

Non-crossing Connectors in the Plane

Jan Kratochvíl

honza@kam.mff.cuni.cz

Torsten Ueckerdt

ueckerdt@googlemail.com

Charles University in Prague,
Prague, Czech Republic

March 1, 2018

Abstract

We consider the non-crossing connectors problem, which is stated as follows: Given n simply connected regions R_1, \dots, R_n in the plane and finite point sets $P_i \subset R_i$ for $i = 1, \dots, n$, are there non-crossing connectors γ_i for (R_i, P_i) , i.e., arc-connected sets γ_i with $P_i \subset \gamma_i \subset R_i$ for every $i = 1, \dots, n$, such that $\gamma_i \cap \gamma_j = \emptyset$ for all $i \neq j$?

We prove that non-crossing connectors do always exist if the regions form a collection of pseudo-disks, i.e., the boundaries of every pair of regions intersect at most twice. We provide a simple polynomial-time algorithm if the regions are axis-aligned rectangles. Finally we prove that the general problem is NP-complete, even if the regions are convex, the boundaries of every pair of regions intersect at most four times and P_i consists of only two points on the boundary of R_i for $i = 1, \dots, n$.

1 Introduction

Connecting points in the plane in a non-crossing way is a natural problem in computational geometry. For example, it is well known that one can always find a non-crossing matching with straight line segments between any set of n red and n blue points in the plane. Related problems ask for a non-crossing path that alternates between red and blue vertices, or two non-crossing spanning trees, one on each color, which minimize the number of crossings between them. We refer to the survey article of Kaneko and Kano [13] for more on red and blue points. Recent investigations consider the problem of finding a non-crossing matching between n ordered pairs of points sets, e.g., between a single point and a vertical line [2, 1], or between two vertical line segments [25].

In this paper, we investigate what happens if we allow general curves instead of just straight line segments. Moreover, we want to connect not only n pairs of points, but n finite sets of points. Of course, we can always find such non-crossing curves, unless two point sets intersect. But if we impose for every point set a region that the corresponding curve must be contained in, then determining whether or not such non-crossing curves exist becomes a non-trivial problem.

Related work. A lot of research has been done for connecting points in the plane with straight line segments in a non-crossing way. For instance, a point set P is *universal* for a class \mathcal{G} of planar graphs if the vertices of every graph $G \in \mathcal{G}$ can be embedded onto the points P such that straight edges do not cross. It has been shown [21, 4] that every set of n points in general position is universal for the class of n -vertex outer-planar graphs. The smallest point set that is universal for the class \mathcal{G}_n of all n -vertex planar graphs consists of at least $1.235n$ [5, 18] and at most $\mathcal{O}(n^2)$ [11, 23] points. Deciding whether a given graph embeds on a given point set, is known to be NP-complete [3].

In many variants edges are allowed to be more flexible than straight line segments. Kaufmann and Wiese [15] show that every n -element point set is universal for \mathcal{G}_n if every edge is a polyline with at most 2 bends. Moreover, some n -element point sets are universal for \mathcal{G}_n if edges bend at most once [9]. If the bending points have to be embedded onto P as well, then universal sets of size $\mathcal{O}(n^2/\log n)$ for 1 bend, $\mathcal{O}(n \log n)$ for 2 bends, and $\mathcal{O}(n)$ for 3 bends are known [7]. A variant with so-called ortho-geodesic edges was studied especially for trees of maximum degree 3 and 4 [6].

Variants where every vertex v has an associated subset P_v of P of its possible positions has been studied for much simpler graphs like matchings and cycles, however for straight edges only. For cycles, deciding whether a set of non-crossing edges exists is known to be NP-complete, even if every P_v is a vertical line segment or every P_v is a disk [20]. For matchings, NP-completeness has been shown if $|P_v| \leq 3$ [2, 1], or P_v is a vertical line segment of unit length [25]. The latter result still holds if every edge $\{u, v\}$ may be a monotone curve within the convex hull of $P_u \cup P_v$ [24].

Our results. In this paper we allow edges to be general curves, which w.l.o.g. can be thought of as polylines of finite complexity, that is with a finite number of bends. Given an N -element point set P , each curve γ_i is asked to go through (connect) a fixed subset P_i of P of two or more points. Any two such curves, called *connectors*, shall be non-crossing, i.e., have empty intersection. In particular, w.l.o.g. P is partitioned into subsets P_1, \dots, P_n and we ask for a set of n non-crossing connectors $\gamma_1, \dots, \gamma_n$ with $P_i \subset \gamma_i$ for $i = 1, \dots, n$. It is easily seen that the order in which γ_i visits the points P_i may be fixed arbitrarily. Indeed, we could ask to embed any planar graph G_i with $|P_i|$ vertices and curved edges onto P_i , even while prescribing the position of every vertex in G_i .

Non-crossing connectors as described above do always exist. But the situation gets non-trivial if we fix subsets R_1, \dots, R_n in the plane, called *regions*, and impose $\gamma_i \subset R_i$ for every $i = 1, \dots, n$. Throughout this paper every R_i is a simply connected region in the plane, which contains P_i .

In Section 3 we prove that non-crossing connectors do always exist if the given regions form a collection of pseudo-disks, i.e., the boundaries of every pair of regions intersect at most twice. In Sections 4 and 5 we consider the computational complexity of deciding whether or not non-crossing connectors exist. In particular, in Section 4 we show that the problem is polynomial if the regions are axis-aligned rectangles, while in Section 5 we prove that the problem is NP-complete, even if the regions are convex, the boundaries of every pair of regions intersect at most four times, and $|P_i| = 2$ for every $i = 1, \dots, n$. We start with some notation in Section 2 and show that the non-crossing connectors problem is in NP.

2 The non-crossing connectors problem

The non-crossing connectors problem is formally defined as follows.

NON-CROSSING CONNECTORS

Given: Collection R_1, \dots, R_n of simply connected subsets of the plane and a finite point set $P_i \subset R_i$, for $i = 1, \dots, n$ with $P_i \cap P_j = \emptyset$ for $i \neq j$.

Question: Is there a collection $\gamma_1, \dots, \gamma_n$ of curves, such that $P_i \subset \gamma_i \subset R_i$ for $i = 1, \dots, n$ and $\gamma_i \cap \gamma_j = \emptyset$ for $i \neq j$?

The boundary of every region R_i is a simple closed curve, denoted by ∂R_i . We assume here and for the rest of the paper that $\partial R_i \cap \partial R_j$ is a finite point set. We may think of $\bigcup_{i=1}^n \partial R_i$ as an embedded planar graph $G = (V, E)$ with vertex set $V = \{p \in \mathbb{R}^2 \mid p \in \partial R_i \cap \partial R_j, i \neq j\}$ and edge set $E = \{e \subset \mathbb{R}^2 \mid e \text{ is a connected component of } \bigcup \partial R_i \setminus V\}$. A point $p \in \partial R_i \cap \partial R_j$ is either a *crossing point* or a *touching point*, depending on whether the cyclic order of edges in ∂R_i and ∂R_j around p is alternating or not. We say that two regions R_i, R_j are k -*intersecting* for $k \geq 0$ if $|\partial R_i \cap \partial R_j| \leq k$ and all these points are crossing points, i.e., w.l.o.g. k is even. A set R_1, \dots, R_n

of regions is k -intersecting if this is the case for any two of them. For example, R_1, \dots, R_n are 0-intersecting if and only if¹ they form a nesting family, i.e., $R_i \cap R_j \in \{\emptyset, R_i, R_j\}$ for $i \neq j$.

Regions R_1, \dots, R_n are called a *collection of pseudo-disks* if and only if they are 2-intersecting. Usually two pseudo-disks may have one touching point. However, this can be locally modified into two crossing points with affecting the existence of non-crossing connectors. Pseudo-disks for example include homothetic copies of a fixed convex point set, but they are not convex in general. A collection of axis-aligned rectangles is always 4-intersecting, but not necessarily 2-intersecting. Finally, if R_1, \dots, R_n are convex polygons with at most k corners, then they are $2k$ -intersecting.

We close this section by showing that the non-crossing connectors problem belongs to NP.

Proposition 2.1 NON-CROSSING CONNECTORS is in NP.

Proof: We reduce our problem to Weak Realizability of Abstract Topological Graphs. Abstract topological graphs (AT-graphs, for short) have been introduced in [17] as triples (V, E, R) where (V, E) is a graph and R is a set of pairs of edges of E . The AT-graph (V, E, R) is weakly realizable if (V, E) has a drawing (not necessarily non-crossing) in the plane such that $ef \in R$ whenever the edges e, f cross in the drawing. Weak realizability of AT-graphs is NP-complete. The NP-hardness was shown in [16], and the NP-membership was shown relatively recently in [22].

We assume the input of the problem be described as a plane graph G (the boundaries of the regions) with the incidence structure of the points of P to the faces of the graph. The size of the input is measured by the number of crossing points of the boundaries plus the number of regions and points to be connected. Strictly speaking G may contain nodeless loops corresponding to the boundaries of regions which are not crossed by any other boundary. Such a case will be handled quickly by the next construction. We refer to Figure 1 for an illustrative example.

Add vertices and edges to G to create a vertex 3-connected supergraph G' of G . This can be achieved for instance by subdividing every edge of G by at most 3 new extra vertices (subdivide nodeless loops by 3 vertices, loops by 2 vertices and simple edges by a single vertex each so that to avoid multiple edges in G'), adding a new vertex in each face adjacent to all vertices of this face (both the original and the new ones) and adding edges connecting the subdividing vertices to create a triangulation. Such a graph G' is 3-connected and hence it has a topologically unique non-crossing drawing in the plane (up to the choice of the outerface and its orientation).




Figure 1: From Non-crossing Connectors to Weak Realizability of AT-graphs.

For every point of P , choose a triangular face of G' in the face of G it lies in, and regard it as a vertex adjacent to the vertices of this triangular face (if more points of P are assigned to the same face of G' , we add edges between them and to the vertices of the face to triangulate it). Call the resulting planar triangulation G'' . As the last step, for every i , add edges that connect the points of P_i by a path, and call the edges of these paths the *connecting edges*. Call the resulting graph \tilde{G} .

¹To be precise, we always require for the “if”-part that $\partial R_i \cap \partial R_j$ is a finite set of crossing points for $i \neq j$.

Now define an AT-graph with underlying graph \tilde{G} by allowing the connecting edges of points P_i to cross anything but other connecting edges and the edges resulting from the boundary of R_i . No other edges are allowed to cross, in particular, no edges of G'' may cross each other. It is straightforward to see that a weak realization of this AT-graph is a collection of non-crossing connectors, and vice versa. Thus the NP-membership follows by the result of Schaefer et al. [22] and the fact that our construction of \tilde{G} is polynomial (if G has c crossing points, a arcs, and f faces, and p is the total number of points in P , then \tilde{G} has at most $c + 3a + f + p$ vertices and at most $3c + 9a + 3f + 4p$ edges). \square

3 Pseudo-Disks

In this section we prove that non-crossing connectors do always exist if the given regions form a collection of pseudo-disks. We begin with an auxiliary lemma.

Lemma 3.1 *Let R, R' be pseudo-disks, $p \in R \setminus R'$ be a point, and $\gamma \subset R \cap R'$ a curve that intersects ∂R exactly twice. Then the connected component of $R \setminus \gamma$ not containing p is completely contained in the interior of R' .*

Proof: Let C be the connected component of $R \setminus \gamma$ not containing p . Let q, r be the intersections of γ with ∂R . Then $\partial R \setminus \{p, q, r\}$ is a set of three disjoint curves. The curve between p and q as well as between p and r contains a point in $\partial R'$, since $p \notin R'$ and $q, r \in \gamma \subset R'$. Because R, R' are pseudo-disks q and r are the only points in $\partial R \cap \partial R'$ and hence the third curve $\delta = C \cap \partial R$ between q and r is completely contained in R' . Since the closed curve $\delta \cup \gamma \subset R'$ is the boundary of C and $\delta \cap \partial R' = \emptyset$, we conclude that C is completely contained in the interior of R' . \square

Theorem 1 *If R_1, \dots, R_n is a collection of pseudo-disks, then non-crossing connectors exist for any finite point sets $P_i \subset R_i$ ($i = 1, \dots, n$) with $P_i \cap P_j = \emptyset$ for $i \neq j$.*

Proof: The proof is constructive. Let R_1, \dots, R_n be a collection of pseudo-disks and P_i a finite subset of R_i for $i = 1, \dots, n$. We assume w.l.o.g. that every ∂R_i is a closed polygonal curve of finite complexity. Moreover, we assume that the regions are labeled from $1, \dots, n$, such that for every $i = 2, \dots, n$ the set $R_i \setminus (R_1 \cup \dots \cup R_{i-1})$ is non-empty, i.e., contains some point p_i . For example, we may order the regions by non-decreasing x -coordinate of their rightmost point. Note that rightmost points of pseudo-disks do not coincide. For simplicity we add p_i to P_i for every $i = 1, \dots, n$ (and denote the resulting point set again by P_i). Clearly, every collection of non-crossing connectors for the new point sets is good for the original point sets, too.

We start by defining a connector γ_1 for (R_1, P_1) arbitrarily, such that $P_1 \subset \gamma_1 \subset R_1$, $\gamma_1 \cap P_i = \emptyset$ for every $i \geq 2$, and γ_1 is a polyline of finite complexity. To keep the number of operations in the upcoming construction finite we consider polylines of finite complexity only. That is, whenever we define a curve we mean a polyline of finite complexity even if we do not explicitly say so.

For $i = 2, \dots, n$ assume that we have non-crossing connectors $\gamma_1, \dots, \gamma_{i-1}$, such that $(\bigcup_{j < i} \gamma_j) \cap (\bigcup_{k \geq i} P_k) = \emptyset$. We want to define a connector γ_i for (R_i, P_i) . The set $R_i \setminus (\bigcup_{j < i} \gamma_j)$ has finitely many connected components $\{C_k\}_{k \in K}$ with $|K| < \infty$. Every point in P_i is contained in exactly one C_k . Let C_0 be the component containing the additional point $p_i \in P_i$. The informal idea is the following. We reroute some of the existing connectors until P_i is completely contained in C_0 . Then, we define a connector γ_i for (R_i, P_i) arbitrarily, such that $P_i \subset \gamma_i \subset C_0$, as well as $\gamma_i \cap P_j = \emptyset$ for $j > i$ and $\gamma_i \cap \gamma_j = \emptyset$ for $j < i$. The reader may consider Figure 2 for an illustration of the upcoming operation. For better readability the parts of connectors are omitted in Figure 2, which

have an endpoint in the interior of R_i . However, those, as well as the point sets P_j with $j > i$, will be circumnavigated by the curve δ .




Figure 2: Rerouting the curve γ bordering the components C and C'' , such that the subset of P_i formerly contained in C is contained in C'' afterwards.

Every connector γ_j for $j < i$ is a simple curve. Thus every connected component of $R_i \setminus (\bigcup_{j < i} \gamma_j)$ contains a point from the boundary of R_i , i.e., the adjacency graph between the components is a tree T on vertex set $\{C_k\}_{k \in K}$, which we consider to be rooted at C_0 . Let $C \neq C_0$ be a component, such that $C \cap P_i \neq \emptyset$ but $C' \cap P_i = \emptyset$ for every descendant C' of C in T . Let γ be the curve in R_i that forms the border between C and its father C'' in T , i.e., γ intersects ∂R_i only at its endpoints and is a subset of some connector γ_{j^*} for (R_{j^*}, P_{j^*}) . In particular, $j^* < i$ and hence $p_i \notin R_{j^*}$. Applying Lemma 3.1 with $p = p_i$ we get that C is contained in the interior of R_{j^*} .

Let $q \notin P_{j^*}$ be any interior point of γ and δ be any curve with endpoint q , such that $(P_i \cap C) \subset \delta \subset (C \cup \{q\}) \subset R_{j^*}$, as well as $\delta \cap (\bigcup_{j < i} \gamma_j) = \{q\}$ and $\delta \cap (\bigcup_{j \neq i} P_j) = \emptyset$. We reroute γ_{j^*} within a small distance around δ . More formally, define a simply connected set $D \supset \delta$ to be a thickening of the curve δ by some small $\varepsilon > 0$, such that D is still contained in $R_{j^*} \setminus (\bigcup_{j \neq i} P_j \cup \bigcup_{j < i} \gamma_j)$. Note that $D \not\subset C$ if P_i (and hence δ) contains points on the boundary of R_i . However, we can ensure that $D \subset R_{j^*}$ since C lies in the interior of R_{j^*} . Moreover, we can choose ε small enough, such that ∂D intersects γ only in two points q_1 and q_2 , which are ε -close to q .

Next, the part of γ between q_1 and q_2 is replaced by the part of ∂D between q_1 and q_2 that runs through C . This rerouting of γ (and implicitly the connector γ_{j^*}) may (or may not) change the subtree of T rooted at C'' . But it does not affect any component of $R_i \setminus \bigcup_{j < i} \gamma_j$ that is not in this subtree. Moreover, C'' is extended by $D \cap C$, which contains all points in $P_i \cap C$. Hence, the so-to-speak total distance of the points in P_i from C_0 in T is decreased. After finitely many steps we have $P_i \subset C_0$ and thus can define the connector γ_i for (R_i, P_i) . \square

4 Axis-Aligned Rectangles

Throughout this section the regions R_1, \dots, R_n are given as axis-aligned rectangles. Whenever we consider some axis-aligned rectangle R we consider it as a closed set, i.e., $R = [x_1, x_2] \times [y_1, y_2]$ for some $x_1 < x_2$ and $y_1 < y_2$. In particular, we always take the closure without explicitly saying so.

Definition 4.1 *Two rectangles R_i, R_j form a cross (and R_i and R_j are called crossing) if they are 4-intersecting, and form a filled cross if additionally both connected components of $R_i \setminus R_j$ contain a point from P_i , and both connected components of $R_j \setminus R_i$ contain a point from R_j .*

Note that if R_i and R_j are crossing, then every connected component C of $(R_i \cup R_j) \setminus (R_i \cap R_j)$ is an axis-aligned rectangle, and hence we consider its closure in the above definition. Obviously, non-crossing connectors do not exist if some pair of rectangles is a filled cross. In other words, the absence of filled crosses is a necessary condition for the existence of non-crossing connectors. Next, we show that this condition is also sufficient.

Theorem 2 *A set of axis-aligned rectangles admits a set of non-crossing connectors if and only if it does not contain a filled cross.*

Proof: The “only if”-part is immediate. We prove the “if”-part by applying Theorem 1. To this end we consider axis-aligned rectangles R_1, \dots, R_n such that no two of them form a filled cross. If there is no cross at all, then the rectangles are pseudo-disks and non-crossing connectors exist by Theorem 1. So assume that some pair of rectangles is a cross, but *not* a filled cross. Consider such a cross $\{R_i = [x_1^i, x_2^i] \times [y_1^i, y_2^i], R_j = [x_1^j, x_2^j] \times [y_1^j, y_2^j]\}$ where $R_i \cap R_j$ is inclusion-minimal among all crosses. W.l.o.g. assume that $R_i \cap R_j = [x_1^j, x_2^j] \times [y_1^i, y_2^i]$ and the connected component $C = [x_1^j, x_2^j] \times [y_2^i, y_2^j]$ of $R_j \setminus R_i$ contains no point from P_j . The situation is illustrated in Figure 3. Figuratively speaking we chop off C (actually a slight superset C' of C) from R_j in order to reduce the total number of crosses of all rectangles. More precisely, choose $\varepsilon > 0$ small enough that $C' := [x_1^j, x_2^j] \times [y_2^i - \varepsilon, y_2^i]$ contains no point from P_j , and that the y -coordinate of no corner of a rectangle $\neq R_i$ lies between $y_2^i - \varepsilon$ and y_2^i . We replace R_j by $\tilde{R}_j := R_j \setminus C'$.




Figure 3: Chopping off C' with $C' \cap P_j = \emptyset$ from the rectangle R_j to obtain \tilde{R}_j .

We claim that every rectangle R_k crosses \tilde{R}_j only if it crosses R_j , too. Indeed, every new intersection point of $\partial\tilde{R}_j$ with ∂R_k lies on the segment $[x_1^j, x_2^j] \times (y_2^i - \varepsilon)$. If $\{R_k, \tilde{R}_j\}$ is a cross, there are two further intersection points of ∂R_k with $\partial\tilde{R}_j$ on the segment $[x_1^j, x_2^j] \times y_1^j$. By the choice of ε and to position of R_i the pair $\{R_i, R_k\}$ would be a cross, too, but with $R_i \cap R_k \subset R_i \cap R_j$, which contradicts the inclusion-minimality of $R_i \cap R_j$ and proves our claim.

We proved that no cross is created by the chopping operation described above, and thus there is no *filled* cross among the rectangles $R_1, \dots, R_{j-1}, \tilde{R}_j, R_{j+1}, \dots, R_n$. Since $\{R_i, \tilde{R}_j\}$ is not a cross while $\{R_i, R_j\}$ is one, it follows that the total number of crosses has decreased by at least one. Repeating the procedure at most $\binom{n}{2}$ times finally results in a collection of axis-aligned rectangles, which are subsets of the original rectangles, and contain no cross at all. By Theorem 1 non-crossing connectors exist for the smaller rectangles, which are good for the original rectangles, too. \square

Corollary 4.1 *It can be tested in $\mathcal{O}(n^2)$ whether or not a set of n axis-aligned rectangles admits a set of non-crossing connectors.*

Proof: By Theorem 2 we only have to check for every pair of rectangles whether they form a cross. If so the answer is 'No' and if not the answer is 'Yes'. \square

5 NP-Completeness

In this section we prove NP-completeness of the non-crossing connectors problem. By Proposition 2.1 the problem is in NP. We prove NP-hardness, even if the regions and their point sets are

very restricted. Let us remark that most of the technicalities in the upcoming reduction, including the use of zones and segment gadgets, are due to the fact that we use *convex* regions only. Dropping convexity but keeping all the other restrictions allows for a much shorter and less technical proof. However, due to space limitations we present only the more technical reduction with convex sets.

Theorem 3 *The non-crossing connectors problem is NP-complete, even if the regions are 4-intersecting convex polygons with at most 8 corners and for every $i = 1, \dots, n$ the set P_i consists of only two points on the boundary of R_i .*

We prove Theorem 3 by a polynomial reduction from PLANAR 3-SAT. In the 3-SAT problem we are given a formula ψ in conjunctive normal form where each clause has at most 3 literals, i.e., positive or negated variables. The *formula graph* G_ψ is the bipartite graph whose vertex set is the union of clauses and variables, and whose edge set E is given by $\{x, c\} \in E$ if and only if variable x appears in clause c . The PLANAR 3-SAT problem is the 3-SAT problem with the additional requirement that G_ψ is a planar graph. It is known [19] that PLANAR 3-SAT is NP-complete, even if every variable appears in at most 3 clauses, i.e., G_ψ has maximum degree 3 [10].

Zones. We start by defining a collection of zones, i.e., polygons in the plane, which will later contain our gadgets. It has been proved several times [14, 8, 12] that a planar graph with maximum degree 3 admits a straight line embedding in which every edge has one of the *basic slopes* $0^\circ, 30^\circ$ or 60° , and such an embedding can be computed in linear time. To be precise, possibly three edges on the outer face will have a bend, i.e., not be straight. We take such an embedding of G_ψ and thicken it so that vertices are represented by disks with diameter ε and edges are rectangles with one side length ε , for some $\varepsilon > 0$ small enough. Let T be an equilateral triangle, smaller than an ε -disk, whose sides have basic slopes and with a tip pointing up. For every clause c let $T(c)$ be a copy of T centered at the position of vertex c in G_ψ . We associate the corners of $T(c)$ with the variables in c in the same clockwise order as the corresponding incident edges at vertex c in G_ψ . If c has size 2, i.e., consists of two variables, then one corner of $T(c)$ is not associated with a variable.

W.l.o.g. every variable x appears at least once positive and once negated in ψ . First, let x be a variable that is contained in only two clauses c_1 and c_2 . We connect the corresponding corners of $T(c_1)$ and $T(c_2)$ by a polyline contained in the thickened edges of G_ψ , such that each of the following holds. See Figure 4 for an illustration.

- a) The polyline consists of only constantly many segments, say at most 10.
- b) Every segment has a basic slope, i.e., $0^\circ, 30^\circ$, or 60° .
- c) The angle between any two consecutive segments is 60° .
- d) The first and last segment is attached to the corresponding corner of $T(c_1)$ and $T(c_2)$, respectively, as depicted for the clause c in Figure 4.

Let x be a variable that is contained in three clauses c_1, c_2, c_3 . Let c_1 and c_2 be the clauses in which x appears with the same sign, i.e., both positive or both negated. We introduce a polyline connecting the corresponding corners of $T(c_1)$ and $T(c_3)$ in the same way as above, i.e., such that **a)–d)** holds. W.l.o.g. assume that one segment s of the polyline is completely contained in the ε -disk for vertex x in G_ψ . We introduce a second polyline, which is contained in the thickened edge $\{x, c_2\}$ in G_ψ , starts at the endpoint of s that is closer to c_3 and ends at the corner of $T(c_2)$ corresponding to x . This polyline shall again satisfy **a)–d)**, where in **d)** only the last segment is considered, while the first segment shall be one half of segment s . See Figure 4 for an illustration.




Figure 4: The definition of polylines and zones.

We denote the two endpoints of every segment by A and B , such that whenever two segments s, s' share an endpoint and are not contained in each other, then this endpoint is denoted differently in s and s' . Finally we define a zone for every clause and every segment of a polyline. The zone $Z(c)$ of a clause c is an equilateral triangle slightly larger than $T(c)$. The zone $Z(s)$ of a segment s is a thin and long rectangle containing s and with two sides parallel to s . How two zones intersect is given in the right of Figure 4. In the particular case that one segment s is contained in another segment s' we define the zones such that $Z(s)$ is contained in $Z(s')$ as well. This completes the definition of zones. In the remainder of this proof we neither need ε -disks, nor polylines any more. We proceed by defining one gadget, i.e., a set of regions, for every zone.

Clause Gadget. We define the clause gadget, which consists of 5 regions as depicted in Figure 5. For better readability not all regions are drawn convex in the figure. However, for the actual reduction we use the combinatorially equivalent *convex* regions depicted in Figure 7.

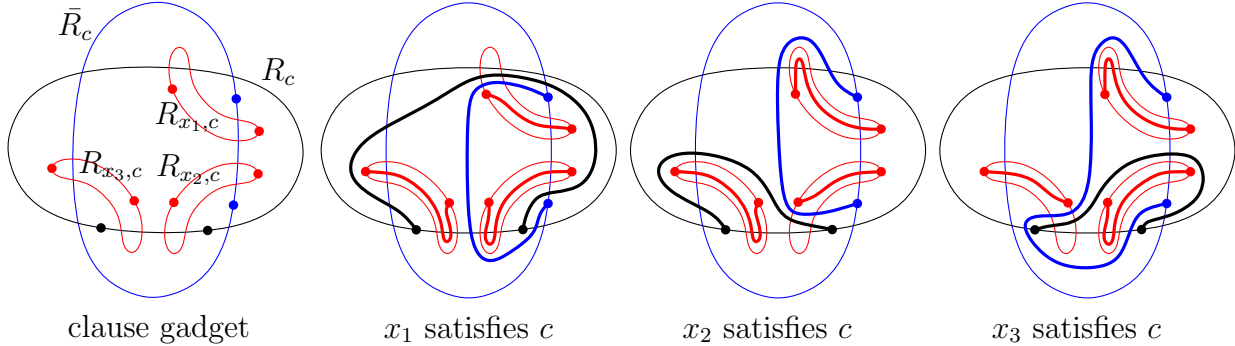


Figure 5: The clause gadget.

For every clause c we define a black region $R_{black}(c)$ and a blue region $R_{blue}(c)$, which are 4-intersecting. The colors are added just for better readability of the figures. We color the 2-element point set corresponding to every region in the same color as the region. Assume c has size 3. We define 3 pairwise disjoint red regions $R_{red}(x, c)$, one for each variable x in c , such that the regions appear in the same clockwise order as the edges $\{x, c\}$ around c in G_ψ . For every pair $\{x, c\}$ of a variable x and a clause c containing x the red region $R_{red}(x, c)$ has one component inside the black region $R_{black}(c)$, which contains both red points, and one outside $R_{black}(c)$, which does not contain a red point. If the connector $\gamma_{red}(x, c)$ for $R_{red}(x, c)$ is completely contained in $R_{black}(c)$, we say that *variable x satisfies clause c* . If the clause has size 2, then only two of the red regions are associated with the variables. Moreover we put the point of the third red region, which is contained in the blue region $R_{blue}(c)$, anywhere outside the black region $R_{black}(c)$ instead of inside $R_{black}(c)$. Hence this “artificial variable” can not satisfy the clause.

It is not difficult to see that the next lemma holds, i.e., Figure 5 verifies the second part of it.

Lemma 5.1 *In every set of non-crossing connectors for the 5 regions of a clause gadget at least one variable satisfies the clause. Moreover, non-crossing connectors do exist as soon as one variable satisfies the clause.*

Segment Gadget. For every segment of a polyline we define another gadget, which is very similar to the clause gadget described above and depicted in Figure 6. Actually, we define two possible gadgets, where the second arises from the first by vertically mirroring the left half of it. For every segment s we again define two 4-intersecting regions, $R_{black}(s)$ is black and $R_{blue}(s)$ is blue. We further define two disjoint red regions $R_{red}(A, s)$ and $R_{red}(B, s)$, associated with the endpoints A, B of the segment s . Again, each of $R_{red}(A, s), R_{red}(B, s)$ is divided into the part inside $R_{black}(s)$, which contains both red points, and the part outside $R_{black}(s)$. We say that *endpoint A or B satisfies the segment s* if the corresponding connector $\gamma_{red}(A, s)$ or $\gamma_{red}(B, s)$ is completely contained in $R_{black}(s)$. We deduce a statement similar to Lemma 5.1.




Figure 6: The two possible segment gadgets.

Lemma 5.2 *In every set of non-crossing connectors for the 4 regions of a segment gadget at least one endpoint satisfies the segment. Moreover, non-crossing connectors do exist as soon as one endpoint satisfies the segment.*

Putting things together. We finally describe the regions corresponding to the formula ψ . For every clause c we define a clause gadget as depicted in the left of Figure 7, which is completely contained in the zone $Z(c)$ corresponding to c . Similarly, for every segment s we define a segment gadget as depicted in the right of Figure 7, which is completely contained in the zone $Z(s)$ corresponding to s . The middle part of a segment gadget, which is highlighted in Figure 7, is stretched such that the gadget reaches from one end of the zone to the other. Note that both gadgets are combinatorially equivalent to the ones in Figure 5 and Figure 6, and consist of solely convex polygons with at most 8 corners.

We let two gadgets intersect as specified in Figure 8. Note that which segment gadget we use for a segment s depends on where the angle of 60° at either end of the zone $Z(s)$ lies. Consider two intersecting gadgets with intersecting red regions R_{red} and R'_{red} , and 4-intersecting black regions R_{black} and R'_{black} . Let γ_{red} and γ'_{red} be the connector for R_{red} and R'_{red} , respectively. It is not difficult to verify that if γ_{red} is contained in R_{black} , then γ'_{red} is not contained in R'_{black} . And similarly if γ'_{red} is contained in R'_{black} , then γ_{red} is not contained in R_{black} . In other words, only one of the two segments/clauses can be satisfied by the endpoint/variable corresponding to the intersection. More formally, we have proven the following lemma.




Figure 7: Clause gadget and segment gadget with solely convex polygons with at most 8 corners.




Figure 8: Overlap of a segment gadget with another segment gadget and a clause gadget.

Lemma 5.3 *If $R_{red}(A, s) \cap R_{red}(B, s') \neq \emptyset$ for two segments s, s' , then A does not satisfy s or B does not satisfy s' . Similarly, if $R_{red}(A, s) \cap R_{red}(x, c)$ for some variable x in a clause c , then A does not satisfy s or x does not satisfy c .*

Figure 9 depicts the mutual overlapping of three segment gadgets. The next lemma can be verified by carefully investigating the intersection pattern of red and black regions.




Figure 9: Three mutually overlapping segment gadgets.

Lemma 5.4 *Suppose $R_{red}(A_1, s_1) \cap R_{red}(A_2, s_2) \cap R_{red}(B_3, s_3) \neq \emptyset$ for three segments s_1, s_2, s_3 . If A_1 satisfies s_1 or A_2 satisfies s_2 , then B_3 does not satisfy s_3 .*

We are now ready to prove Theorem 3.

Proof of Theorem 3: Let ψ be a 3-SAT formula, and G_ψ be planar with maximum degree 3. We define an instance \mathcal{I} of the non-crossing connectors problem as described above, consisting of linearly many regions, which are convex polygons with at most 8 corners and constantly many different slopes. We claim that ψ is satisfiable if and only if non-crossing connectors exist for \mathcal{I} .

First consider a set of non-crossing connectors for \mathcal{I} . By Lemma 5.1 every clause is satisfied by at least one variable. If variable x satisfies clause c , we set x to true if x is positive in c and false if x is negated in c . In case some variable has not received a truth value, we choose one arbitrarily. If this truth assignment is consistent, i.e., no variable is set to true and false at the same time, then it clearly satisfies formula ψ . So consider a variable x receiving a truth assignment from clause c , i.e., x satisfies c . We show that x does not satisfy a clause c' , in which x appears with the opposite sign. Consider the chain of segments s_1, \dots, s_k such that $R_{red}(A_1, s_1) \cap R_{red}(x, c) \neq \emptyset$, $R_{red}(B_k, s_k) \cap R_{red}(x, c') \neq \emptyset$, and $R_{red}(B_i, s_i) \cap R_{red}(A_{i+1}, s_{i+1}) \neq \emptyset$, for every $i = 1, \dots, k-1$. By Lemma 5.3 A_1 does not satisfy s_1 . Then by Lemma 5.2 B_1 satisfies s_1 . Then again by Lemma 5.3, or may be Lemma 5.4, A_2 does not satisfy s_2 . and by Lemma 5.2 B_2 satisfies s_2 . Iterating this pattern yields that B_k satisfies s_k and thus x does not satisfy c' , which is what we wanted to prove.

Secondly, we consider a satisfying truth assignment for formula ψ and want to conclude that there is a collection of non-crossing connectors for \mathcal{I} . We define the connectors for the clause gadget similarly to Figure 5, such that x satisfies c if and only if the variable x is assigned true and appears positive in c , or is assigned false and appears negated in c . Since ψ is satisfied, every clause c has at least one such variable x , so by Lemma 5.1 such non-crossing connectors do exist. Then we define the remaining connectors along the chain of segments starting at a clause c that is satisfied by x and ending at a clause c' that is not satisfied by x . By a reasoning similar to the previous one, we can construct non-crossing connectors this way for the entire instance \mathcal{I} . \square

6 Conclusions

In this paper we investigated the computational complexity of the non-crossing connectors problem, i.e., given pairwise disjoint finite point sets P_1, \dots, P_n and simply connected regions $R_i \supset P_i$ for every $i = 1, \dots, n$, is there a set of pairwise disjoint curves $\gamma_1, \dots, \gamma_n$, called connectors, such that $P_i \subset \gamma_i \subset R_i$ for every $i = 1, \dots, n$. We proved that the existence of non-crossing connectors can be tested in polynomial time if the regions are pseudo-disks or axis-aligned rectangles. It might be worthwhile to derive from our proofs polynomial-time algorithms to actually compute non-crossing connectors. We proved that the problem is NP-complete for 4-intersecting convex regions, even if every P_i consists of 2 elements. However, we do not know the complexity in case each R_i is the convex hull of the corresponding P_i . Moreover, it is interesting to consider other sets of regions, like ellipsoids or isosceles triangles with horizontal bases.

Instead of fixing the position of the points in P_i one can consider a set of possible positions for every such point. From previous results [2, 1] it follows that this variant is NP-complete if every point has at most 3 possible positions, but this proof again does not work if $R_i = \text{conv}(P_i)$ for $i = 1, \dots, n$. What if we want to connect pairs of vertical straight segments by non-crossing curves, each within the convex hull of the corresponding segments?

Furthermore, we could allow $P_i \cap P_j \neq \emptyset$ for $i \neq j$ and allow connectors γ_i, γ_j to intersect in $P_i \cap P_j$. If $|P_i| = 2$ for every i , this corresponds to drawing a given planar graph with fixed vertex positions and curved edges, each lying within a prescribed region. In this variant non-crossing connectors sometimes do not exist even when regions are pseudo-disks.

Acknowledgements

We would like to thank Maria Saumell, Stefan Felsner and Irina Mustata for fruitful discussions.

References

- [1] G. Aloupis, J. Cardinal, S. Collette, E. Demaine, M. Demaine, M. Dulieu, R. Fabila-Monroy, V. Hart, F. Hurtado, S. Langerman, M. Saumell, C. Seara, and P. Taslakian, *Non-crossing matchings of points with geometric objects*, Computational Geometry: Theory and Applications, to appear.
- [2] ———, *Matching points with things*, LATIN 2010: Theoretical Informatics, Lecture Notes in Computer Science, vol. 6034, Springer Berlin / Heidelberg, 2010, pp. 456–467.
- [3] S. Cabello, *Planar embeddability of the vertices of a graph using a fixed point set is NP-hard*, In 20th European Workshop on Computational Geometry (EWCG, 2004.
- [4] N. Castañeda and J. Urrutia, *Straight line embeddings of planar graphs on point sets*, Proceedings of the 8th Canadian Conference on Computational Geometry, Carleton University Press, 1996, pp. 312–318.
- [5] M. Chrobak and H. Karloff, *A lower bound on the size of universal sets for planar graphs*, SIGACT News **20** (1989).
- [6] E. Di Giacomo, F. Frati, R. Fulek, L. Grilli, and M. Krug, *Orthogeodesic point-set embedding of trees*, Proceedings of the 19th International Symposium on Graph Drawing (GD 2011), 2011.
- [7] V. Dujmović, W. Evans, S. Lazard, W. Lenhart, G. Liotta, D. Rappaport, and S. Wismath, *On point-sets that support planar graphs*, Proceedings of the 19th International Symposium on Graph Drawing (GD 2011), 2011.
- [8] V. Dujmović, M. Suderman, and D. R. Wood, *Drawings of planar graphs with few slopes and segments*, Computational Geometry Theory and Applications 38:194212, Springer, 2005, pp. 122–132.
- [9] H. Everett, S. Lazard, G. Liotta, and S. Wismath, *Universal sets of n points for one-bend drawings of planar graphs with n vertices*, Discrete & Computational Geometry **43** (2010), 272–288.
- [10] M. R. Fellows, J. Kratochvíl, M. Middendorf, and F. Pfeiffer, *The complexity of induced minors and related problems*, Algorithmica **13** (1995), 266–282.
- [11] H. Fraysseix, J. Pach, and R. Pollack, *How to draw a planar graph on a grid*, Combinatorica **10** (1990), 41–51.
- [12] E. R. Gansner, Y. F. Hu, M. Kaufmann, and S. G. Kobourov, *Optimal polygonal representation of planar graphs*, Algorithmica (2011), to appear.
- [13] A. Kaneko and M. Kano, *Discrete geometry on red and blue points in the plane – a survey*, in Discrete and Computational Geometry, The Goodman-Pollack Festschrift, Springer, 2003, pp. 551–570.
- [14] G. Kant, *Hexagonal grid drawings*, Graph-Theoretic Concepts in Computer Science, Lecture Notes in Computer Science, vol. 657, 1993, pp. 263–276.
- [15] M. Kaufmann and R. Wiese, *Embedding vertices at points: Few bends suffice for planar graphs*, in Graph Drawing (Proc. GD '99), LNCS 1731, 2002, pp. 165–174.

- [16] J. Kratochvíl, *String graphs II. Recognizing string graphs is NP-hard*, Journal of Combinatorial Theory, Series B **52** (1991), no. 1, 67 – 78.
- [17] J. Kratochvíl, A. Lubiw, and J. Nešetřil, *Noncrossing subgraphs in topological layouts*, SIAM Journal on Discrete Mathematics **4** (1991), no. 2, 223–244.
- [18] M. Kurowski, *A 1.235 lower bound on the number of points needed to draw all n -vertex planar graphs*, Information Processing Letters **92** (2004), no. 2, 95 – 98.
- [19] D. Lichtenstein, *Planar formulae and their uses*, SIAM Journal on Computing **11** (1982), no. 2, 329–343.
- [20] M. Löffler, *Existence of simple tours of imprecise points*, Proceedings of 23rd Annual European Workshop on Computational Geometry, 2007.
- [21] J. Pach, P. Gritzmann, B. Mohar, and R. Pollack, *Embedding a planar triangulation with vertices at specified points*, American Mathematical Monthly **98** (1991), 165–166.
- [22] M. Schaefer, E. Sedgwick, and D. Štefankovič, *Recognizing string graphs in NP*, Journal of Computer and System Sciences **67** (2003), no. 2, 365 – 380.
- [23] W. Schnyder, *Embedding planar graphs on the grid*, Proceedings of the first annual ACM-SIAM symposium on Discrete algorithms, SODA '90, 1990, pp. 138–148.
- [24] B. Speckmann, personal communication, 2011.
- [25] K. Verbeek, *Non-crossing paths with fixed endpoints*, Master's Thesis, Eindhoven, 2008.