Inserting One Edge into a Simple Drawing Is Hard

Alan Arroyo[1], Fabian Klute[2], Irene Parada[3], Raimund Seidel[4], Birgit Vogtenhuber[5](✉), and Tilo Wiedera[6]

[1] IST Austria, Klosterneuburg, Austria
alanmarcelo.arroyoguevara@ist.ac.at
[2] Utrecht University, Utrecht, The Netherlands
f.m.klute@uu.nl
[3] TU Eindhoven, Eindhoven, The Netherlands
i.m.de.parada.munoz@tue.nl
[4] Universität des Saarlandes, Saarbrücken, Germany
rseidel@cs.uni-saarland.de
[5] Graz University of Technology, Graz, Austria
bvogt@ist.tugraz.at
[6] Osnabrück University, Osnabrück, Germany
tilo.wiedera@uos.de

Abstract. A *simple drawing* $D(G)$ of a graph G is one where each pair of edges share at most one point: either a common endpoint or a proper crossing. An edge e in the complement of G can be *inserted* into $D(G)$ if there exists a simple drawing of $G + e$ extending $D(G)$. As a result of Levi's Enlargement Lemma, if a drawing is rectilinear (pseudolinear), that is, the edges can be extended into an arrangement of lines (pseudolines), then any edge in the complement of G can be inserted. In contrast, we show that it is NP-complete to decide whether one edge can be inserted into a simple drawing. This remains true even if we assume that the drawing is pseudocircular, that is, the edges can be extended to an arrangement of pseudocircles. On the positive side, we show that, given an arrangement of pseudocircles $\mathcal{A}$ and a pseudosegment σ, it can be decided in polynomial time whether there exists a pseudocircle Φ_σ extending σ for which $\mathcal{A} \cup \{\Phi_\sigma\}$ is again an arrangement of pseudocircles.

A. Arroyo—Funded by the Marie Skłodowska-Curie grant agreement No 754411.
F. Klute—This work was conducted while Fabian Klute was a member of the "Algorithms and Complexity Group" at TU Wien.
I. Parada—Partially supported by the Austrian Science Fund (FWF): W1230 and within the collaborative DACH project *Arrangements and Drawings* as FWF project I 3340-N35.
B. Vogtenhuber—Partially supported by Austrian Science Fund (FWF) within the collaborative DACH project *Arrangements and Drawings* as FWF project I 3340-N35.
T. Wiedera—Supported by the German Research Foundation (DFG) grant CH 897/2-2.

I. Adler and H. Müller (Eds.): WG 2020, LNCS 12301, pp. 325–338, 2020.
https://doi.org/10.1007/978-3-030-60440-0_26

1 Introduction

A *simple drawing* of a graph G (also known as *good drawing* or as *simple topological graph* in the literature) is a drawing $D(G)$ of G in the plane such that every pair of edges shares at most one point that is either a proper crossing or a common endpoint. In particular, no tangencies between edges are allowed, edges must not contain any vertices in their relative interior, and no three edges intersect in the same point. Simple drawings have received a great deal of attention in various areas of graph drawing, for example in connection with two long-standing open problems: the crossing number of the complete graph [30] and Conway's thrackle conjecture [26].

In this work, we study the problem of inserting an edge into a simple drawing of a graph. Given a simple drawing $D(G)$ of a graph $G = (V, E)$ and an edge e of the complement $\overline{G}$ of G we say that e can be *inserted* into $D(G)$ if there exists a simple drawing of $G' = (V, E \cup \{e\})$ that contains $D(G)$ as a subdrawing.

A *pseudoline arrangement* is an arrangement of simple biinfinite arcs, called *pseudolines*, such that every pair of pseudolines intersects in a single point that is a proper crossing. Similarly, an *arrangement of pseudocircles* is an arrangement of simple closed curves, called *pseudocircles*, such that every pair of pseudocircles intersects in either zero or two points, where in the latter case, both intersection points are proper crossings. A simple drawing $D(G)$ is called *pseudolinear* if the drawing of every edge can be extended to a pseudoline such that the extended drawing forms a pseudoline arrangement. Likewise, $D(G)$ is called *pseudocircular* if the drawing of every edge can be extended to a pseudocircle such that the extended drawing forms an arrangement of pseudocircles.

Pseudoline arrangements were introduced by Levi [24] in 1926 and have since been extensively studied; see for example [13]. One of the most fundamental results on pseudoline arrangements, nowadays well known as Levi's Enlargement Lemma, stems from Levi's original paper[1]. It states that, for any given pseudoline arrangement $\mathcal{L}$ and any two points p and q not on the same pseudoline of $\mathcal{L}$, it is always possible to insert a pseudoline through p and q into $\mathcal{L}$ such that the resulting arrangement is again a valid pseudoline arrangement.

From Levi's Enlargement Lemma, it immediately follows that given any pseudolinear drawing $D(G)$ and any set E^* of edges from $\overline{G}$, it is always possible to insert all edges from E^* into $D(G)$ such that the resulting drawing is again pseudolinear. To the contrary, as shown by Kynčl [23], this is in general not the case for simple drawings, not even if G is a matching plus two isolated vertices which are the endpoints of the edge to be inserted [22]. The latter implies that an analogous statement to Levi's Enlargement Lemma is not true for arrangements of pseudosegments (simple arcs that pairwise intersect at most once). Moreover, Arroyo, Derka, and Parada [2] recently showed that given a simple drawing $D(G)$ and a set E^* of edges from $\overline{G}$, it is NP-complete to decide whether E^* can be inserted into $D(G)$ (such that the resulting drawing is again simple). However,

[1] Also known as Levi's Extension Lemma. Several different proofs of Levi's Enlargement Lemma have been published since then [3, 14, 31–33].

the cardinality of E^* required for their hardness proof is linear in the size of the constructed graph. The main open problem posed in [2] is the complexity of deciding whether one single given edge e of $\overline{G}$ can be inserted into $D(G)$.

In this work, we show that this decision problem is NP-complete, even if G is a matching plus two isolated vertices which are the endpoints of e. This implies that, given an arrangement $\mathcal{S}$ of pseudosegments and two points p and q not on the same pseudosegment, it is NP-complete to decide whether it is possible to insert a pseudosegment from p to q into $\mathcal{S}$ such that the resulting arrangement is again a valid arrangement of pseudosegments (Sect. 2). On the positive side, we observe that the decision problem is fixed-parameter tractable (FPT) in the number of crossings of the drawing (Sect. 4).

Snoeyink and Hershberger [32] showed the following analogon to Levi's Enlargement Lemma for arrangements of pseudocircles: For any arrangement $\mathcal{A}$ of pseudocircles and any three points p, q, and r, not all of them on one pseudocircle of $\mathcal{A}$, there exists a pseudocircle Φ through p, q, and r such that $\mathcal{A} \cup \{\Phi\}$ is again an arrangement of pseudocircles. Refining our hardness proof, we show that the edge-insertion decision problem remains NP-complete when $D(G)$ is a pseudocircular drawing, regardless of whether the resulting drawing is required to be again pseudocircular or allowed to be any simple drawing. This holds even if we are in addition given an arrangement of pseudocircles extending $D(G)$. On the positive side, we show that, given an arrangement $\mathcal{A}$ of pseudocircles and a pseudosegment σ, it can be decided in polynomial time whether there exists an extension Φ_σ of σ to a simple closed curve such that $\mathcal{A} \cup \{\Phi_\sigma\}$ is again an arrangement of pseudocircles (Sect. 3).

One of the implications of the results presented in this paper concerns so-called saturated drawings [22]. A simple drawing $D(G)$ of a graph G is called *saturated* if no edge e from $\overline{G}$ can be inserted into $D(G)$. It is known that there are saturated simple drawings with a linear number of edges [16]. A natural question is to determine the complexity of deciding whether a simple drawing is saturated. Our hardness result implies that the straight-forward idea of testing whether $D(G)$ is saturated by checking for every edge in $\overline{G}$ whether it can be inserted into $D(G)$ is not feasible unless P = NP.

The problem of inserting an edge (or multiple edges or a star) into a planar graph has been extensively studied in the contexts of determining the crossing number of the resulting graph [6,29] and of finding a drawing of the resulting graph in which the original planar graph is drawn crossing-free and the drawing of the resulting graph has as few crossings as possible [10,11,15,28]. In relation to our work, a main difference is that we consider inserting edges into some given non-plane drawing of a graph. Furthermore, the question considered in this paper is strongly related to work on extending partial representations of graphs. Here, we are usually given a representation of a part of the graph G and are asked to extend it into a full representation of G such that the partial representation is a sub-representation of the full one. Recent years have seen a plethora of results in this topic [1,4,5,7–9,12,17–21,25,27].

Proofs of statements marked with $\star$ are deferred to the full version of this work.

2 Inserting One Edge into a Simple Drawing Is Hard

Theorem 1. *Given a simple drawing $D(G)$ of a graph $G = (V, E)$ and an edge uv of $\overline{G}$, it is* NP*-complete to decide whether uv can be inserted into $D(G)$, even if $V \setminus \{u, v\}$ induces a matching in G and u and v are isolated vertices.*

It is straightforward to verify that the problem is in NP (see Arroyo et al. [2] for a combinatorial description of our problem using the dual of the planarization of the drawing). We show NP-hardness via a reduction from 3SAT. Let $\phi(x_1, \ldots x_n)$ be a 3SAT-formula with *variables* $x_1, \ldots, x_n$ and set of *clauses* $\mathcal{C} = \{C_1, \ldots, C_m\}$. An occurrence of a variable x_i in a clause $C_j \in \mathcal{C}$ is called a *literal*. For convenience, we assume that in $\phi(x_1, \ldots, x_n)$, each clause has three (not necessarily different) literals. In a preprocessing step, we eliminate clauses with only positive or only negative literals via the transformation from Lemma 1.

Lemma 1 ($\star$). *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* `false` *is the constant value false):*

$$x_i \vee x_j \vee x_k \Rightarrow \begin{cases} x_k \vee y \vee \texttt{false} & (i) \\ x_i \vee x_j \vee \neg y & (ii) \end{cases} \quad \neg x_i \vee \neg x_j \vee \neg x_k \Rightarrow \begin{cases} \neg x_i \vee \neg x_j \vee y & (iii) \\ \neg x_k \vee \neg y \vee \texttt{false} & (iv) \end{cases}$$

After the preprocessing, we have a *transformed* 3SAT-formula where each clause is of one of the following four types: Type (i) two positive literals and one constant `false`; Type (ii) one negative and two positive literals; Type (iii) one positive and two negative literals, and finally, Type (iv) two negative literals and one constant `false`.

Given a transformed 3SAT-formula $\phi = \phi(x_1, \ldots, x_n)$ with set of clauses $\mathcal{C} = \{C_1, \ldots, C_m\}$, satisfiability of ϕ will correspond to being able to insert a given edge uv into a simple drawing D of a matching constructed from the formula ϕ. The main idea of the reduction is that the variable and clause gadgets in D act as "barriers" inside a simple closed region R of D, in which we need to insert a simple arc γ from one side to the other to connect u and v. Crossing a barrier in some way imposes constraints on how or whether we can cross other barriers afterwards.

To simplify the description, we first focus our attention to the inside of the simple closed region R. We assume that γ cannot cross the boundary of R. In the following we use two lines, named λ and μ, to bound the regions in which a variable and clause gadget will be placed. Particularly, these lines will be identified with opposite segments on R's boundary.

Variable Gadget. A variable gadget W is bounded from above by a horizontal line λ and from below by a horizontal line μ. Additionally, it contains a vertical segment κ between λ and μ, a set P of pairwise non-crossing arcs (parts of later-defined edges), each with one endpoint on κ and the other endpoint on μ, and a set N of pairwise non-crossing arcs, each with one endpoint on κ and the other endpoint on λ. On κ, all the endpoints of arcs in P lie above all the endpoints of arcs in N, implying that every arc in P crosses every arc in N. Finally, we choose two points u and v such that u is to the left of all arcs in W and v is to the right of them; see Fig. 1 for an illustration. The arcs in P and N correspond to positive and negative appearances of the variable, respectively.

Lemma 2 (⋆). *Let W be a variable gadget. Any arc between the horizontal lines λ and μ that connects u and v crosses either all arcs in P or all arcs in N.*

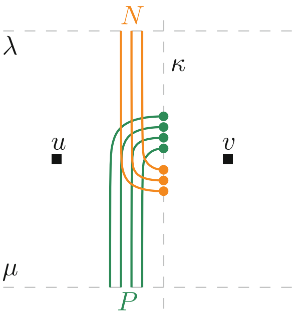

Fig. 1. Variable gadget. The orange arcs belong to N, the green ones to P. (Color figure online)

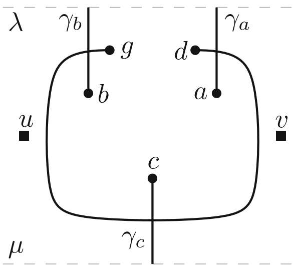

Fig. 2. Clause gadget.

Clause Gadget. Similar to a variable gadget, a clause gadget K is bounded from above and below by two horizontal lines λ and μ, respectively. Additionally, it contains three horizontal arcs (parts of later-defined edges) γ_a, γ_b, and γ_c, where the former two have one endpoint on λ and the latter has one endpoint on μ. On λ, the endpoint of γ_a lies to the right of the one of γ_b. The other endpoints of γ_a, γ_b, and γ_c are called a, b, and c, respectively. None of these three arcs cross. Moreover, K contains two points d and g and an edge dg that crosses γ_a, γ_c, and γ_b in that order when traversed from d to g. Notice that we do not require any specific rotation of the crossings of dg with γ_a and γ_b (where the rotation is the clockwise order of the endpoints of the crossing arcs). However, to simplify the description, we assume that the rotations of the crossings are as in Fig. 2. The rotation of the crossing of dg with γ_c is forced by the order of the crossings along dg. Finally, we again choose two points u and v such that u is to the left of all arcs in K and v is to the right of them; see Fig. 2 for an illustration.

Lemma 3 (⋆). *Let K be a clause gadget. Any arc uv between the horizontal lines λ and μ that connects u and v crosses either dg twice or at least one of the arcs γ_a, γ_b, and γ_c.*

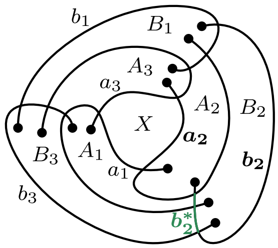

Fig. 3. The simple drawing $\circledcirc$ presented in [22].

The Reduction. Let $\phi(x_1, \ldots, x_n)$ be a transformed 3SAT-formula with clause set $\mathcal{C} = \{C_1, \ldots, C_m\}$ (each clause being of one of the four types identified above). To build our reduction we need one more gadget. First, we introduce the following simple drawing introduced by Kynčl et al. [22, Figure 11] and depicted in Fig. 3. Here, we denote this drawing by $\circledcirc$. Following the notation by Kynčl et al., we denote its six arcs by a_1, a_2, a_3, b_1, b_2, and b_3; and its eight cells by X, A_1, A_2, A_3, B_1, B_2, B_3, and Y; see Fig. 3 for an illustration. The core property $\mathcal{P}$ of $\circledcirc$ is that it is not possible to insert an edge between a point in cell X and another point in cell Y such that the result is a simple drawing [22, Lemma 15].

For our reduction, we first choose two arbitrary points u and v in the cells X and B_2 and insert them as vertices into $\circledcirc$. Let $\circledcirc'$ be the obtained drawing. Further, let b_2^* be the part of the arc b_2 between the crossing point of b_2 and a_2 and the crossing point of b_2 and b_3, see again Fig. 3.

Lemma 4 ($\star$). *The edge uv cannot be inserted into $\circledcirc'$ without crossing b_2^*.*

The final piece we need for our reduction is a set F of $m^I + m^{IV} + 4$ arcs that we insert into $\circledcirc'$, where m^I is the number of clauses of Type (i) and m^{IV} the number of clauses of Type (iv). For an arc $f \in F$ we will place one of its endpoints on a horizontal line κ_F inside A_2 and the other one inside B_2. The only crossings of f with $\circledcirc'$ are with the arcs a_2, a_1, b_3, and b_2, in that order, when traversing f from its endpoint on κ_F to its endpoint in B_2. Furthermore, when f is traversed in that direction, it crosses from A_2 to A_1, from A_1 to B_3, from B_3 to Y, and from Y to B_2.

Consider the $m^I + m^{IV} + 4$ endpoints on κ_F sorted from left to right. We denote by f_j the arc in F incident with the j-th such endpoint. When traversing b_2 from its endpoint in A_2 to its endpoint in B_1, the crossings of arcs in F with b_2 appear in the same order as their endpoints on κ_F. More precisely, the crossings of b_2, when b_2 is traversed in that direction, are with a_2, a_1, b_3, f_1, f_2, $\ldots$, $f_{|F|}$, and b_1, in that order.

The arcs f_{m^I+1}, f_{m^I+2}, f_{m^I+3}, and f_{m^I+4} will behave differently than the other arcs in F. In the following, we denote these four arcs by r_2, r_1, ℓ_1, and ℓ_2, respectively. There are only two crossings between arcs in F, namely, between r_1 and r_2, and between ℓ_1 and ℓ_2, and both these crossings are inside B_2. These four crossing arcs divide B_2 into three regions. Let R denote the region with b_2^* on its boundary; let R_r denote the (other) region incident with the crossing between r_1 and r_2; and let R_ℓ denote the (other) region incident with the crossing between ℓ_1 and ℓ_2. Arcs r_1, r_2, ℓ_1, and ℓ_2 must be drawn such that the vertex v lies in R; see the red arcs in Fig. 4 for an illustration. The precise endpoints of the edges in $F \setminus \{r_1, r_2, \ell_1, \ell_2\}$ will be fixed when we insert the clause gadgets.

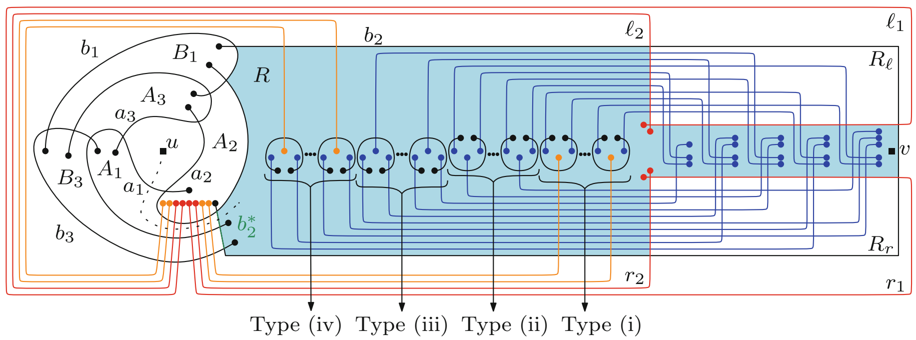

Fig. 4. Illustration of the reduction. (Color figure online)

Lemma 5 ($\star$). *The edge uv cannot be inserted into $\circledcirc'$ without crossing every arc in F in the closure of A_1 or of B_3.*

It remains to insert inside R the clause and variable gadgets and precisely define the endpoints of arcs in $F \setminus \{\ell_1, \ell_2, r_1, r_2\}$. For simplicity, we first insert the variable gadgets and then the clause gadgets. The idea is that each clause and variable gadget is inserted in R separating b_2^* from v. This is done by identifying the endpoints that were lying on λ or μ with points on ℓ_1, ℓ_2, r_1, r_2, or b_2. As a result, Lemmas 2 and 3 can be applied to the arc that we insert connecting u and v in the final drawing, since it has to cross b_2^* by Lemma 4.

We now insert the variable gadgets into R. Let $W^{(i)}$ be the variable gadget corresponding to variable x_i. For a gadget $W^{(i)}$, the arcs in N are drawn such that the endpoints on λ lie on the part of ℓ_1 that bounds R. The arcs in P are drawn similarly, but with the endpoints on μ lying on the part of r_1 that bounds R. Moreover, we identify vertex v in the gadget with vertex v in $\circledcirc'$. Gadgets corresponding to different variables are inserted without crossing each other. We now specify how they are inserted relative to each other. As we traverse ℓ_1 from its endpoint on κ_F to its endpoint in R, we encounter the endpoints of arcs in $W^{(i)}$ before the endpoints of arcs in $W^{(i+1)}$. Analogously, as we traverse r_1 from its endpoint on κ_F to its endpoint in R, we encounter the endpoints of arcs in $W^{(i)}$ before the endpoints of arcs in $W^{(i+1)}$. See Fig. 4 for an illustration.

In a similar way we insert the clause gadgets. Let $K^{(j)}$ be the clause gadget corresponding to clause C_j. If C_j is of Type (i), $K^{(j)}$ is inserted such that the endpoints on λ lie on the part of ℓ_2 that bounds R. If C_j is the j'-th clause of Type (i), we identify c with the endpoint of the arc $f_{j'}$. Similarly, if C_j is of Type (iv), $K^{(j)}$ is inserted such that the endpoints on λ lie on the part of r_2 that bounds R. If C_j is the j'-th clause of Type (iv), we identify c with the endpoint of the arc $f_{m^I+4+j'}$. If C_j is of Type (ii), $K^{(j)}$ is inserted such that the endpoints on λ lie on the part of ℓ_2 that bounds R and the endpoint on μ lies on the part of r_2 that bounds R. Similarly, if C_j is of Type (iii), $K^{(j)}$ is inserted such that the endpoint on μ lies on the part of ℓ_2 that bounds R and the endpoints on λ lie on the part of r_2 that bounds R. The crossings in R of arcs from different clause gadgets are of arcs with an endpoint in r_2 with arcs in $\{f_j : 1 \leq j \leq m^I\}$.

We now specify how different clause gadgets are inserted relative to each other. As we traverse ℓ_2 from its endpoint on κ_F to its endpoint in R, we first encounter the endpoints of arcs corresponding to Type (iii) clauses, followed by the ones corresponding to Type (ii) clauses, and finally the ones corresponding to Type (i) clauses. Analogously, as we traverse r_2 from its endpoint on κ_F to its endpoint in R, we first encounter the endpoints of arcs corresponding to Type (iv) clauses, followed by the ones corresponding to Type (iii) clauses, and finally the ones corresponding to Type (ii) clauses. Moreover, as we traverse ℓ_2 and r_2 in the specified directions, the endpoints of arcs corresponding to the j'-th clause of a certain type are encountered before the endpoints of arcs corresponding to the $(j'-1)$-st clause of this type. An illustration can be found in Fig. 4.

Finally, we connect arcs from variable and clause gadgets inside the regions R_ℓ and R_r. This is done such that if a literal in a clause is x_k then the corresponding arc in the clause gadget, that has an endpoint on ℓ_2, is connected with an arc in N of the gadget $W^{(k)}$, that has an endpoint on ℓ_1. Thus, these connections can lie in R_ℓ. Analogously, if a literal in a clause is $\neg x_k$ then the corresponding arc in the clause gadget, that has an endpoint on r_2, is connected with an arc in P of the gadget $W^{(k)}$, that has an endpoint on r_1. Thus, these connections can lie in R_r. Since, without loss of generality, we can assume that R_ℓ and R_r are convex regions and the endpoints we want to connect are pairwise distinct points on the boundaries of those regions, the connections can be drawn as straight-line segments. (For clarity, in Fig. 4, these connections have one bend per arc.) Therefore, there is at most one crossing between each pair of connecting arcs.

Each connecting arc is concatenated with the arcs in a variable and in a clause gadget that it joins. These concatenated arcs are edges in our drawing that have one endpoint in a variable gadget and the other one in a clause gadget. By construction, each such edge corresponds to a literal in the formula ϕ and each pair of them crosses at most once. Similarly, the arcs in $F \setminus \{\ell_1, \ell_2, r_1, r_2\}$ have one endpoint in a clause gadget and also define edges in our final drawing that we denote by the same names as the corresponding arcs.

We now have all the pieces that constitute our final drawing. It consists of (i) the simple drawing $\circledcirc'$; (ii) the edges $f_i \in F$ drawn as the described arcs (with their endpoints as vertices); (iii) the edges corresponding to literals (with their endpoints as vertices); and (iv) the edges dg in each clause gadget (with d and g as vertices). Observe that the constructed drawing is a simple drawing, as it is the drawing of a matching (plus the vertices u and v) and, by construction, any two edges cross at most once.

It remains to show that the presented construction is a valid reduction.

Lemma 6 (⋆). *The above construction is a poly-time reduction from* 3SAT *to the problem of deciding whether an edge can be inserted into a simple drawing.*

Remarks and Extensions. As our reduction from 3SAT constructs a simple drawing $D(G)$ of a matching, the general problem is NP-hard even if G is as sparse as possible. We remark that if we do not require G to be a matching, our variable gadget can be simplified by identifying all the vertices on κ and removing the crossings between edges in N and P. Moreover, from the constructed drawing $D(G)$, one can produce an equivalent instance that is connected: This is done by inserting an apex vertex into an arbitrary cell of the drawing, and then subdividing its incident edges so that the resulting drawing D^* is simple. If uv can be inserted into $D(G)$ then it can be inserted also into D^*. Finally, it is possible to show that the simple drawings produced by our reduction are pseudocircular implying the following result.

Corollary 1 (⋆). *Given a pseudocircular drawing $D(G)$ of a graph $G = (V, E)$ and an edge uv of $\overline{G}$, it is NP-complete to decide whether uv can be inserted into $D(G)$, even if an arrangement of pseudocircles extending the drawing of the edges in $D(G)$ is provided.*

3 Extending an Arrangement of Pseudocircles Is Easy

In the previous section we proved that deciding whether an edge can be inserted into a pseudocircular drawing such that the result is a simple (or a pseudocircular) drawing is hard. In this section we focus on extending arrangements instead of drawings of graphs. Snoeyink and Hershberger [32] showed that given an arrangement $\mathcal{A}$ of pseudocircles and three points, not all three on the same pseudocircle, one can find a pseudocircle Φ through the three points such that $\mathcal{A} \cup \{\Phi\}$ is again an arrangement of pseudocircles. Now, given any arrangement $\mathcal{A}$ and a pseudosegment σ intersecting each pseudocircle in $\mathcal{A}$ at most twice, it is not always possible to extend σ to a pseudocircle $\Phi_\sigma \supset \sigma$ such that $\mathcal{A} \cup \{\Phi_\sigma\}$ is again an arrangement of pseudocircles. Two examples are shown in Figs. 5 and 6. In either, any pseudocircle Φ_σ extending σ crosses one red or blue pseudocircle at least four times. However, we show in the following that the extension decision question can be answered in polynomial time:

Theorem 2. *Given an arrangement $\mathcal{A}$ of n pseudocircles and a pseudosegment σ intersecting each pseudocircle in $\mathcal{A}$ at most twice, it can be decided in time polynomial in n whether there exists an extension of σ to a pseudocircle Φ_σ such that $\mathcal{A} \cup \{\Phi_\sigma\}$ is an arrangement of pseudocircles.*

Proof. Throughout this proof we write $\overline{R} := \mathbb{R}^2 \setminus R$ for the *complement* of a set $R \subseteq \mathbb{R}^2$. An arrangement (of pseudocircles) partitions the plane into *vertices* (0-dimensional cells), *edges* (1-dimensional cells), and *faces* (2-dimensional cells). Since tangencies are not allowed, all vertices are proper crossings. Two arrangements are *combinatorially equivalent* (or, *isomorphic*) if the corresponding cell complexes are isomorphic, that is, if there is an incidence- and dimension-preserving bijection between their cells. By possibly transforming $\mathcal{A}$ into an isomorphic arrangement while preserving the incidences of σ, we can assume without loss of generality that σ is a horizontal segment. Let u and v be the left and right endpoints of σ, respectively. Further, we can assume that u is incident with the unbounded cell and that the intersection points of σ with the pseudocircles in $\mathcal{A}$ are all proper crossings. Our algorithm aims to compute a pseudocircle $\Phi_\sigma = \sigma \cup \sigma'$ such that $\mathcal{A} \cup \{\Phi_\sigma\}$ is an arrangement of pseudocircles, or determine that no such σ' exists. We call σ' an *extension* of σ.

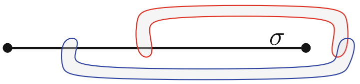

Fig. 5. Obstruction where all pseudocircles intersect σ twice.

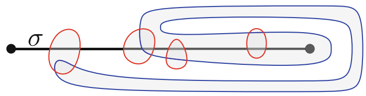

Fig. 6. Obstruction where one pseudocircle intersects σ only once.

We partition the set of pseudocircles of $\mathcal{A}$ into three sets $\mathcal{C}_0$, $\mathcal{C}_1$, and $\mathcal{C}_2$, where for each $i \in \{0, 1, 2\}$, $\mathcal{C}_i$ is the set of pseudocircles in $\mathcal{A}$ crossing σ exactly i times. Note that u lies outside all pseudocircles $\phi \in \mathcal{A}$ while v lies outside of all $\phi \in \mathcal{C}_0 \cup \mathcal{C}_2$ and inside all $\phi \in \mathcal{C}_1$, that is, each $\phi \in \mathcal{C}_1$ separates u and v. Further, an extension σ' must not cross any $\phi \in \mathcal{C}_2$, it needs to cross every $\phi \in \mathcal{C}_1$ exactly once, and it can cross each $\phi \in \mathcal{C}_0$ either twice or not at all.

The idea is to construct a finite sequence $R_0 \subset R_1 \subset \ldots$ of closed subsets of $\mathbb{R}^2$, each consisting of cells of $\mathcal{A} \cup \sigma$ that cannot be reached by σ'. Each set R_i will be a simply connected closed region of $\mathbb{R}^2$ with both u and v on its boundary. Further, for each R_i and each $\phi \in \mathcal{C}_0$, we will maintain the invariant that $\mathrm{int}(\phi) \cap \overline{R_i}$ is either a connected region or empty, where $\mathrm{int}(\phi)$ denotes the interior of the bounded area enclosed by ϕ. (Note that $\mathrm{int}(\phi) \cap \overline{R_i}$ is connected if and only if $R_i \setminus \mathrm{int}(\phi)$ is connected.) The construction will either end by determining that σ cannot be extended, or with a set R_m such that routing σ' closely along the boundary of R_m gives a valid extension of σ.

Let R_0' be the union of σ and all the closed disks bounded by the pseudocircles in $\mathcal{C}_2$ and consider the faces induced by R_0'. Since u is incident with the unbounded cell of R_0', and since σ' must not intersect the interior of R_0', σ' cannot reach any bounded face of R_0'. Let R_0 be the closure of the union of these bounded faces and σ. We may assume that $v \in \partial R_0$, as otherwise no extension σ' exists and we are done.

To see that the invariant holds for R_0, assume that there exists a pseudocircle $\phi \in \mathcal{C}_0$ such that $R_0 \setminus \mathrm{int}(\phi)$ is not connected. As ϕ does not intersect σ, there exists a component D of $R_0 \setminus \mathrm{int}(\phi)$ that is disjoint from σ. Further, as $\mathrm{int}(\phi)$ is simply connected, $D \cap \partial R_0 \neq \emptyset$. Moreover, any point x on $\partial D \cap \partial R_0$ lies on some circle $\phi_x \in \mathcal{C}_2$. On the other hand, any path from a point of σ to x must enter and leave $\mathrm{int}(\phi)$ and hence intersect ϕ at least twice. As ϕ_x intersects σ twice and lies in R_0, we get that ϕ_x intersects ϕ in at least four points, a contradiction.

For the iterative step, consider the arrangement $\mathcal{A}_i^\phi$ formed by ∂R_i and a pseudocircle $\phi \in \mathcal{C}_0 \cup \mathcal{C}_1$, and the cells of it that lie in $\overline{R_i}$. If $\phi \in \mathcal{C}_1$ and an extension σ' exists, then the only two such cells that can be intersected by σ' are the ones incident to u and v, respectively. Similarly, if $\phi \in \mathcal{C}_0$, then σ' can only intersect the cell(s) incident to u and v, plus the (by the invariant) unique cell $\text{int}(\phi) \cap \overline{R_i}$. In both cases, all other cells of this arrangement should be added to the forbidden area. We denote all cells $\mathcal{A}_i^\phi \cap \overline{R_i}$ that can possibly be intersected by σ' as *reachable* (by σ') and all other cells as *unreachable* (by σ').

Assume that there exists some pseudocircle $\phi \in \mathcal{C}_0 \cup \mathcal{C}_1$ such that the arrangement $\mathcal{A}_i^\phi$ of ϕ and ∂R_i contains unreachable cells. Then we obtain R'_{i+1} by adding all those cells to R_i. If v lies in a bounded face of $\overline{R'_{i+1}}$, then no extension σ' exists and we are done. Otherwise, $R_{i+1} = R'_{i+1}$ is a simply connected region that has both u and v on its boundary. It remains to show that the invariant is still maintained for R_{i+1}.

Lemma 7 **($\star$).** *If R_i fulfills the invariant and u and v both lie in the unbounded region of R'_{i+1} then R_{i+1} also fulfills the invariant.*

Now assume that both u and v lie on the boundary of all sets R_i constructed in this way. Then the iterative process stops with a set R_m where for each $\phi \in \mathcal{C}_0 \cup \mathcal{C}_1$, all cells in the arrangement $\mathcal{A}_m^\phi$ of ϕ and ∂R_m that are contained in $\overline{R_m}$ are reachable by σ'. Note that $m = O(n^4)$ as $\mathcal{A}$ has $\Theta(n^4)$ cells, as in every iteration i, at least one cell of $\mathcal{A}$ has been added to R_i, and as each cell of $\mathcal{A}$ is added at most once. Consider a path P from u to v in $\overline{R_m}$ that is routed closely along the boundary ∂R_m (note that there are two different such paths). Then for any $\phi \in \mathcal{C}_1$, P intersects exactly two cells of $\mathcal{A}_m^\phi$, namely, the ones incident to u and v, respectively. Hence P crosses ϕ exactly once. Similarly, for any $\phi \in \mathcal{C}_0$, the path P intersects at most three cells of $\mathcal{A}_m^\phi$, namely, the one(s) incident to u and v plus possibly the cell $\text{int}(\phi) \cap \overline{R_m}$, which is one cell by the invariant. Hence P crosses ϕ at most twice. Thus $\sigma' = P$ is a valid extension for σ, which completes the correctness proof.

Note that computing R_0 and σ' (in case that the algorithm didn't terminate with a negative answer before) can be done in poly-time. Also, for each R_i and each $\phi \in \mathcal{C}_0 \cup \mathcal{C}_1$, the set of unreachable cells of $\mathcal{A}_i^\phi$ can be determined in poly-time. As we have $O(n^4)$ iteration steps, we can hence compute R_m from R_0 (or determine that σ is not extendible) in poly-time, which concludes the proof.

As an immediate consequence of Theorem 2 we have the following result:

Corollary 2. *Given an arrangement $\mathcal{A}$ of n pseudocircles and a pseudosegment σ, it can be decided in polynomial time whether σ can be extended to a pseudocircle $\Phi_\sigma \supset \sigma$ such that $\mathcal{A} \cup \{\Phi_\sigma\}$ is an arrangement of pseudocircles.*

4 FPT-Algorithm for Bounded Number of Crossings

In this section we show that for drawings with a bounded number of crossings it can be decided in FPT-time whether an edge can be inserted. Given a simple drawing $D(G)$ with k crossings, one can construct a *kernel* of size $O(k)$ by exhaustively removing isolated vertices and uncrossed edges from $D(G)$. For a simple drawing $D(G)$ of a graph $G = (V, E)$ and $e \in E$, let $D(G - e)$ be the subdrawing of $D(G)$ without the drawing of e. Similarly, for an isolated vertex $u \in V$ let $D(G - u)$ be the subdrawing of $D(G)$ without the drawing of u.

Observation 1. *Given a simple drawing $D(G)$ of a graph $G = (V, E)$ and an isolated vertex $w \in V$, an edge uv of $\overline{G}$ can be inserted into $D(G)$ if and only if uv can be inserted into $D(G - w)$.*

Lemma 8. (⋆). *Given a simple drawing $D(G)$ of a graph $G = (V, E)$ and an edge $e \in E$ that is uncrossed in $D(G)$, an edge uv of $\overline{G}$ can be inserted into $D(G)$ if and only if uv can be inserted into $D(G - e)$.*

Theorem 3. (⋆). *Given a simple drawing $D(G)$ of a graph $G = (V, E)$ and an edge uv of $\overline{G}$, there is an* FPT*-algorithm in the number k of crossings in $D(G)$ for deciding whether uv can be inserted into $D(G)$.*

Acknowledgements. This work was started during the 6th Austrian-Japanese-Mexican-Spanish Workshop on Discrete Geometry in Juni 2019 in Austria. We thank all the participants for the good atmosphere as well as discussions on the topic. Also, we thank Jan Kynčl for sending us remarks on the first arXiv version of this work and an anonymous referee for further helpful comments.

References

1. Angelini, P., et al.: Testing planarity of partially embedded graphs. ACM Trans. Algorithms **11**(4), 32:1–32:42 (2015). https://doi.org/10.1145/2629341
2. Arroyo, A., Derka, M., Parada, I.: Extending simple drawings. In: Archambault, D., Tóth, C.D. (eds.) GD 2019. LNCS, vol. 11904, pp. 230–243. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-35802-0_18
3. Arroyo, A., McQuillan, D., Richter, R.B., Salazar, G.: Levi's lemma, pseudolinear drawings of K_n, and empty triangles. J. Graph Theory **87**(4), 443–459 (2018). https://doi.org/10.1002/jgt.22167
4. Bagheri, A., Razzazi, M.: Planar straight-line point-set embedding of trees with partial embeddings. Inf. Process. Lett. **110**(12–13), 521–523 (2010). https://doi.org/10.1016/j.ipl.2010.04.019
5. Brückner, G., Rutter, I.: Partial and constrained level planarity. In: Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017), pp. 2000–2011. SIAM (2017). https://doi.org/10.1137/1.9781611974782.130
6. Cabello, S., Mohar, B.: Adding one edge to planar graphs makes crossing number and 1-planarity hard. SIAM J. Comput. **42**(5), 1803–1829 (2013). https://doi.org/10.1137/120872310
7. Chaplick, S., Dorbec, P., Kratochvíl, J., Montassier, M., Stacho, J.: Contact representations of planar graphs: extending a partial representation is hard. In: Kratsch, D., Todinca, I. (eds.) WG 2014. LNCS, vol. 8747, pp. 139–151. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-12340-0_12
8. Chaplick, S., Fulek, R., Klavík, P.: Extending partial representations of circle graphs. J. Graph Theory **91**(4), 365–394 (2019). https://doi.org/10.1002/jgt.22436

9. Chaplick, S., Guśpiel, G., Gutowski, G., Krawczyk, T., Liotta, G.: The partial visibility representation extension problem. Algorithmica **80**(8), 2286–2323 (2017). https://doi.org/10.1007/s00453-017-0322-4
10. Chimani, M., Gutwenger, C., Mutzel, P., Wolf, C.: Inserting a vertex into a planar graph. In: Proceedings of the 20th ACM-SIAM Symposium on Discrete Algorithms (SODA 2009), pp. 375–383. SIAM (2009). https://doi.org/10.1137/1.9781611973068.42
11. Chimani, M., Hlinený, P.: Inserting multiple edges into a planar graph. In: Proceedings of the 32nd International Symposium on Computational Geometry (SoCG 2016), vol. 51, pp. 30:1–30:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016). https://doi.org/10.4230/LIPIcs.SoCG.2016.30
12. Da Lozzo, G., Di Battista, G., Frati, F.: Extending upward planar graph drawings. In: Friggstad, Z., Sack, J.-R., Salavatipour, M.R. (eds.) WADS 2019. LNCS, vol. 11646, pp. 339–352. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-24766-9_25
13. Felsner, S., Goodman, J.E.: Pseudoline arrangements. In: Tóth, C.D., O'Rourke, J., Goodman, J.E. (eds.) Handbook of Discrete and Computational Geometry, 3rd edn., pp. 125–157. CRC Press, Boca Raton. (2017)
14. Grünbaum, B.: Arrangements and spreads. AMS (1972)
15. Gutwenger, C., Mutzel, P., Weiskircher, R.: Inserting an edge into a planar graph. Algorithmica **41**(4), 289–308 (2005). https://doi.org/10.1007/s00453-004-1128-8
16. Hajnal, P., Igamberdiev, A., Rote, G., Schulz, A.: Saturated simple and 2-simple topological graphs with few edges. J. Graph Algorithms Appl. **22**(1), 117–138 (2018). https://doi.org/10.7155/jgaa.00460
17. Jelínek, V., Kratochvíl, J., Rutter, I.: A Kuratowski-type theorem for planarity of partially embedded graphs. Comput. Geom. Theory Appl. **46**(4), 466–492 (2013). https://doi.org/10.1016/j.comgeo.2012.07.005. Special Issue on the 27th Annual Symposium on Computational Geometry (SoCG'11)
18. Klavík, P., Kratochvíl, J., Krawczyk, T., Walczak, B.: Extending partial representations of function graphs and permutation graphs. In: Epstein, L., Ferragina, P. (eds.) ESA 2012. LNCS, vol. 7501, pp. 671–682. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33090-2_58
19. Klavík, P., et al.: Extending partial representations of proper and unit interval graphs. Algorithmica **77**(4), 1071–1104 (2016). https://doi.org/10.1007/s00453-016-0133-z
20. Klavík, P., Kratochvíl, J., Otachi, Y., Saitoh, T.: Extending partial representations of subclasses of chordal graphs. Theor. Comput. Sci. **576**, 85–101 (2015). https://doi.org/10.1016/j.tcs.2015.02.007
21. Klavík, P., Kratochvíl, J., Otachi, Y., Saitoh, T., Vyskočil, T.: Extending partial representations of interval graphs. Algorithmica **78**(3), 945–967 (2016). https://doi.org/10.1007/s00453-016-0186-z
22. Kynčl, J., Pach, J., Radoičić, R., Tóth, G.: Saturated simple and k-simple topological graphs. Comput. Geom. Theory Appl. **48**(4), 295–310 (2015). https://doi.org/10.1016/j.comgeo.2014.10.008
23. Kynčl, J.: Improved enumeration of simple topological graphs. Discrete Comput. Geom. **50**(3), 727–770 (2013). https://doi.org/10.1007/s00454-013-9535-8
24. Levi, F.: Die Teilung der projektiven Ebene durch Gerade oder Pseudogerade. Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig, Mathematisch-Physische Klasse **78**, 256–267 (1926). (in German)

25. Mchedlidze, T., Nöllenburg, M., Rutter, I.: Extending convex partial drawings of graphs. Algorithmica **76**(1), 47–67 (2015). https://doi.org/10.1007/s00453-015-0018-6
26. Pach, J., Brass, P., Moser, W.O.J.: Research Problems in Discrete Geometry. Springer, New York (2005). https://doi.org/10.1007/0-387-29929-7
27. Patrignani, M.: On extending a partial straight-line drawing. Int. J. Found. Comput. Sci. **17**(5), 1061–1070 (2006). https://doi.org/10.1142/S0129054106004261
28. Radermacher, M., Rutter, I.: Inserting an edge into a geometric embedding. In: Biedl, T., Kerren, A. (eds.) GD 2018. LNCS, vol. 11282, pp. 402–415. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-04414-5_29
29. Riskin, A.: The crossing number of a cubic plane polyhedral map plus an edge. Studia Scientiarum Mathematicarum Hungarica **31**(4), 405–414 (1996)
30. Schaefer, M.: Crossing Numbers of Graphs. CRC Press, Boca Raton (2018)
31. Schaefer, M.: A proof of Levi's extension lemma. ArXiv e-Prints (2019)
32. Snoeyink, J., Hershberger, J.: Sweeping arrangements of curves. In: Discrete and Computational Geometry: Papers from the DIMACS Special Year, vol. 6, pp. 309–350. DIMACS/AMS (1991). https://doi.org/10.1090/dimacs/006/21
33. Sturmfels, B., Ziegler, G.M.: Extension spaces of oriented matroids. Discrete Comput. Geom. **10**(1), 23–45 (1993). https://doi.org/10.1007/BF02573961