection of $\rho[s, t]$ and $\mathrm{low}(\mathcal{U}_{i-1})$ other than t by u. Then $\rho(u, t)$ lies above $\mathrm{low}(\mathcal{U}_{i-1})$ and so $\rho[u, t]$ together with $\mathrm{low}(\mathcal{U}_i)$ forms a ρ-top bigon, a contradiction with invariant (I1) for $\mathcal{U}_{i-1}$. This concludes the proof of Theorem 1.

2.2 Complexity of extending the arrangement

We now discuss how the proof of Theorem 1 can be turned into an algorithm. The outline of the algorithm can already be seen in the proof of Theorem 1. Two main issues remain unresolved. Firstly, how to represent the arrangement, and secondly, how to efficiently find a pseudosegment γ_i of $\mathcal{A}' \setminus \mathcal{U}_{i-1}$ intersecting $\mathrm{low}(\mathcal{U}_{i-1})$ at least twice. We explain what representation we use and how to implement the algorithm in time linear in the number of incidences between pseudosegments and their endpoints or intersection points.

2.2.1 Representation of a monotone arrangement

We assume we have an arrangement $\mathcal{A}$ of n pseudosegments. Combinatorially, the only points of interest on the pseudosegments are the endpoints of the pseudosegments and the intersection points of at least two pseudosegments. From now on, in the description of the algorithm, we will call these points of interest *vertices*. Let m be the total number of incidences between the pseudosegments and the vertices.

For every vertex a, we denote by H_a^+ and H_a^- the open halfplane to the right and left, respectively, of the vertical line containing a. Let x_a be the x-coordinate of a and let x_a^+ and x_a^- be x-coordinates that are slightly larger and slightly smaller, respectively, than x_a so that there are no vertices with x-coordinates in the intervals (x_a^-, x_a) and (x_a, x_a^+). For

the input and output we use the following representation of arrangements of monotone pseudosegments:

- Every pseudosegment corresponds to a unique label
- Every vertex a has its own structure containing its x- and y-coordinate and a list of incident pseudosegments together with the information whether a is a left or right endpoint or an inner point of the pseudosegment. Furthermore, for every pseudosegment α that has a as its left endpoint, the structure of a contains the pair $(\alpha, f_a(\alpha))$ where $f_a(\alpha)$ is the label of the closest pseudosegment below α that intersects both a and H_a^+ (or $f_a(\alpha) = \text{NULL}$ if α is the bottommost pseudosegment with this property). If all pseudosegments incident with a have a as their left endpoint, the structure of a contains the label $f(a)$ of the closest pseudosegment or vertex below a (or $f(a) = \text{NULL}$ if there is none).
- The array L of vertices is ordered in ascending order first by their x-coordinates and then by their y-coordinates. Note that we cannot assume that no two vertices have the same x-coordinate because if we tried to "rotate" the arrangement a bit to avoid this, in the rotated arrangement we could weave a monotone pseudosegment in between the two rotated points, but the corresponding pseudosegment in the original arrangement would not be monotone.

The main operation we use on this representation is *Sweep*, during which we sweep the plane from left to right and maintain the vertical order of pseudosegments in linear time. We do it by processing the vertices in L in the given order. At each vertex a of L we reverse the vertical order of all pseudosegments that contain a as an inner point, we delete all pseudosegments that contain a as their right endpoint, and, finally, we insert each pseudosegment α that has a as its left endpoint to an appropriate position in the order determined by $f_a(\alpha)$ and $f(a)$. We process all vertices with the same x-coordinate together. For any vertex a, we compute the vertical order at coordinates x_a^- and x_a^+, not at the point a itself. Overall, the procedure takes time $O(m)$.

Now we discuss several aspects of this representation. The coordinates of the vertices are not strictly necessary. Instead, it would be sufficient to have a partition of the set of vertices into several classes, a linear ordering of the classes, and a linear ordering of each class. Here each class of the partition would represent a set of vertices with the same x-coordinate. Alternatively, the lexicographic order of L is not strictly needed if the coordinates of the vertices are given. In this case, we could compute the ordering in time $O(m \log(m))$.

2.2.2 Initialization and auxiliary structures

In the input, we need to know the position of the points a and b. We assume we are given their (x, y)-coordinates and also which vertices or pseudosegments are right below a and b. We also assume that $\mathcal{A} = \mathcal{A}[a, b]$, otherwise we would first shorten the arrangement.

Firstly, we construct $\mathcal{A}' = \mathcal{A} \cup \{\tau_1, \tau_2, \tau_3\}$ from $\mathcal{A}$. To add τ_1, we first create a new vertex a if a was not a vertex before. That is straightforward, because in that case there are no incidences with a and to $f(a)$ we assign the vertex or pseudosegment that is right below a, which is a part of the input. Next, we compute the vertical order of pseudosegments at the coordinate x_a^+. The new pseudosegment τ_1 then intersects all pseudosegments in this order above a. It is straightforward to add all these intersections and the new pseudosegment to the representation. We add τ_2 and τ_3 similarly.

We now create a more detailed representation of the arrangement that will help us describe an efficient algorithm. We create the following data structures.

- We define two types of incidences between vertices and pseudosegments. A "+"-incidence $(p, \alpha)^+$ is an incidence between a vertex p and a pseudosegment α that intersect H_p^+. Similarly, a "−"-incidence $(p, \alpha)^-$ is an incidence between a vertex p and a pseudosegment α that intersect H_p^-. We note that if p is an interior point of α, then the incidence between p and α is of both types, "+" and "−". We define the (partial) *horizontal order* on these incidences first by comparing the x-coordinates of their vertices, and in addition, for every pair of incidences $(p, \alpha)^-$ and $(p, \beta)^+$ sharing the same vertex p, the "−"-incidence comes before the "+"-incidence.

- For every pseudosegment α we create an array of "+" and "−"-incidences containing α, linearly ordered according to their horizontal order. For example, a pseudosegment β starting in p, ending in q and containing one more vertex r in the middle is represented by the array $[(p, \alpha)^+, (r, \alpha)^-, (r, \alpha)^+, (q, \alpha)^-]$.

- We use the original structure for vertices but we add two lists into it. For each vertex p we add a list $V^-(p)$ containing pointers to all the "−"-incidences between p and pseudosegments ordered from top to bottom at the x_p^- coordinate. Similarly, we add a list $V^+(p)$ containing pointers to all the "+"-incidences between p and pseudosegments ordered from top to bottom at the x_p^+ coordinate. To each incidence of either form we also add pointers to its location in $V^+(p)$ and $V^-(p)$.

This can be done by a sweep from left to right, so in time $O(m)$.

Note that given two incidences between a point p and two pseudosegments α, β incident with p, we can access any cyclical interval of pseudosegments around p in between α and β in time linear in the size of the interval.

The main part that remains is to describe how to represent the lower envelope in each step of the algorithm, how to select the pseudosegment γ_i effectively, and then how to update the lower envelope. We maintain the following auxiliary structures.

- A linked list X of pointers to the "+" and "−"-incidences on $\mathrm{low}(\mathcal{U}_{i-1})$ linearly ordered according to their horizontal order. To each vertex p we add pointers to the incidences in this list that contain p.

- A linked list Y that for each pseudosegment α from $\mathcal{A}' \setminus \mathcal{U}_{i-1}$ contains the list of all intersections between $\mathrm{low}(\mathcal{U}_{i-1})$ and α that are either proper intersections or

endpoints of α where α continues below low($\mathcal{U}_{i-1}$). From the proof of Theorem 1, we know that any item from Y contains at most two intersections. Together with these intersections, we also store the pointers to their positions in X. For each pseudosegment, we also include a pointer to this list.

- A set of pseudosegments Z from $\mathcal{A}' \setminus \mathcal{U}_{i-1}$ intersecting the lower envelope of $\mathcal{U}_{i-1}$ twice.
- For each $\alpha \in \mathcal{A}'$ we remember whether it is in $\mathcal{U}_{i-1}$.

These structures can be constructed for $\mathcal{U}_0$ in time $O(n)$ because points a and b are incident to at most n pseudosegments and all of the other at most $O(n)$ vertices on τ_1, τ_2, τ_3 are simple intersections of only two pseudosegments.

2.2.3 Implementation of the inductive step

In each inductive step, we select γ_i as an arbitrary pseudosegment from Z and delete it from Z. Let x and y be its intersections with the lower envelope of $\mathcal{U}_{i-1}$ (we get them from Y). Let $(x, \alpha)^+$ be the "+"-incidence between x and low($\mathcal{U}_{i-1}$) and $(y, \beta)^-$ be the "−"-incidence between y and low($\mathcal{U}_{i-1}$). We go along the boundary of the bigon formed by low($\mathcal{U}_{i-1}$) and γ_i and update the data structures X, Y and Z for $\mathcal{U}_{i-1}$ to X, Y and Z for $\mathcal{U}_i$. The implementation is straightforward using the described data structures: we delete the intersection information of low($\mathcal{U}_{i-1}[x, y]$) with incident pseudosegments below it and add the intersection information of $\gamma_i[x, y]$ with incident pseudosegment below it. At point x we need to delete the intersection information with pseudosegments from the incidences from $V^+(x)$ between $(x, \alpha)^+$ and $(x, \gamma_i)^+$. At point y we need to delete the intersection information with pseudosegments from the incidences from $V^-(y)$ between $(y, \beta)^-$ and $(y, \gamma_i)^-$. As noted before we can access any cyclical interval of pseudosegments around any vertex in time linear in the size of the interval. Furthermore, all bigons in the algorithm are disjoint, so during the algorithm we spend at most a constant amount of time for each incidence between a pseudosegment and a vertex.

At the end we get the list X describing the lower envelope of $\mathcal{U}$ and the list Y that contains all the the intersection of low($\mathcal{U}$) with pseudosegments that continue below low($\mathcal{U}$).

2.2.4 Adding the new pseudosegment into the arrangement

We have the linked lists X and Y describing low($\mathcal{U}$) and its intersections with $\mathcal{A}' \setminus \mathcal{U}$ that continue below low($\mathcal{U}$). The new pseudosegment lies slightly below low($\mathcal{U}$). Thus, it intersect exactly the same pseudosegments that are described in Y. The intersections with the new pseudosegment will be slightly below and slightly to the left or right of the intersections with low($\mathcal{U}$). Let p be an old intersection of low($\mathcal{U}$) with a pseudosegment γ from Y with "+" and "−"-incidences $(p, \alpha)^-$ and $(p, \beta)^+$ on low($\mathcal{U}$). To decide whether the new intersection p' below p should be to the right or to the left of p, we just need to find γ either below α in $V(p)^-$ or below β in $V(p)^+$. Given all this information it is straightforward to add the new pseudosegment into the original representation of $\mathcal{A}$ to update the arrangement.

3 Cylindrically monotone arrangements

Recall that we represent the cylinder as the surface $S^1 \times \mathbb{R}$. If we imagine cutting and unrolling the cylinder, we can represent it in the plane as a vertical strip whose left and right sides represent the same vertical line of the cylinder. We can also select the orientation so that the counter-clockwise direction on the cylinder corresponds to the left-to-right direction in the plane. For points $a, b \in S^1$ we denote by $[a, b]$ the counter-clockwise circular arc in S^1 from a to b and we call it an *interval*. For points c and d on the cylinder we call a monotone curve starting in c and ending in d on the cylinder *left-oriented* or *right-oriented* if its projection by the canonical projection map π to S^1 is a clockwise circular arc or a counter-clockwise circular arc, respectively.

3.1 Normal cylindrically monotone arrangements

We prove Corollary 3 using the alternative characterization of normal cylindrically monotone arrangements provided by Proposition 4. Then we prove the proposition itself.

3.1.1 Proof of Corollary 3

Let $\mathcal{A}$ be a normal arrangement of cylindrically monotone pseudosegments. Let a, b be a pair of points on the cylinder but not on the same vertical line and not on the same pseudosegment of $\mathcal{A}$. We need to show that $\mathcal{A}$ is (a, b)-extendable in the family of all normal arrangements of cylindrically monotone pseudosegments. That is, we need to find a cylindrically monotone simple curve γ with endpoints a, b that intersects every curve of $\mathcal{A}$ at most once such that $\mathcal{A} \cup \{\gamma\}$ is still normal.

Let σ be the orthogonal projection from $S^1 \times \mathbb{R}$ to S^1. Let f be the homeomorphism from Proposition 4. We define a homeomorphism $g\colon S^1 \times \mathbb{R} \to S^1 \times \mathbb{R}$ as $g(x, y) = (f(x), y)$. The image $g(\mathcal{A}) = \{g[\alpha];\, \alpha \in \mathcal{A}\}$ of $\mathcal{A}$ under g is a homeomorphic arrangement of cylindrically monotone pseudosegments whose orthogonal projections to S^1 have lengths smaller than π. Furthermore, g maps every vertical line onto a vertical line. Thus, extending $\mathcal{A}$ by a pseudosegment from a to b is equivalent to extending $g(\mathcal{A})$ by a pseudosegment from $g(a)$ to $g(b)$.

We look at the orthogonal projections $\sigma(g(a)), \sigma(g(b))$ of $g(a), g(b)$ to S^1. Without loss of generality we may assume that the length of $[\sigma(g(a)), \sigma(g(b))]$ is smaller than π. Since orthogonal projections of all pseudosegments of $g(A)$ to S^1 have lengths smaller than π, then

$$\{g[\alpha] \cap ([\sigma(g(a)), \sigma(g(b))] \times \mathbb{R})\,;\, \alpha \in \mathcal{A}\} \setminus \{\emptyset\}$$

is an arrangement of cylindrically monotone pseudosegments on the cylindrical strip

$$[\sigma(g(a)), \sigma(g(b))] \times \mathbb{R};$$

note that we remove the empty set since it corresponds to all pseudosegments that have an empty intersection with the cylindrical strip. This arrangement is homeomorphic to an arrangement of monotone pseudosegments in the plane (by unrolling the strip to

the plane). Thus, by Theorem 1, we can extend the arrangement in the plane by a monotone pseudosegment β connecting the images of $g(a)$ and $g(b)$ after unrolling. Let β' be the cylindrically monotone curve on the cylinder corresponding to β. Clearly, β' is a curve connecting $g(a)$ and $g(b)$ and intersecting every pseudosegment from $g(\mathcal{A})$ at most once. Moreover, since the length of $[\sigma(g(a)), \sigma(g(b))]$ is smaller than π, the resulting arrangement is still normal. Thus, $\gamma = g^{-1}(\beta')$ is the desired curve extending the original arrangement $\mathcal{A}$.

Remark. Let us note that we could also find the extending curve γ without the help of Proposition 4 by "cutting" the cylinder along the vertical lines passing through a and b and showing that one of the resulting parts corresponds to the planar version.

3.1.2 Proof of Proposition 4

Let $\mathcal{C}$ be a normal system of n arcs in the unit circle S^1. Without loss of generality, we assume that their endpoints are pairwise distinct. Let $V = \{v_1, v_2, \ldots, v_{2n}\}$ be the set of the $2n$ endpoints of the arcs from $\mathcal{C}$, labeled arbitrarily.

Our first goal is to define, for each endpoint $v \in V$, its "antipodal" point v', so that these new $2n$ points, together with the original $2n$ endpoints from V, are pairwise distinct, no arc from $\mathcal{C}$ contains both v and v', and for every pair $v, w \in V$ with $v \neq w$, the interval $[v, v']$ contains exactly one point from $\{w, w'\}$; that is, the endpoints of the intervals $[v, v']$ and $[w, w']$ alternate on the circle.

Let $\mathcal{C}_0 = \mathcal{C}$. In $2n$ steps, we will define a sequence of points $v'_1, v'_2, \ldots, v'_{2n}$ and a sequence of sets $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_{2n}$. We will make sure the following properties are satisfied:

(P1) Each $\mathcal{C}_i$ is a set of $n + 2i$ intervals; in particular, $\mathcal{C}_i = \mathcal{C} \cup \bigcup_{j=1}^{i} \{[v_j, v'_j), [v'_j, v_j)\}$.

(P2) The points $v_1, v_2, \ldots, v_{2n}, v'_1, v'_2, \ldots, v'_i$ are pairwise distinct.

(P3) No pair of the intervals from $\mathcal{C}_i$ with nonempty intersection covers the whole circle S^1.

By the assumption, all three properties (P1)–(P3) are satisfied for $\mathcal{C}_0$.

Let $i \in \{1, \ldots, 2n\}$. Assume that we have defined $\mathcal{C}_{i-1}$ so that (P1)–(P3) are satisfied for $\mathcal{C}_{i-1}$.

Let

$$J_i = \bigcup \{I \in \mathcal{C}_{i-1}; v_i \in I\} \quad \text{and} \quad J'_i = \bigcap \{S^1 \setminus I; I \in \mathcal{C}_{i-1} \wedge v_i \in I\}.$$

That is, J_i is the union of all intervals from $\mathcal{C}_{i-1}$ that contain v_i, and J'_i is the complement $S^1 \setminus J_i$. We claim that J'_i is an interval of positive length. Indeed, by property (P3) and by the 1-dimensional Helly theorem applied to the intervals $S^1 \setminus I$ where $I \in \mathcal{C}_{i-1}$ and $v_i \in I$, these intervals intersect at a common point. In fact, by properties (P1) and (P2), the endpoints of these intervals are pairwise disjoint, and so their common intersection J'_i is a nontrivial interval.

Let v'_i be an arbitary point from the interior of J'_i that is distinct from all the points $v_1, v_2, \ldots, v_{2n}, v'_1, v'_2, \ldots, v'_{i-1}$, and let $\mathcal{C}_i = \mathcal{C}_{i-1} \cup \{[v_i, v'_i), [v'_i, v_i)\}$. Clearly, $\mathcal{C}_i$ satisfies

properties (P1) and (P2) by construction. Now we verify property (P3). Suppose that one of the intervals $[v_i, v_i'), [v_i', v_i)$, together with some interval $I \in \mathcal{C}_{i-1}$, cover the whole circle. If $v_i \in I$, then $I \subseteq J_i$ and none of the intervals $[v_i, v_i'), [v_i', v_i)$ covers the complement J_i' by the choice of v_i'. If $v_i \notin I$, then some neighborhood $N(v_i)$ of v_i is disjoint with I, but none of the intervals $[v_i, v_i'), [v_i', v_i)$ contains $N(v_i)$. Therefore, property (P3) is satisfied also for $\mathcal{C}_i$.

It remains to define the desired homeomorphism of the circle. Let $w_1, w_2, \ldots, w_{4n}$ be a relabeling of the points $v_1, v_2, \ldots, v_{2n}, v_1', v_2', \ldots, v_{2n}'$ in the counter-clockwise cyclic order around the circle. By property (P3), every pair of points $\{v_i, v_i'\}$ is relabeled as a pair $\{w_j, w_{j+2n}\}$ for some j. We will map the $4n$ points w_j to the vertices of a regular $4n$-gon, keeping their circular ordering.

Define the homeomorphism $f : S^1 \to S^1$ as follows. For every $j \in \{1, 2, \ldots, 4n\}$, let $f(w_j) = e^{2\pi ij/(4n)}$, and interpolate f in each of the intervals $[w_j, w_{j+1}]$ by an arbitrary homeomorphism with the interval $[f(w_j), f(w_{j+1})]$. By the previous observation, each pair $\{v_i, v_i'\}$ is mapped by f to a pair of antipodal points on the circle. Therefore, by property (P3), every arc from $\mathcal{C}$ is mapped by f to an interval of length at most $\pi \cdot (1 - 1/(2n))$.

3.2 General cylindrically monotone arrangements

In this subsection, when we say that an arrangement is extendable, we mean that it is extendable in the family of all cylindrically monotone arrangements.

Unlike normal arrangements, general arrangements of cylindrically monotone pseudosegments are not always extendable, as stated in Proposition 5. Moreover, Theorem 6 states that in general, deciding whether a cylindrically monotone arrangement is (a, b)-extendable is NP-hard.

3.2.1 Proof of Proposition 5

Figure 8 provides an example of an arrangement $\mathcal{A}$ of five pseudosegments and a pair of points a, b such that every cylindrically monotone curve connecting a and b has to intersect some pseudosegment of $\mathcal{A}$ twice. This means that $\mathcal{A}$ is not (a, b)-extendable. This example is the smallest that we know of where the points a and b do not lie on any of the curves of $\mathcal{A}$.

3.2.2 Proof of Theorem 6

Our proof is heavily inspired by a recent proof by Arroyo et al. [4] showing that deciding the extendability of a simple drawing of a graph by one edge between a prescribed pair of points is NP-hard. We adapt their reduction to the more restricted setting of cylindrically monotone drawings.

We prove the NP-hardness by a reduction from 3-SAT. Let $\phi(x_1, \ldots, x_n)$ be a 3-SAT formula with variables $x_1, \ldots, x_n$ and a set of clauses $C = \{C_1, \ldots, C_m\}$. A *literal* is an occurrence of a variable in a clause.

We use the following lemma by Arroyo et al. [4].

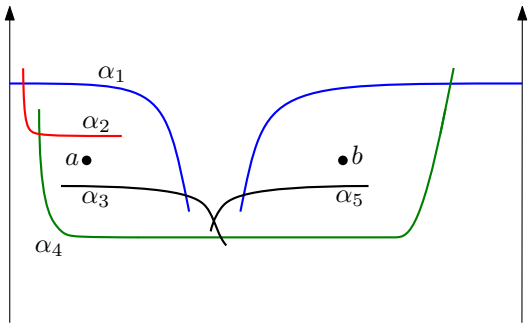


Figure 8: An arrangement of five cylindrically monotone pseudosegments that is not (a, b)-extendable.

Lemma 7 ([4])**.** *The following transformation of a clause with only positive or only negative literals, respectively, preserves the satisfiability of the clause (y is a new variable and* $\texttt{false}_\texttt{p}$, $\texttt{false}_\texttt{n}$ *are constants with value false):*

$$\begin{aligned} x_i \vee x_j \vee x_k &\longmapsto \big(x_k \vee y \vee \texttt{false}_\texttt{n}\big) \wedge \big(x_i \vee x_j \vee \neg y\big) \\ \neg x_i \vee \neg x_j \vee \neg x_k &\longmapsto \big(\neg x_i \vee \neg x_j \vee y\big) \wedge \big(\neg x_k \vee \neg y \vee \texttt{false}_\texttt{p}\big) \end{aligned}$$

We use the transformation from this lemma on every clause containing only positive or only negative literals. We obtain a new 3-SAT formula with the same satisfiability and a modified set of clauses and variables (and new constants $\texttt{false}_\texttt{p}$, $\texttt{false}_\texttt{n}$). For simplicity, we abuse the notation and use the same notation for the new formulas, variables, and clauses. Furthermore, we consider each occurrence of $\texttt{false}_\texttt{p}$ to be a positive literal and each occurrence of $\texttt{false}_\texttt{n}$ to be a negative literal. The new formula $\phi(x_1, \ldots, x_n)$ has exactly three literals in each clause and no clause has only positive or only negative literals.

We construct an arrangement $\mathcal{A}$ of cylindrically monotone pseudosegments and a pair of points a, b, such that $\mathcal{A}$ is (a, b)-extendable if on only if $\phi(x_1, \ldots, x_n)$ is satisfiable. Since 3-SAT is NP-hard the theorem will follow.

We start with an arrangement $\mathcal{A}_1$ of cylindrically monotone pseudosegments depicted in Figure 9. Assume that γ is a cylindrically monotone curve starting in a and ending in b that intersects every pseudosegment of $\mathcal{A}$ at most once. No such γ can be left-oriented because it would cross some pseudosegment of $\mathcal{A}_1$ twice. Hence, γ must be right-oriented. Furthermore, it cannot cross the "right part" of the pseudosegment α (the right component of α in the figure) since it is forced to cross the "left part" of α, and so it has to go above the right part of α and follow the marked direction (green arrows) in Figure 9.

For every occurrence of a literal (even constant one), we will add a new pseudosegment into $\mathcal{A}_1$ starting in variable or constant gadgets and ending in clause gadgets (we will define these gadgets later). Pseudosegments corresponding to positive literals will be left-oriented, and we call them *positive pseudosegments*. Pseudosegments corresponding to negative literals will be right-oriented, and we call them *negative pseudosegments*. All of these new pseudosegments will be horizontal everywhere except for the regions X_{neq} and X_{pos}.

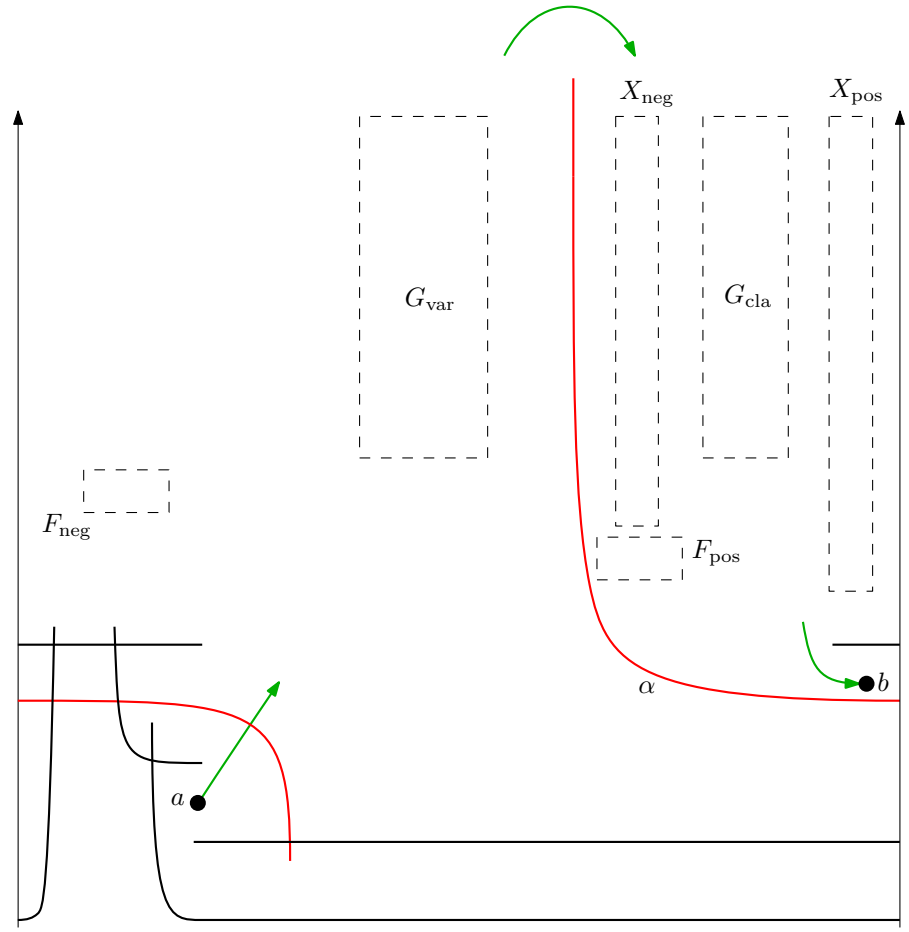


Figure 9: An arrangement $\mathcal{A}_1$ of cylindrically monotone pseudosegments (depicted by solid black and red lines). Any monotone curve joining a and b and extending A_1 would have to be drawn in a way suggested by the (green) arrows. Regions $F_{\text{neg}}, F_{\text{pos}}, G_{\text{var}}, G_{\text{cla}}$ illustrate the approximate positioning of gadgets from the proof: negative constant gadgets, positive constant gadgets, variable gadgets, and clause gadgets, respectively. Regions X_{neg} and X_{pos} illustrate the only areas where pseudosegments corresponding to the negative and positive literals cross, respectively.

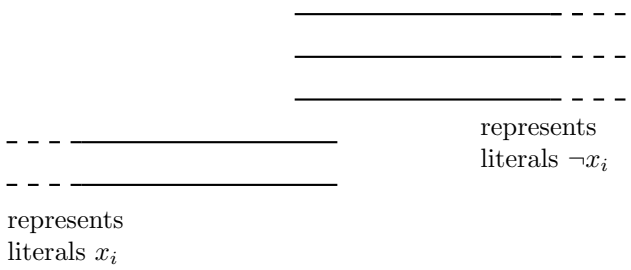


Figure 10: The variable gadget of a variable x_i. Every pseudosegment corresponds to an occurrence of one literal. Pseudosegments corresponding to negative literals continue to the right. Pseudosegments corresponding to positive literals continue to the left.

For each variable x_i we construct the corresponding *variable gadget*. It consists of a neighborhood of starting points of all pseudosegments corresponding to literals x_i or $\neg x_i$, arranged so that the positive pseudosegments continue to the left and are below the negative pseudosegments, that continue to the right. Locally, the variable gadget looks as depicted in Figure 10. Globally, we order the variable gadgets and put them into the region G_{var} depicted in Figure 9, so that the variable gadget corresponding to a variable x_i is above and to the right of the variable gadgets corresponding to variables $x_1, \ldots, x_{i-1}$.

For the constants $\texttt{false}_\texttt{p}$ and $\texttt{false}_\texttt{n}$ we construct two *constant gadgets*. The constant gadget of each of the two constants consists of a neighborhood of starting points of all pseudosegments corresponding to the literals equal to this constant. The starting points of these pseudosegments are arranged on a vertical line and the pseudosegments continue to the left in the case of $\texttt{false}_\texttt{p}$ or to the right in the case of $\texttt{false}_\texttt{n}$. Globally, we put the constant gadget of $\texttt{false}_\texttt{p}$ into the region F_{pos} as depicted in Figure 9 and the constant gadget of $\texttt{false}_\texttt{n}$ into F_{neg}.

For each clause we construct a *clause gadget*. It consists of a neighborhood of the three remaining endpoints of the pseudosegments corresponding to the literals in the clause together with one new auxiliary pseudosegment intersecting the three pseudosegments. There are two types of clause gadgets: one type for clauses with exactly one positive literal and the second type for clauses with exactly two positive literals. The

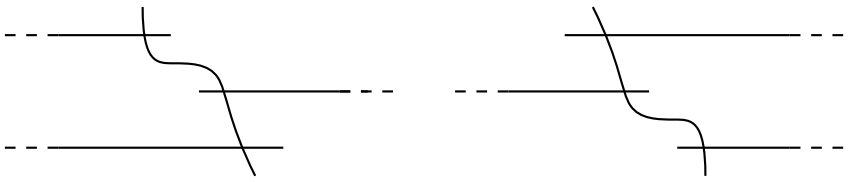


Figure 11: Clause gadgets. Left: a gadget of a clause with one positive literal. Right: a gadget of a clause with two positive literals. Pseudosegments corresponding to negative literals continue to the left. Pseudosegments corresponding to positive literals continue to the right.

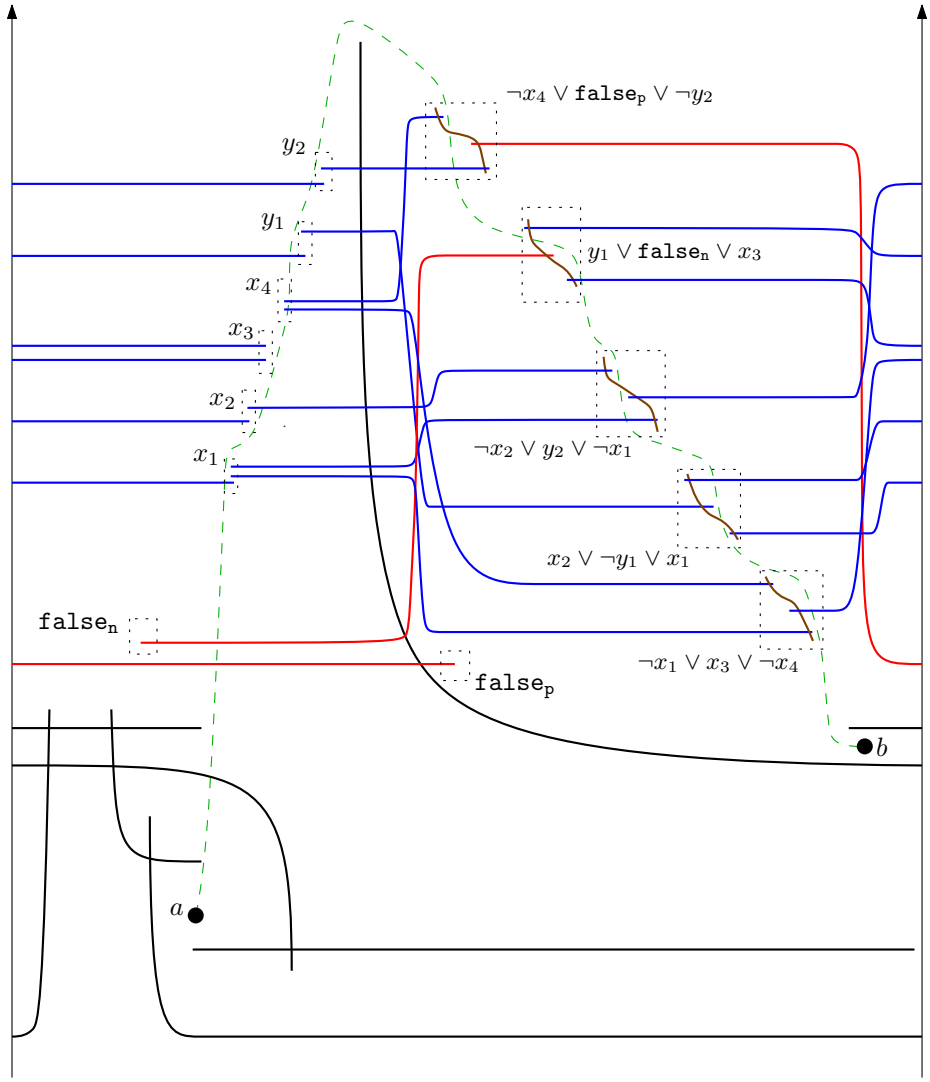


Figure 12: Reduction of the satisfiability of the formula $(\neg x_1 \vee x_3 \vee \neg x_4) \wedge (x_2 \vee \neg y_1 \vee x_1) \wedge (\neg x_2 \vee y_2 \vee \neg x_1) \wedge (y_1 \vee \texttt{false}_\texttt{n} \vee x_3) \wedge (\neg x_4 \vee \texttt{false}_\texttt{p} \vee \neg y_2)$ to the problem of (a, b)-extendability of an arrangement of cylindrically monotone pseudosegments.
Pseudosegments corresponding to the constant literals are red, pseudosegments corresponding to remaining literals are blue, auxiliary pseudosegments of clause gadgets are brown, and pseudosegments of the starting arrangement $\mathcal{A}_1$ are black. A clause is satisfiable if and only if we can extend the arrangement by a new monotone pseudosegment connecting a and b. In the figure, one such pseudosegment is drawn by the dashed green curve corresponding to the satisfying assignment $x_1, y_1, y_2 \longrightarrow \texttt{false}$ and $x_2, x_3, x_4 \longrightarrow \texttt{true}$. Colors refer only to the electronic version.

positive pseudosegments continue to the right, and the negative pseudosegments continue to the left. In the case of a clause with one positive literal, the endpoint of the positive pseudosegment is below and to the right of the endpoint of one negative pseudosegment, and above and to the left of the endpoint of the other negative pseudosegment. The other case is analogous. Locally, the two types look as depicted in Figure 11. Globally, we order the clause gadgets and put them into the region G_{cla} depicted in Figure 9, so that the clause gadget corresponding to a clause C_i is below and to the right of the clause gadgets $C_1, \ldots, C_{i-1}$.

It remains to connect the parts of pseudosegments from variable and constant gadgets with the parts from clause gadgets. We extend the pseudosegments from their endpoints horizontally until they reach region X_{neg} or X_{pos}. Negative pseudosegments reach the region X_{neg}, positive pseudosegments reach X_{pos}. Then we join the corresponding parts inside X_{neg} and X_{pos} in such a way that each pair of pseudosegment cross at most once; see Figure 12.

Clearly, this construction is polynomial since we need only the combinatorial description of the drawing.

We now show the correctness of the construction.

Assume that γ is a cylindrically monotone curve joining a and b that intersects every pseudosegment of $\mathcal{A}$ at most once. Split γ into two parts at the point where it goes above α in between the regions G_{var} and G_{cla}; denote by γ_1 the left part and by γ_2 the right part. By construction, the curve γ_1 has to intersect all pseudosegments corresponding to $\texttt{false}_{\texttt{p}}$ and $\texttt{false}_{\texttt{n}}$. Furthermore, by construction, for each x_i the curve γ_1 has to intersect all pseudosegments corresponding to the positive literals of x_i or all pseudosegments corresponding to the negative literals of x_i. In the former case assign `false` to x_i, in the latter assign `true`. We claim that this is a satisfying assignment. Otherwise, some clause would not be satisfied. Consider the three pseudosegments corresponding to the literals of an unsatisfied clause (all of them are evaluated to `false`). By our construction of the clause gadget, γ_2 has to cross one of these pseudosegments. But γ_1 already crossed this pseudosegment since it crossed all pseudosegments corresponding to literals evaluated to `false`; a contradiction.

On the other hand, assume that there is a satisfying assignment of $\phi(x_1, \ldots, x_n)$. If x_i is evaluated to `false` we draw γ_1 to the left of the variable gadget corresponding to x_i, so that γ_1 does not intersect any pseudosegment corresponding to a negative literal $\neg x_i$. In the other case, we draw γ_1 to the right of the variable gadget corresponding to x_i, so that γ_1 does not intersect any pseudosegment corresponding to a positive literal x_i. This can clearly be done independently for all x_i. We draw γ_2 so that it intersects pseudosegments only in the clause gadgets. Since each clause K is satisfied, there exists a pseudosegment ρ corresponding to a literal of K that is not intersected by γ_1. Then γ_2 can be drawn through the clause gadget of K by intersecting only ρ and (possibly) the auxiliary pseudosegment of the clause gadget. Hence, no pseudosegment that is crossed by γ_2 is crossed by γ_1 and so γ is a valid (a, b)-extension.

Acknowledgements

We thank the anonymous referees for helpful comments. Supported by project 23-04949X of the Czech Science Foundation (GAČR). Jan Soukup also received support from the grant SVV-2023-260699 of Charles University.

References

[1] O. Aichholzer, J. Orthaber and B. Vogtenhuber, Towards crossing-free Hamiltonian cycles in simple drawings of complete graphs, *Comput. Geom. Topol.* **3** (2024), no. 2, Art. 5, 30 pp.

[2] O. Aichholzer, J. Orthaber and B. Vogtenhuber, Separable drawings: extendability and crossing-free Hamiltonian cycles, *Proceedings of the 32nd International Symposium on Graph Drawing and Network Visualization (GD 2024)*, Leibniz International Proceedings in Informatics (LIPIcs) 320, 34:1–34:17, Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2024).

[3] A. Arroyo, J. Bensmail and R. B. Richter, Extending drawings of graphs to arrangements of pseudolines, *J. Comput. Geom.* **12** (2021), no. 2, 3–24.

[4] A. Arroyo, F. Klute, I. Parada, B. Vogtenhuber, R. Seidel and T. Wiedera, Inserting one edge into a simple drawing is hard, *Discrete Comput. Geom.* **69** (2023), no. 3, 745–770.

[5] A. Arroyo, D. McQuillan, R. B. Richter and G. Salazar, Levi's Lemma, pseudolinear drawings of K_n, and empty triangles, *J. Graph Theory* **87** (2018), no. 4, 443–459.

[6] Roger B. Eggleton, Crossing numbers of graphs, PhD thesis, University of Calgary, 1973.

[7] S. Felsner and J. E. Goodman, Pseudoline arrangements, *Handbook of Discrete and Computational Geometry*, Third edition, Edited by Jacob E. Goodman, Joseph O'Rourke and Csaba D. Tóth, Discrete Mathematics and its Applications (Boca Raton), CRC Press, Boca Raton, FL, 2018. ISBN: 978-1-4987-1139. Electronic version: http://www.csun.edu/~ctoth/Handbook/HDCG3.html (accessed May 2023).

[8] R. Fulek and A. J. Ruiz-Vargas, Topological graphs: empty triangles and disjoint matchings, *Proceedings of the twenty-ninth annual symposium on Computational geometry (SoCG '13)*, 259–266, ACM, New York, 2013.

[9] J. E. Goodman, Proof of a conjecture of Burr, Grünbaum, and Sloane, *Discrete Math.* **32** (1980), no. 1, 27–35.

[10] J. Kynčl, Improved enumeration of simple topological graphs, *Discrete Comput. Geom.* **50**(3) (2013), 727–770.

[11] J. Kynčl, J. Pach, R. Radoičić and G. Tóth, Saturated simple and k-simple topological graphs, *Comput. Geom.* **48** (2015), no. 4, 295–310.

[12] J. Kynčl and J. Soukup, Extending simple monotone drawings, *Combinatorial algorithms*, *Lecture Notes in Comput. Sci.* 15885, 17–30, Springer, 2025.

[13] F. Levi, Die Teilung der projektiven Ebene durch Gerade oder Pseudogerade, *Berichte Math.-Phys. Kl. Sächs. Akad. Wiss. Leipzig* **78** (1926), 256–267.

[14] M.C. Lin and J. L. Szwarcfiter, Characterizations and recognition of circular-arc graphs and subclasses: a survey, *Discrete Math.* **309** (2009), no. 18, 5618–5635.

[15] M. Schaefer, A proof of Levi's Extension Lemma, arXiv:1910.05388v1 (2019).

[16] J. Snoeyink and J. Hershberger, Sweeping arrangements of curves, *Discrete and computational geometry (New Brunswick, NJ, 1989/1990)*, *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.* 6, 309–349, Amer. Math. Soc., Providence, RI, 1991.