Extending Drawings of Graphs to Arrangements of Pseudolines

Alan Arroyo
IST Austria, Klosterneuburg, Austria
https://alanarroyo.github.io/
alanmarcelo.arroyoguevara@ist.ac.at

Julien Bensmail
Université Côte d'Azur, CNRS, Inria, I3S, Sophia-Antipolis, France
julien.bensmail.phd@gmail.com

R. Bruce Richter
Department of Combinatorics and Optimization, University of Waterloo, Canada
brichter@uwaterloo.ca

Abstract

In the recent study of crossing numbers, drawings of graphs that can be extended to an arrangement of pseudolines (pseudolinear drawings) have played an important role as they are a natural combinatorial extension of rectilinear (or straight-line) drawings. A characterization of the pseudolinear drawings of K_n was found recently. We extend this characterization to all graphs, by describing the set of minimal forbidden subdrawings for pseudolinear drawings. Our characterization also leads to a polynomial-time algorithm to recognize pseudolinear drawings and construct the pseudolines when it is possible.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms; Mathematics of computing → Graphs and surfaces

Keywords and phrases graphs, graph drawings, geometric graph drawings, arrangements of pseudolines, crossing numbers, stretchability

Digital Object Identifier 10.4230/LIPIcs.SoCG.2020.9

Related Version A full version of the paper is available at [4], https://arxiv.org/abs/1804.09317.

Funding *Alan Arroyo*: Supported by CONACYT. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 754411.
Julien Bensmail: ERC Advanced Grant GRACOL, project no. 320812.
R. Bruce Richter: Supported by NSERC grant number 50503-10940-500.

1 Introduction

Since 2004, geometric methods have been used to make impressive progress for determining the crossing number of (certain classes of drawings of) the complete graph K_n. In particular, drawings that extend to straight lines, or, more generally, arrangements of pseudolines, have been central to this work, spurring interest in such drawings for arbitrary graphs, not just complete graphs [2, 5, 6, 7, 12].

In particular, for pseudolinear drawings, it is now known that, for $n \geq 10$, a pseudolinear drawing of K_n has more than

$$H(n) := \frac{1}{4}\left\lfloor\frac{n}{2}\right\rfloor\left\lfloor\frac{n-1}{2}\right\rfloor\left\lfloor\frac{n-2}{2}\right\rfloor\left\lfloor\frac{n-3}{2}\right\rfloor$$

crossings [1, 13]. The number $H(n)$ is conjectured by Harary and Hill to be the smallest number of crossings over all topological drawings of K_n; that is, the crossing number $\mathrm{cr}(K_n)$ is conjectured to be $H(n)$.

36th International Symposium on Computational Geometry (SoCG 2020).
Editors: Sergio Cabello and Danny Z. Chen; Article No. 9; pp. 9:1–9:14
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

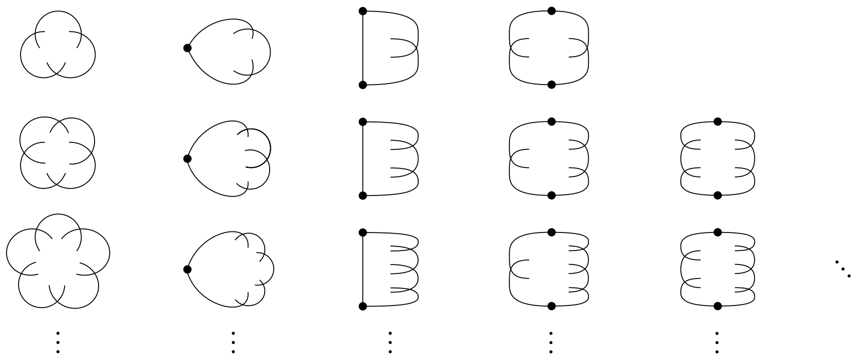

Figure 1 Obstructions to pseudolinearity.

A *pseudoline* is the image ℓ of a continuous injection from the real numbers $\mathbb{R}$ to the plane $\mathbb{R}^2$ such that $\mathbb{R}^2 \setminus \ell$ is not connected. An *arrangement of pseudolines* is a set Σ of pseudolines such that, if ℓ, ℓ' are distinct elements of Σ, then $|\ell \cap \ell'| = 1$ and the intersection is a crossing point. More on pseudolines and their importance for studying geometric drawings of graphs can be found in [10, 11].

A drawing D of a graph G is *pseudolinear* if there is an arrangment of pseudolines consisting of a different pseudoline ℓ_e for each edge e of G and such that $D[e] \subseteq \ell_e$.

In the study of crossing numbers, restricting the drawing to either straight lines or pseudolines yields the rectilinear crossing number $\overline{\mathrm{cr}}(K_n)$ or the pseudolinear crossing number $\widetilde{\mathrm{cr}}(K_n)$, respectively. Clearly $\overline{\mathrm{cr}}(K_n) \geq \widetilde{\mathrm{cr}}(K_n)$ and the geometric methods prove that $\widetilde{\mathrm{cr}}(K_n) > H(n)$, for $n \geq 10$.

A *good drawing* is one where no edge self-intersects and any two edges share at most one point – either a crossing or a common end point – and no three edges share a common crossing. One somewhat surprising result is from Aichholzer et al.: *a good drawing of K_n in the plane is homeomorphic to a pseudolinear drawing if and only if it does not contain a non-planar drawing of K_4 whose crossing is incident with the unbounded face of the K_4* [2]. There are equivalent characterizations in [5, 6]. These conditions can be shown to be equivalent to not containing the *B-configuration* depicted as the third drawing of the first row of Figure 1.

Twenty-five years earlier, Thomassen proved a similar theorem for drawings in which each edge is crossed only once [16]. The B- and W-configurations are shown as the third and fourth drawings in the first row of Figure 1. Thomassen's theorem is: *if D is a planar drawing of a graph G in which each edge is crossed at most once, then D is homeomorphic to a rectilinear drawing of G if and only if D contains no B- or W-configuration.*

Thomassen presented in [16] the *clouds* (first column in Figure 1) as an infinite family of drawings that are minimally non-pseudolinear.

Shortly after Thomassen's paper, Bienstock and Dean proved that if $\mathrm{cr}(G) \leq 3$, then $\overline{\mathrm{cr}}(G) = \mathrm{cr}(G)$ [8]. They also exhibited examples based on overlapping W-configurations to show the result fails for $\mathrm{cr}(G) = 4$; such graphs can have arbitrarily large rectilinear crossing number.

Despite the existence of infinitely many obstructions to pseudolinearity, we characterize them all.

▶ **Theorem 1.** *A good drawing of a graph G is pseudolinear if and only if it does not contain one of the infinitely many obstructions shown in Figure 1.*

The drawings in Figure 1 are obtained from the *clouds* (first column) by replacing at most two crossings by vertices. The formal statement of Theorem 1 is Theorem 15 in Section 6; also a more general version of this statement, Theorem 2, is discussed below. That there is a result such as ours is somewhat surprising, because stretching an arrangement of pseudolines to a rectilinear drawing has been shown by Mnëv [14, 15] to be $\exists\mathbb{R}$-hard. In particular, recognizing a drawing as being homeomorphic to a rectilinear drawing is NP-hard.

The natural setting for our characterization is strings embedded in the plane. An *arc* σ is the image $f([0,1])$ of the compact interval $[0,1]$ under a continuous map $f:[0,1]\to\mathbb{R}^2$. Let $S(\sigma)=\{p\in\sigma\ :\ |f^{-1}(p)|\geq 2\}$ be the set of self-intersections of σ. A *string* is an arc σ for which $S(\sigma)$ is finite. If $S(\sigma)=\emptyset$, then σ is *simple*.

An intersection point between of two strings σ and σ' is *ordinary* if it is either an endpoint of σ or σ', or is a *crossing* (a crossing is a non-tangential intersection point in $\sigma\cap\sigma'$ that is not an end of σ or σ'). A set Σ of strings is *ordinary* if Σ is finite and any two strings in Σ have only finitely many intersections, all of which are ordinary. All the sets of strings considered in this paper are ordinary.

If Σ is an ordinary set of strings, then its *planarization* $G(\Sigma)$ is the plane graph obtained from Σ by inserting vertices at each crossing between strings and also at the endpoints of every string in Σ. To keep track of the information given by the strings, we will always assume that each string Σ has a different color and that each edge in $G(\Sigma)$ inherits the color of the string including it.

If Σ is an ordinary set of strings, then, for a cycle C in $G(\Sigma)$ (which is a simple closed curve in $\mathbb{R}^2$) and a vertex $v\in V(C)$, v is a *rainbow* for C if all the edges incident with v and drawn in the closed disk bounded by C (including the two edges of C at v) have different colours. The reader can verify that, for each drawing in Figure 1, if we let Σ be the edges of the drawing, then the unique cycle in $G(\Sigma)$ has at most two rainbows. Our main result characterizes these cycles as the only possible obstructions:

▶ **Theorem 2.** *An ordinary set of strings Σ can be extended to an arrangement of pseudolines if and only if every cycle C of $G(\Sigma)$ has at least three rainbows.*

Henceforth, we define any cycle C in $G(\Sigma)$ with at most two rainbows as an *obstruction*. A set of strings is *pseudolinear* if it has an extension to an arrangement of pseudolines.

Theorem 2 is our main contribution. In the next section, we show that the presence of an obstruction implies the set of ordinary strings is not pseudolinear. The converse is proved in Section 4 by extending, one small step at a time, the strings in Σ to get closer to an arrangement of pseudolines. After each extension, we must show that no obstruction has been introduced. This involves dealing with cycles in $G(\Sigma)$ that have precisely three rainbows (that we refer as *near-obstructions*). In Section 3 we show the key lemma that if G has two such near-obstructions that intersect nicely at a vertex v, then G has an obstruction. In Section 5 we present a polynomial-time algorithm for detecting obstructions and we argue why the proof of Theorem 2 implies a polynomial-time algorithm for extending a pseudolinear set of strings. Finally, in Section 6, we show how Theorem 1 follows from Theorem 2 and we present some concluding remarks.

2 A set of strings with an obstruction is not extendible

Let us start by showing the easy direction of Theorem 2:

▶ **Lemma 3.** *If the underlying graph $G(\Sigma)$ of a set Σ of strings has an obstruction, then Σ is not pseudolinear.*

Suppose that C is a cycle of $G(\Sigma)$ for some set of strings Σ. We define $\delta(C)$ as the set of vertices of C for which their two incident edges in C have different colours. In a set Σ of simple strings where no two intersect twice, $|\delta(C)| \geq 3$ for every cycle C of $G(\Sigma)$.

▶ **Lemma 4.** *Let Σ be a set of simple strings where every pair intersect at most once. Suppose that C is an obstruction with $|\delta(C)|$ as small as possible. Let $S = x_0, x_1, \ldots, x_\ell$ be a path of $G(\Sigma)$ representing a subsegment of some string $\sigma \in \Sigma$ such that $x_0x_1 \in E(C)$, $x_1 \in \delta(C)$ and x_1 is not a rainbow of C. Then $V(C) \cap V(S) = \{x_0, x_1\}$.*

Proof. By way of contradiction, suppose that there is a vertex $x_r \in V(C) \cap V(S)$ with $r \geq 3$. Assume that $r \geq 3$ is as small as possible. Let P be the subpath of S connecting x_1 to x_r. Since $x_0x_1 \in E(C)$ and $x_1 \in \delta(C)$ and $P \subseteq \sigma$, $x_1x_2 \notin E(C)$. Because x_1 is not a rainbow for C and no two strings tangentially intersect at x_1, the edge x_1x_2 is drawn in the closed disk bounded by C. By choice of r, P is an arc connecting x_1 to x_r in the interior of C.

Let C_1 and C_2 be the cycles obtained from the union of P and one of the two xy-subpaths in C. We may assume that $x_0x_1 \in E(C_1)$. Let $\rho(C)$ be either $\delta(C)$ or the set of rainbows in C. For $i = 1, 2$, let $Q_i = V(C_i) \setminus V(P)$. Then $\rho(C) \cap Q_i = \rho(C_i) \cap Q_i$. We see that $\rho(C_1) \setminus Q_1 \subseteq \{x_r\}$ and $\rho(C_2) \setminus Q_2 \subseteq \{x_1, x_r\}$.

For $\rho = \delta$, $|\delta(C_2)| \geq 3$, so $|\delta(C) \cap Q_2| \geq 1$. Since $x_1 \notin \delta(C_1)$, $|\delta(C_1)| \leq |\delta(C_1) \cap Q_2| + |\{x_r\}| \leq |\delta(C)| - 2 + |\{x_r\}| < |\delta(C)|$. Likewise, $|\delta(C) \cap Q_1| \geq 2$ and $x_1 \in \delta(C) \cap \delta(C_2)$. Therefore, $|\delta(C_2)| \leq |\delta(C)| - 2 + |\{x_r\}| < |\delta(C)|$. Thus, neither C_1 nor C_2 is an obstruction.

Now taking ρ to be the set of rainbows, the preceding paragraph shows $|\rho(C_1)| \geq 3$ and $|\rho(C_2)| \geq 3$. Therefore, $|\rho(C) \cap Q_1| = |\rho(C_1) \cap Q_1| \geq 2$ and $|\rho(C) \cap Q_2| = |\rho(C_2) \cap Q_2| \geq 1$. Thus, $|\rho(C)| \geq 3$, a contradiction. ◀

Proof of Lemma 3. By way of contradiction, suppose that Σ is pseudolinear and that $G(\Sigma)$ has an obstruction C.

Consider an extension of Σ to an arrangement of pseudolines, and then cut off the two infinite ends of each pseudoline to obtain a set of strings Σ' extending Σ, and in which every pair of strings in Σ' cross once. In $G(\Sigma')$, there is a cycle C' that represents the same simple closed curve as C. Because C' is obtained from subdividing some edges of C and the colours of a subdivided edge are the same, C' has fewer than three rainbows. Therefore, we may assume that $\Sigma = \Sigma'$ and $C = C'$. Now, the ends of every string in Σ are degree-1 vertices in the outer face of $G(\Sigma)$.

As every string in Σ is simple and no two strings intersect more than once, $|\delta(C)| \geq 3$. We will assume that C is chosen to minimize $|\delta(C)|$.

Since C is an obstruction, there exists $x_1 \in \delta(C)$ such that x_1 is not a rainbow in C. Consider a neighbour x_0 of x_1 in C. Let $S = x_0, x_1, \ldots x_\ell$ be the path obtained by traversing the string σ extending x_0x_1, such that x_ℓ is an end of σ. By Observation 4, $V(S) \cap V(C) = \{x_0, x_1\}$, and because x_ℓ is in the outer face of C, the segment of σ from x_1 to x_ℓ has its relative interior in the outer face of C.

However, since x_1 is not a rainbow, there exists a string $\sigma' \in \Sigma$ including two edges at x_1 drawn in the disk bounded by C. Thus, σ and σ' tangentially intersect at x_1, a contradiction. ◀

3 The key lemma

In this section we present the key lemma used in the proof of Theorem 2.

A plane graph G is *path-partitioned* if for $m \geq 1$, there exists a colouring $\chi : E(G) \to \{1, \ldots, m\}$ such that for each $i \in \{1, \ldots, m\}$, the edges in $\chi^{-1}(i)$ induce a path $P_i \subseteq G$ where any two distinct paths P_i and P_j do not tangentially intersect. Indeed, every underlying planar graph $G(\Sigma)$ of a set of simple strings Σ is path-partitioned. Moreover, every path-partitioned plane graph can be obtained by subdividing a planarization of an ordinary set of simple strings. To extend the previously introduced notation we refer to each P_i as a string. The concepts of rainbow and obstruction naturally extend to the context of path-partitioned plane graphs.

Suppose that G is a path-partitioned plane graph. Given $v \in V(G)$, a *near-obstruction at* v is a cycle C with at most three rainbows and such that v is a rainbow of C. Understanding how near-obstructions behave is the key ingredient needed in the proof of Theorem 2:

▶ **Lemma 5.** *Let G be a path-partitioned plane graph and let $v \in V(G)$. Suppose that C_1 and C_2 are two near-obstructions at v such that the union of the closed disks bounded by C_1 and C_2 contains a small open ball centered at v. Suppose that one of the following two holds:*

1. *no obstruction of G contains v; or*
2. *the two edges of C_1 incident with v are the same as the two edges of C_2 incident with v.*

Then G has an obstruction not including v.

Given a plane graph G, a cycle $C \subseteq G$ and a vertex $v \in V(C)$, *the edges at v inside C* are the edges of G incident with v drawn in the disk bounded by C.

▶ **Useful Fact.** *Let G be planar path-partitioned graph. Suppose that for two cycles C and C', $v \in V(C) \cap V(C')$ is a vertex such that the edges at v inside C' are also edges at v inside C. If v is a rainbow for C, then v is a rainbow for C'.*

Proof of Lemma 5. By way of contradiction, suppose that G has no obstruction not including v. The "small ball" hypothesis implies that v is not in the outer face of the subgraph $C_1 \cup C_2$.

We claim that $|V(C_1) \cap V(C_2)| \geq 2$. Suppose not. Then C_1 and C_2 are edge-disjoint and $V(C_1) \cap V(C_2) = \{v\}$. For $i = 1, 2$, let e_i and f_i be the edges of C_i at v and let Δ_i be the closed disk bounded by C_i. From the "small ball" hypothesis it follows that (i) Δ_1 contains the edges e_2 and f_2; and (ii) the points near v in the exterior of Δ_2 are contained in Δ_1. These two properties imply that the path $C_2 - \{e_2, f_2\}$ intersects C_1 at least twice, and hence, $|V(C_1) \cap V(C_2)| \geq 2$.

From the last paragraph we know that $C_1 \cup C_2$ is 2-connected, and hence the outer face of $C_1 \cup C_2$ is bounded by a cycle C_{out}. We will assume that

(*) the cycles C_1 and C_2 satisfying the hypothesis of Lemma 5 are chosen so that the number of vertices of G in the disk bounded by C_{out} is minimal.

The Useful Fact applied to $C = C_{out}$ and to each $C' \in \{C_1, C_2\}$, shows that every vertex that is a rainbow in C_{out} is also a rainbow in each of the cycles in $\{C_1, C_2\}$ containing it. We can assume that C_{out} is not an obstruction or else we are done. We may relabel C_1 and C_2 so that two of the rainbows of C_{out}, say p and q, are also rainbows in C_1. Neither p nor q is v because $v \notin V(C_{out})$. Because C_1 is a near-obstruction, p, q and v are the only rainbows of C_1.

Since $v \notin V(C_{out})$, by following C_1 in the two directions starting at v, we find a path $P_v \subseteq C_1$ containing v in which only the ends u and w of P_v are in C_{out} (note that $u \neq v$ because $\{p, q\} \subseteq V(C_1) \cap V(C_{out})$). As v is in the interior face of C_{out}, P_v is also in the interior of C_{out}. Let Q^1_{out}, Q^2_{out} be the uw-paths of C_{out}. One of the two closed disks bounded by $P_v \cup Q^1_{out}$ and $P_v \cup Q^2_{out}$ contains C_1. By symmetry, we may assume that C_1 is contained in the first disk. Since $C_{out} \subseteq C_1 \cup C_2$, this implies that Q^2_{out} is a subpath of C_2.

Our desired contradiction will be to find three rainbows in C_2 distinct from v. We find the first: let $C_1 - (P_v)$ be the uw-path in C_1 distinct from P_v. The disk bounded by $(C_1 - (P_v)) \cup Q^2_{out}$ contains the one bounded by C_1. The Useful Fact applied to $C = (C_1 - (P_v)) \cup Q^2_{out}$ and $C' = C_1$ implies that each vertex in $C_1 - (P_v)$ that is rainbow in $(C_1 - (P_v)) \cup Q^2_{out}$ is also rainbow in C_1. Since C_1 has at most two rainbows in $C_1 - (P_v)$, namely p and q, $(C_1 - (P_v)) \cup Q^2_{out}$ has a third rainbow r_1 in the interior of Q^2_{out} (else $(C_1 - (P_v)) \cup Q^2_{out}$ is an obstruction and we are done). Note that r_1 is also a rainbow for C_2.

To find another rainbow in C_2, consider the edge e_u of C_2 incident to u and not in Q^2_{out}. We claim that either u is a rainbow in C_2 or that e_u is not included in the closed disk bounded by $P_v \cup Q^2_{out}$. Seeking a contradiction, suppose that u is not a rainbow of C_2 and that e_u is included in the disk. Then we can find two edges in the rotation at u, included in the disk bounded by $P_v \cup Q^2_{out}$, that belong to the same string σ. The vertex u is a rainbow in C_1, as else, we would find a string σ' with two edges inside $Q^1_{out} \cup P_v$, showing that σ and σ' tangentially intersect at u. As p and q are the only rainbows of C_1 in C_{out}, u is one of p and q. Therefore u is a rainbow in C_{out}, and hence, a rainbow in C_2, a contradiction.

If u is a rainbow in C_2, then this is the desired second one. Otherwise, e_u is not in the closed disk bounded by $P_v \cup Q^2_{out}$. Let $P_u \subseteq C_2$ be the path starting at u, continuing on e_u and ending on the first vertex u' in P_v that we encounter. Let C_u be the cycle consisting of P_u and the uu'-subpath uP_vu' of P_v.

▷ Claim 6. If P_u does not have a rainbow of C_u in its interior, then either C_u is an obstruction not containing v or:

(a) C_u and C_2 are near-obstructions at v satisfying the same conditions as C_1 and C_2 in Lemma 5; and

(b) the closed disk bounded by the outer cycle of $C_u \cup C_2$ contains fewer vertices than the disk bounded by C_{out}.

Proof. Suppose that all the rainbows of C_u are located in uP_vu'. If z is a rainbow of C_u, then $z \in \{u, v, u'\}$, as otherwise z is a rainbow of C_1 distinct from p, q and v, a contradiction. Thus, if $v \notin V(C_u)$, then C_u is the desired obstruction. We may assume that $v \in V(C_u)$.

If $u' = w$, then $C_2 = P_u \cup Q^2_{out}$, violating the assumption that $v \in V(C_2)$. Thus $u' \neq w$. If $u' = v$, then the rainbows of C_u are included in $\{u, u'\}$, and hence C_u is an obstruction. However, the existence of C_u shows that both alternatives (1) and (2) in Lemma 5 fail: condition (1) fails because C_u contains v and (2) fails because the edge of P_u incident with v is in $E(C_2) \setminus E(C_1)$. Thus $u' \neq v$.

The previous two paragraphs show that C_u is a near-obstruction at v with rainbows u, v and u'. Since the interior of C_u near v is the same as the interior of C_1 near v, the pair (C_u, C_2) satisfies the "small ball" hypothesis. Thus, (a) holds.

Let C'_{out} be the outer cycle of $C_u \cup C_2$. From the fact that $C_u \cup C_2 \subseteq C_1 \cup C_2$ it follows that the disk bounded by C_{out} includes the disk bounded by C'_{out}.

Since $p, q \in V(C_{out})$, p and q are in the disk bounded by C_{out}. If both p and q are in C_2, then p, q and r_1 are rainbows in C_2, and also distinct from v, contradicting that C_2 is a near-obstruction for v. If, say $p \notin V(C_2)$, then p is not in the disk bounded by C'_{out}, which implies (b). ◁

From Claim 6(b) and assumption (*) either C_u is the desired obstruction or P_u contains a rainbow r_2 of C_2 in its interior. We assume the latter as else we are done.

In the same way, the last rainbow r_3 comes by considering the edge of $C_2 - Q^2_{out}$ incident with w. It follows that v, r_1, r_2 and r_3 are four different rainbows in C_2, contradicting the fact that C_2 is a near-obstruction. ◀

4 Proof of Theorem 2

In this section we prove that a set of strings with no obstructions can be extended to an arrangement of pseudolines.

Proof of Theorem 2. It was shown in Observation 3 that the existence of obstructions implies non-extendibility. For the converse, suppose that Σ is a set of strings for which $G(\Sigma)$ has no obstructions.

We start by reducing to the case where the point set $\bigcup\Sigma$ is connected: iteratively add a new string in a face of $\bigcup\Sigma$ connecting two connected components of $\bigcup\Sigma$. No obstruction is introduced at each step (obstructions are cycles), and, eventually, the obtained set $\bigcup\Sigma$ is connected. An extension of the new set of strings contains an extension for the original set, thus we may assume that $\bigcup\Sigma$ is connected.

Our proof is algorithmic, and consists of repeatedly applying one of the three steps described below.

- **Disentangling Step.** If a string $\sigma \in \Sigma$ has an end a with degree at least 2 in $G(\Sigma)$, then we slightly extend the a-end of σ into one of the faces incident with a.
- **Face-Escaping Step.** If a string $\sigma \in \Sigma$ has an end a with degree 1 in $G(\Sigma)$, and is incident with an interior face, then we extend the a-end of σ until it intersects some point in the boundary of this face.
- **Exterior-Meeting Step.** Assuming that all the strings in Σ have their two ends in the outer face and these ends have degree 1 in $G(\Sigma)$, we extend the ends of two disjoint strings so that they meet in the outer face.

Each of these three steps either increases the number of pairs of strings that intersect, or increase the number crossings (recall that a crossing between σ and σ' is a non-tangential intersection point in $\sigma \cap \sigma'$ that is not an end of σ or σ'). Moreover, these steps can be performed as long as not all the strings have their ends in the outer face and they are pairwise crossing (in this case we extend their ends to infinity to obtain the desired arrangement of pseudolines). Henceforth, we will show that, if performed correctly, none of these steps introduces an obstruction. The proof for each step can be read independently.

▶ **Lemma 7** (Disentangling Step). *Suppose that $\sigma \in \Sigma$ has an end a with degree at least 2 in $G(\Sigma)$. Then we can extend the a-end of σ into one of the faces incident to a without creating an obstruction.*

Proof. A pair of different edges f and f' in $G(\Sigma)$ incident with a are *twins* if they belong to the same string in Σ. The edge $e \subseteq \sigma$ incident with a has no twin.

The fact that no pair of strings tangentially intersect at a tells us that if (f_1, f_1') and (f_2, f_2') are pairs of twins, then f_1, f_2, f_1', f_2' occur in this cyclic order for either the clockwise or counterclockwise rotation at a. Thus, we may assume that the counterclockwise rotation at a restricted to the twins and e is $e, f_1, \dots, f_t, f_1', \dots, f_t'$, where (f_i, f_i') is a twin pair for $i = 1, \dots, t$.

To avoid tangential intersections, the extension of σ at a must be in the angle between f_t and f_1' not containing e. Let $e_1, \dots, e_k$ be the counterclockwise ordered list of non-twin edges at a having an end in this angle (as depicted in Figure 2). We label $e_0 = f_t$ and $e_{k+1} = f_1'$. If there are no twins, then let $e_0 = e_{k+1} = e$.

Let us consider all the possible extensions: for $i \in \{0, \dots, k\}$, let Σ_i be the set of strings obtained from Σ by slightly extending the a-end of σ into the face containing the angle between e_i and e_{i+1}. Let α_i be the new edge at a extending σ in Σ_i (see α_0 in Figure 2).

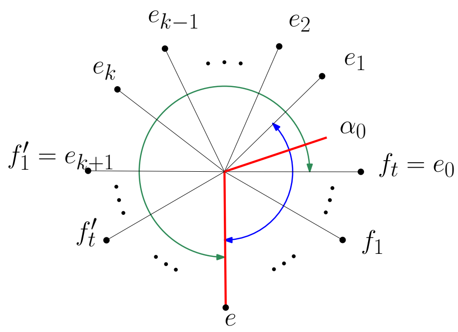

Figure 2 Substrings included in the disk bounded by C_0.

Seeking a contradiction, suppose that, for each $i \in \{0, ..., k\}$, $G(\Sigma_i)$ contains an obstruction C_i. Since α_i contains a degree-1 vertex, α_i is not in C_i. Hence C_i is a cycle of $G(\Sigma)$. Thus C_i is not an obstruction in $G(\Sigma)$ that becomes one in $G(\Sigma_i)$. This conversion has a simple explanation: in $G(\Sigma)$, C_i has exactly three rainbows, and one of them is a. After α_i is added, a is not a rainbow in C_i (witnessed by the edges e and α_i included in the new version of σ).

Recall from Section 3 that a *near-obstruction at a* is a cycle with exactly three rainbows, and one of them is a. Each of C_0, C_1,...,C_k is a near-obstruction at a in $G(\Sigma)$.

For a cycle $C \subseteq G$, let $\Delta(C)$ denote the closed disk bounded by C. Both e and α_0 are in $\Delta(C_0)$. Thus, either $\Delta(C_0) \supseteq \{e, f_1, f_2, \ldots, f_t, e_1\}$ (blue bidirectional arrow in Figure 2) or $\Delta(C_0) \supseteq \{f_t, e_1, \ldots, e_k, f'_1, f'_2, \ldots, f'_t, e\}$ (green bidirectional arrow). We rule out the latter situation as the second list contains f_t and f'_t, and this would imply that a is not a rainbow for C_0 in $G(\Sigma)$.

We just showed that $\{e, e_0, e_1\} \subseteq \Delta(C_0)$. By symmetry, $\{e_k, e_{k+1}, e\} \subseteq \Delta(C_k)$. Consider the largest index $i \in \{0, 1, \ldots, k-1\}$ for which $\{e, e_0, \ldots, e_{i+1}\} \subseteq \Delta(C_i)$. By the choice of i, and because $\{e, \alpha_{i+1}\} \subseteq \Delta(C_{i+1})$, $\{e, f'_t, \ldots, f'_1, e_k, \ldots, e_i\} \subseteq \Delta(C_{i+1})$. However, by applying Lemma 5 to the pair C_i and C_{i+1}, we obtain that $G(\Sigma)$ has an obstruction, a contradiction. ◀

▶ **Lemma 8** (Face-Escaping Step). *Suppose that there is a string σ that has an end a with degree 1 in $G(\Sigma)$, and a is incident to an interior face F. Then there is an extension σ' of σ from its a-end to a point in the boundary of F such that the set $(\Sigma \setminus \{\sigma\}) \cup \{\sigma'\}$ has no obstruction.*

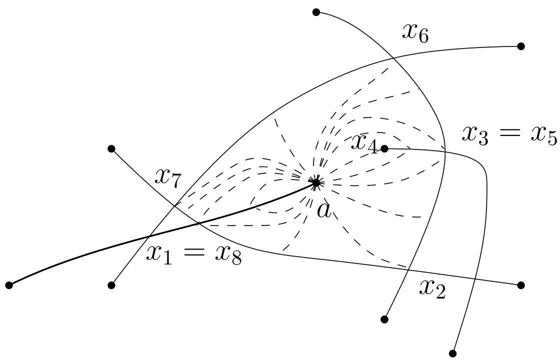

Figure 3 All possible extensions in the Face-Escaping Step.

Proof. Let W be the closed boundary walk $(x_0, e_1, x_1, e_2, \dots, e_n, x_n)$ of F such that $x_0 = x_n = a$ and F is to the left as we traverse W (see Figure 3 for an illustration with $n = 9$). For $i = 1, \dots, n$ we let m_i be a point in the relative interior of e_i, and let P be the list of non-necessarily distinct points m_1, x_1, m_2, $x_2 \dots, m_n$, which are the potential ends for all the different extensions. For each $p \in P$, let Σ_p be the set of strings obtained from Σ by extending the a-end of σ by adding an arc α_p connecting a to p in F (see Figure 3). We assume that every two distinct arcs α_p and $\alpha_{p'}$ are internally disjoint.

Let f_p be the edge $e_1 \cup \alpha_p$ in $G(\Sigma_p)$; f_p has ends x_1 and p. Also, let $\sigma^p = \sigma \cup \alpha_p$. Seeking a contradiction, suppose that each $G(\Sigma_p)$ has an obstruction.

▷ Claim 9. Let $p \in P$. Then there exists an obstruction C_p in $G(\Sigma_p)$ including f_p. Moreover,
(1) if $p \in \sigma$, then C_p can be chosen so that all its edges are included in σ^p; and
(2) if $p \notin \sigma$, then every obstruction includes f_p.

Proof. First, if $p \in \sigma$, then the string σ^p self-intersects at p; thus σ^p has a simple close curve including f_p. In this case let C_p be the cycle in $G(\Sigma_p)$ representing this simple closed curve without rainbows, and thus (1) holds.

Second, assume that $p \notin \sigma$ and let C_p be any obstruction of $G(\Sigma_p)$. For (2), we will show that $f_p \in E(C_p)$.

Seeking a contradiction, suppose that $f_p \notin E(C_p)$.

If $p = m_i$ for $i \in \{1, \dots, n\}$, since m_i is the only vertex whose rotation in $G(\Sigma)$ differs from its rotation in $G(\Sigma_{m_i})$, $m_i \in V(C_p)$. Consider the cycle C of $G(\Sigma)$ obtained from C_p by replacing the subpath (x_{i-1}, m_i, x_i) by the edge $x_{i-1}x_i$. For each vertex $v \in V(C)$ the colors of the edges of $G(\Sigma)$ at v included in the disk bounded by C are the same as in $G(\Sigma_p)$ for the disk bounded by $V(C_p)$. Thus, C is an obstruction for $G(\Sigma)$, a contradiction.

Suppose now that p is one of $x_1, \dots, x_{n-1}$. The only vertex in $G(\Sigma)$ whose rotation is different in $G(\Sigma_p)$ is p. Therefore, p is a point that is a rainbow for C_p in $G(\Sigma)$, but not a rainbow in $G(\Sigma_p)$, witnessed by two edges included in σ^p. Since at least one of the two witnessing edges is in $G(\Sigma)$, $p \in \sigma$. This contradicts the assumption that $p \notin \sigma$. Hence $f_p \in E(C_p)$. ◁

Henceforth we assume that, for $p \in P$, C_p is an obstruction in $G(\Sigma_p)$ as in Claim 9.

More can be said about the obstructions in $G(\Sigma_p)$, but for this we need some terminology. If we orient an edge e in a plane graph, then the *sides* of e are either the points near e that are to the right of e, or the points near e to the left of e. For any cycle C of G through e, exactly one side of e lies inside C. This is the side of e *covered* by C. For the next claim and in the rest of the proof we will assume that for $p \in P$, f_p is oriented from x_1 to p.

▷ Claim 10. For $p \in P$ with $p \notin \sigma$, every obstruction in $G(\Sigma_p)$ covers the same side of f_p.

Proof. Suppose that for $p \in P$ there are obstructions C_p and C'_p covering both sides of f_p. Let G' be the plane graph obtained from $G(\Sigma_p)$ by subdividing f_p, and let v be the new degree-2 vertex inside f_p.

We consider the edge-colouring χ induced by the strings in Σ_p. Let χ' be a new colouring obtained from χ by replacing the colour of the edge vp by a new colour not used in χ. It is a routine exercise to verify that (i) χ' induces a path-partition in G' (defined in Section 3); and (ii) C_p and C'_p are near-obstructions for v with respect to χ'. By applying Lemma 5 to $C_1 = C_p$ and $C_2 = C'_p$, we obtain an obstruction in G' not containing v. However, this implies the existence of an obstruction in $G(\Sigma)$, a contradiction. ◁

Recall that the boundary walk of F is $W = (x_0, e_1, \dots, e_n, x_n)$, with $x_0 = x_n = a$. Since x_1 and x_{n-1} are in σ, the extreme obstructions C_{x_1} and C_{x_2} cover the right of f_{x_1} and the left of $f_{x_{n-1}}$, respectively. Thus, there are two consecutive vertices x_{i-1}, x_i in $W - a$, such that the interior of $C_{x_{i-1}}$ covers the right of $f_{x_{i-1}}$ and the interior of C_{x_i} covers the left of f_{x_i}. Moreover, we may assume that the interior of C_{m_i} includes the left of f_{m_i} (otherwise we reflect our drawing).

The next claim (proved in the full version of this paper [4]) is the last ingredient to obtain a final contradiction.

▷ Claim 11. Exactly one of the following holds:

(a) $x_{i-1} \in \sigma$, $m_i \notin \sigma$ and $G(\Sigma_{m_i})$ has an obstruction covering the side of f_{m_i} not covered by C_{m_i}; or

(b) $x_{i-1} \notin \sigma$ and $G(\Sigma_{x_{i-1}})$ has an obstruction covering the side of $f_{x_{i-1}}$ not covered by $C_{x_{i-1}}$.

Claims 10 and 11 contradict each other. Thus, for some $p \in P$, $G(\Sigma_p)$ has no obstructions. ◀

▶ **Lemma 12** (Exterior-Meeting Step). *If all the strings in Σ have their ends on the outer face of $G(\Sigma)$ and the ends have degree 1 in $G(\Sigma)$, then we can extend a pair disjoint strings so that they intersect without creating an obstruction.*

Proof. By following the outer boundary of $\bigcup \Sigma$, we obtain a simple closed curve $\mathcal{O}$ containing all the ends of the strings in Σ, but otherwise disjoint from $\bigcup \Sigma$.

Suppose σ_1, σ_2 are two disjoint strings in Σ. For $i = 1, 2$, let a_i, b_i be the ends of σ_i; since σ_1 and σ_2 do not cross, we may assume that these ends occur in the cyclic order a_1, b_1, b_2, a_2. We extend the a_i-ends of σ_1 and σ_2 so that they meet in a point p in the outer face, and so that all the ends of σ_1 and σ_2 remain incident with the outer face (Figure 4). Let Σ' be the obtained set of strings.

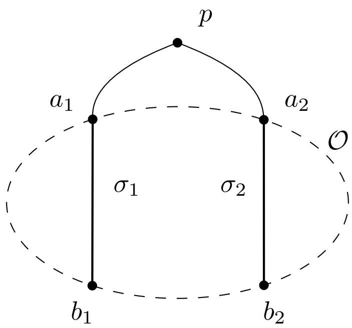

Figure 4 Exterior-Meeting Step.

Seeking a contradiction, suppose that $G(\Sigma')$ has an obstruction C. Since $G(\Sigma)$ has no obstruction, $p \in V(C)$. Our contradiction will be to find three rainbows in C. Note that p is a rainbow. To obtain a second rainbow, traverse C starting from p towards a_1. Let d_1 be the first vertex during our traversal that is not in the extended σ_1, and let c_1 be its neighbour in σ_1, one step before we reach d_1. Since b_1 has degree one, $c_1 \neq b_1$.

▷ Claim 13. The cycle C has a rainbow included in the closed disk Δ_1 bounded by σ_1 and the $a_1 b_1$-arc of $\mathcal{O}$ disjoint from σ_2.

Proof. First, suppose that $d_1 \notin \Delta_1$. In this case, c_1 is a rainbow because otherwise there would be a string σ that tangentially intersects σ_1 at c_1. Thus, if $d_1 \notin \Delta_1$, then c_1 is the desired rainbow.

Second, suppose that $d_1 \in \Delta_1$. Let P_1 be the path of C starting at c_1, continuing on the edge $c_1 d_1$, and ending at the first vertex we encounter in σ_1. Since the cycle C' enclosed by $P_1 \cup \sigma_1$ is not an obstruction, there is one rainbow of C' that is an interior vertex of P_1; this is the desired rainbow of C. This concludes the proof of Claim 13. ◁

Considering σ_2 instead of σ_1, Claim 13 yields a third rainbow in C inside an analogous disk Δ_2 disjoint from Δ_1, contradicting that C is an obstruction. Hence Lemma 12 holds. ◀

We iteratively apply the Disentangling Step, Face-Escaping Step or Exterior-Meeting Step without creating obstructions. Each step increases the number of pairwise intersecting strings in Σ until we reach a stage where the strings are pairwise intersecting and all of them have their two ends in the unbounded face. From this we extend them into an arrangement of pseudolines. This concludes the proof of Theorem 2. ◀

5 Finding obstructions and extending strings in polynomial time

We start this section by describing an algorithm to detect obstructions. Henceforth, we assume that the input to the problem is the planarization $G(\Sigma)$ of an ordinary set of s strings Σ. For the running-time analysis, we assume that n and m are the number of vertices and edges in $G(\Sigma)$, respectively. Since $G(\Sigma)$ is planar, $m = O(n)$. Moreover, if Σ is pseudolinear, then $n \leq \binom{s}{2} + 2s = \binom{s+2}{2} - 1$. At the end of this section we explain how to extend Σ (if possible) in polynomial time.

Recall that each string in Σ receives a different colour; this induces an edge-colouring on $G(\Sigma)$ where each string is a monochromatic path. An *outer-rainbow* is a vertex $x \in V(G(\Sigma))$ incident with the outer face and for which the edges incident with x have different colours. Next we describe the basic operation in our obstruction-detecting algorithm.

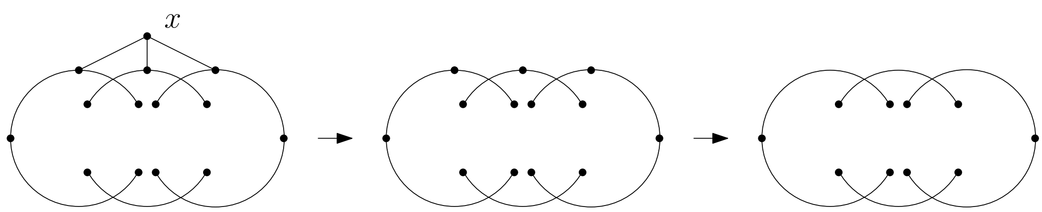

Figure 5 From Σ to $\Sigma - x$.

Outer-rainbow deletion. Given an outer-rainbow $x \in V(G(\Sigma))$, the instance $G(\Sigma - x)$ is defined by: first, removing x and the edges incident to x; second, suppressing the degree-2 vertices incident with edges of the same colour; and third, removing remaining degree-0 vertices (Figure 5 illustrates this process). Edge colours are preserved.

It is easy to verify that $G(\Sigma - x)$ is the planarization of an arrangement of strings. The colours removed by this operation are those belonging to strings that are paths of length 1 in $G(\Sigma)$ incident with x. Our obstruction-detecting algorithm relies on the following property:

(**) if x is an outer-rainbow of $G(\Sigma)$, then there is an obstruction in $G(\Sigma)$ not including x if and only if there is an obstruction in $G(\Sigma - x)$.

This property holds because cycles in $G(\Sigma) - x$ and in $G(\Sigma - x)$ are in 1-1 correspondence: two cycles correspond to each other if they are the same simple closed curve. This correspondence is obstruction-preserving.

Let us now describe the two subroutines in our algorithm. For this, we remark that an outer-rainbow of $G(\Sigma)$ is a rainbow for any cycle containing it.

Algorithm 1 Subroutine for detecting an obstruction through two outer-rainbows x and y.

(1) Find a cycle C through x and y whose edges are incident with the outer face of $G(\Sigma)$. If no such C exists, then output *No obstruction through x and y*. Else, go to Step 2.

(2) Find whether there is a third outer-rainbow $z \in V(C) \setminus \{x, y\}$. If such z exists, update $G(\Sigma) \longleftarrow G(\Sigma - z)$ and go to Step 1. If no such z exists, output C.

Correctness and running-time of Algorithm 1: If an obstruction through x and y exists, then x and y are in the same block (some authors use the term "biconnected component"). Since x and y are incident with the outer face, the outer boundary of the block containing x and y is the cycle C from Step 1. This C can be found by considering outer boundary walk W of $G(\Sigma)$ and then by finding whether x and y belong to the same non-edge block of W. Finding W is $O(m)$ and computing the blocks of W via a DFS takes $O(m)$ time.

In Step 2, if there is a third outer rainbow z in C, then no obstruction through x and y contains z. Property (**) justifies the update that takes $O(m)$ time.

A full run from Step 1 to Step 2 takes $O(m)$. Moving from Step 2 to Step 1 occurs $O(n)$ times. Thus, the total time for Algorithm 1 is $O(mn) = O(n^2)$.

Algorithm 2 Subroutine for detecting an obstruction through a single outer-rainbow x.

(1) Find a cycle C through x whose edges are incident with the outer face of $G(\Sigma)$. If no such C exists, output *No obstruction through x*. Else, go to Step 2.

(2) Find whether there is an outer-rainbow y in $V(C) \setminus \{x\}$. If no such y exists, output C. Else, apply Algorithm 1 to x and y; if there is an obstruction C' through x and y, then output C'. Else, update $G(\Sigma) \longleftarrow G(\Sigma - y)$ and go to Step 1.

Correctness and running-time of Algorithm 2: If $G(\Sigma)$ has an obstruction through x, then there is a non-edge block in $G(\Sigma)$ containing x. The outer boundary of this block is a cycle C through x having all edges incident with the outer face. As in Algorithm 1, Step 1 takes $O(m)$ time.

Detecting the existence of y in Step 2 is $O(m)$ because to detect rainbows in C, each edge incident with a vertex in $V(C)$ is verified at most twice. The update in Step 2 is justified by Property (**). Since Step 2 may use Algorithm 1, Step 2 takes $O(n^2)$ time. As moving from Step 2 to Step 1 occurs $O(n)$ times, the total running-time for Algorithm 2 is $O(n^3)$.

We are now ready for the algorithm to detect obstructions.

Algorithm 3 Detecting obstructions in $G(\Sigma)$.

(1) Find a cycle C having all edges incident with the outer face. If no such C exists, output *No obstruction*. Else, go to step 2.

(2) Find whether there is an outer rainbow $x \in V(C)$. If not, output C. Else apply Algorithm 2 to x; if there is an obstruction C' through x, output C'. Else, update $G(\Sigma) \longleftarrow G(\Sigma - x)$ and go to Step 1.

Correctness and running-time of Algorithm 3: If $G(\Sigma)$ has an obstruction, then it has a non-trivial block whose outer boundary is a cycle C as in Step 1. As before, C and x as in Step 2 can be found in $O(m)$ steps. If C has not outer rainbow x, then C is an obstruction; Property (**) justifies the update in Step 2.

Since Step 2 may use Algorithm 2, a full run of Steps 1 and 2 takes $O(n^3)$ time. Since Step 2 goes to Step 1 $O(n)$ times, the running-time of Algorithm 3 is $O(n^4)$.

Algorithm 3 and the constructive proof of Theorem 2 imply the following result (proved in the full version of this paper [4]).

▶ **Theorem 14.** *There is a polynomial-time algorithm to recognize and extend an ordinary set of strings that are extendible to an arrangement of pseudolines.*

6 Concluding remarks

In this work we characterized in Theorem 2 sets of strings that can be extended into arrangements of pseudolines. Moreover, we showed that the obstructions to pseudolinearity can be detected in $O(n^4)$ time, where n is the number of vertices in the planarization of the set of strings.

An easy consequence of Theorem 2 is the following (presented before as Theorem 1). We prove this result in the full version of this paper [4].

▶ **Theorem 15.** *Let D be a non-pseudolinear good drawing of a graph H. Then there is a subset S of edge-arcs in $\{D[e] : e \in E(H)\}$, such that each $\sigma \in S$ has a substring $\sigma' \subseteq \sigma$ for which $\bigcup_{\sigma \in S} \sigma'$ is one of the drawings represented in Figure 1.*

Theorem 2 can also be applied to find a short proof that pseudolinear drawings of K_n are characterized by forbidding the B-configuration (see Theorem 2.5.1 in [3]). This implies the characterizations of pseudolinear drawings of K_n presented in [2, 5, 6].

A drawing is *stretchable* if it is homeomorphic to a rectilinear drawing. There are pseudolinear drawings that are not stretchable. For instance, consider the Non-Pappus configuration in Figure 6. Nevertheless, as an immediate consequence of Thomassen's main result in [16], pseudolinear and stretchable drawings are equivalent, under the assumption that every edge is crossed at most once.

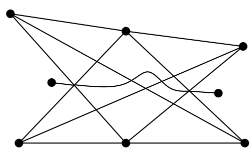

Figure 6 Non-Pappus configuration.

▶ **Corollary 16.** *A drawing D of a graph in which every edge is crossed at most once is stretchable if and only if it is pseudolinear.*

Proof. If a drawing D is stretchable then clearly it is pseudolinear. To show the converse, suppose that D is pseudolinear. Then D does not contain any obstruction, and in particular, neither of the B- and W-configurations in Figure 1 occurs in D. This condition was shown in [16] to be equivalent to being stretchable. ◀

One can construct more general examples of pseudolinear drawings that are not stretchable by considering non-strechable arrangements of pseudolines. However, such examples seem to inevitably have some edge with multiple crossings. This leads to a natural question.

▶ **Question 17.** *Is it true that if D is a pseudolinear drawing in which every edge is crossed at most twice, then D is stretchable?*

We believe that there are other instances where pseudolinearity characterizes stretchability of drawings. A drawing is *near planar* if the removal of one edge produces a planar graph. One instance, is the following result by Eades et al. that can be translated to the language of pseudolines:

▶ **Theorem 18** ([9]). *A drawing of a near-planar graph is stretchable if and only if the drawing induced by the crossed edges is pseudolinear.*

References

1 Bernardo M Ábrego and Silvia Fernández-Merchant. A lower bound for the rectilinear crossing number. *Graphs and Combinatorics*, 21(3):293–300, 2005.

2 Oswin Aichholzer, Thomas Hackl, Alexander Pilz, Birgit Vogtenhuber, and G Salazar. Deciding monotonicity of good drawings of the complete graph. In *Encuentros de Geometría Computacional*, pages 33–36. ., 2015.

3 Alan Arroyo. *On Geometric Drawings of Graphs*. PhD thesis, University of Waterloo, 2018.

4 Alan Arroyo, Julien Bensmail, and R Bruce Richter. Extending drawings of graphs to arrangements of pseudolines. *arXiv preprint*, 2018. `arXiv:1804.09317`.

5 Alan Arroyo, Dan McQuillan, R Bruce Richter, and Gelasio Salazar. Levi's lemma, pseudolinear drawings of, and empty triangles. *Journal of Graph Theory*, 87(4):443–459, 2018.

6 Martin Balko, Radoslav Fulek, and Jan Kynčl. Crossing numbers and combinatorial characterization of monotone drawings of K_n. *Discrete & Computational Geometry*, 53(1):107–143, 2015.

7 József Balogh, Jesús Leaños, Shengjun Pan, R Bruce Richter, and Gelasio Salazar. The convex hull of every optimal pseudolinear drawing of K_n is a triangle. *Australasian Journal of Combinatorics*, 38:155, 2007.

8 Daniel Bienstock and Nathaniel Dean. Bounds for rectilinear crossing numbers. *Journal of Graph Theory*, 17(3):333–348, 1993.

9 Peter Eades, Seok-Hee Hong, Giuseppe Liotta, Naoki Katoh, and Sheung-Hung Poon. Straight-line drawability of a planar graph plus an edge. *arXiv preprint*, 2015. `arXiv:1504.06540`.

10 Stefan Felsner. *Geometric graphs and arrangements: some chapters from combinatorial geometry*. Springer Science & Business Media, 2012.

11 Stefan Felsner and Jacob E Goodman. Pseudoline arrangements. In *Handbook of Discrete and Computational Geometry*, pages 125–157. Chapman and Hall/CRC, 2017.

12 César Hernández-Vélez, Jesús Leaños, and Gelasio Salazar. On the pseudolinear crossing number. *Journal of Graph Theory*, 84(3):155–162, 2017.

13 László Lovász, Katalin Vesztergombi, Uli Wagner, and Emo Welzl. Convex quadrilaterals and k-sets. *Contemporary Mathematics*, 342:139–148, 2004.

14 Nicolai E Mnëv. Varieties of combinatorial types of projective configurations and convex polytopes. *Doklady Akademii Nauk SSSR*, 283(6):1312–1314, 1985.

15 Nikolai E Mnëv. The universality theorems on the classification problem of configuration varieties and convex polytopes varieties. In *Topology and geometry – Rohlin seminar*, pages 527–543. Springer, 1988.

16 Carsten Thomassen. Rectilinear drawings of graphs. *Journal of Graph Theory*, 12(3):335–341, 1988.