An Algorithmic Meta-Theorem for Graph Modification to Planarity and FOL

Fedor V. Fomin
Department of Informatics, University of Bergen, Norway
fedor.fomin@ii.uib.no

Petr A. Golovach
Department of Informatics, University of Bergen, Norway
petr.golovach@ii.uib.no

Giannos Stamoulis
Department of Informatics and Telecommunications,
National and Kapodistrian University of Athens, Greece
Inter-university Postgraduate Programme "Algorithms, Logic, and Discrete Mathematics" (ALMA), Athens, Greece
giannos95@gmail.com

Dimitrios M. Thilikos
LIRMM, Univ. Montpellier, CNRS, Montpellier, France
sedthilk@thilikos.info

Abstract

In general, a *graph modification problem* is defined by a graph modification operation $\boxtimes$ and a target graph property $\mathcal{P}$. Typically, the modification operation $\boxtimes$ may be vertex removal, edge removal, edge contraction, or edge addition and the question is, given a graph G and an integer k, whether it is possible to transform G to a graph in $\mathcal{P}$ after applying k times the operation $\boxtimes$ on G. This problem has been extensively studied for particilar instantiations of $\boxtimes$ and $\mathcal{P}$. In this paper we consider the general property $\mathcal{P}_\phi$ of being planar and, moreover, being a model of some First-Order Logic sentence ϕ (an FOL-sentence). We call the corresponding meta-problem Graph $\boxtimes$-Modification to Planarity and ϕ and prove the following algorithmic meta-theorem: there exists a function $f : \mathbb{N}^2 \to \mathbb{N}$ such that, for every $\boxtimes$ and every FOL sentence ϕ, the Graph $\boxtimes$-Modification to Planarity and ϕ is solvable in $f(k, |\phi|) \cdot n^2$ time. The proof constitutes a hybrid of two different classic techniques in graph algorithms. The first is the *irrelevant vertex technique* that is typically used in the context of Graph Minors and deals with properties such as planarity or surface-embeddability (that are *not* FOL-expressible) and the second is the use of *Gaifman's Locality Theorem* that is the theoretical base for the meta-algorithmic study of FOL-expressible problems.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms; Theory of computation → Parameterized complexity and exact algorithms

Keywords and phrases Graph modification Problems, Algorithmic meta-theorems, First Order Logic, Irrelevant vertex technique, Planar graphs, Surface embeddable graphs

Digital Object Identifier 10.4230/LIPIcs.ESA.2020.51

Funding The research has been supported by the Research Council of Norway and the French Ministry of Europe and Foreign Affairs via the Franco-Norwegian project PHC AURORA 2019.
Fedor V. Fomin: Research Council of Norway via the project "MULTIVAL".
Petr A. Golovach: Research Council of Norway via the project "MULTIVAL".
Dimitrios M. Thilikos: DEMOGRAPH (ANR-16-CE40-0028) and ESIGMA (ANR-17-CE23-0010).

28th Annual European Symposium on Algorithms (ESA 2020).
Editors: Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders; Article No. 51; pp. 51:1–51:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The term *algorithmic meta-theorems* was coined by Grohe in his seminal exposition in [20] in order to describe results providing general conditions, typically of logical and/or combinatorial nature, that automatically guarantee the existence of certain types of algorithms for wide families of problems. Algorithmic meta-theorems reveal deep relations between logic and combinatorial structures, which is a fundamental issue of computational complexity. Such theorems not only yield a better understanding of the scope of general algorithmic techniques and the limits of tractability but often provide (or induce) a variety of new algorithmic results. The archetype of algorithmic meta-theorems is Courcelle's theorem [5,6] stating that all graph properties expressible in Monadic Second-Order Logic (in short, *MSOL-expressible properties*) are fixed-parameter tractable when parameterized by the size of the sentence and the treewidth of the graph.

Our meta-theorem belongs to the intersection of two algorithmic research directions: Deciding First-Order Logic properties of sparse graphs and graph planarization algorithms.

FOL-expressible properties on sparse graphs. For graph properties expressible in first-order logic (in short *FOL-expressible properties*), a rich family of algorithmic meta-theorems, were developed within the last decades. Each of these meta-theorems can be stated in the following form: for a graph class $\mathcal{C}$, deciding FOL-expressible properties is fixed-parameter tractable on $\mathcal{C}$, i.e. there is an algorithm running in $f(|\phi|) \cdot n^{\mathcal{O}(1)}$ time where $|\phi|$ is the size of the the input FOL-sentence ϕ and n is the number of vertices of the input graph. The starting point in the chain of such meta-theorems is the work of Seese [33] for $\mathcal{C}$ being the class of graphs of bounded degree [33]. The first significant extension of Seese's theorem was obtained by Frick and Grohe [16] for the class $\mathcal{C}$ of graphs of bounded local treewidth [16]. The class of graphs of bounded local treewidth contains graphs of bounded degree, planar graphs, graphs of bounded genus, and apex-minor-free graphs. The next step was done by Flum and Grohe [12], who extended these results up to graph classes excluding some minor. Dawar, Grohe, and Kreutzer [9] pushed the tractability border up to graphs locally excluding a minor. Further extension was due to Dvořák, Král, and Thomas, who proved tractability for the class $\mathcal{C}$ of being locally bounded expansion [11]. Finally, Grohe, Kreutzer, and Siebertz [22] established fixed-parameter tractability for classes that are effectively nowhere dense. In some sense, the result of Grohe et al. is the culmination of this long line of meta-theorems, because for somewhere dense graph classes closed under taking subgraphs deciding first-order properties is unlikely to be fixed-parameter tractable [11, 26].

Notice that the above line of results also sheds some light on graph modification problems. In particular, since many modification operations are FOL-expressible, in some situations when the target property $\mathcal{P}$ is FOL-expressible, the above meta-algorithmic results can be extended to graph modification problems. As a concrete example, consider the problem of removing at most k vertices to obtain a graph of degree at most 3. All vertices of the input graph of degree at least $4 + k$ should be deleted, so we delete them and adapt the parameter k accordingly. In the remaining graph all vertices are of degree at most $3 + k$ and the property of removing at most k vertices from such a graph to obtain a graph of degree at most 3 is FOL-expressible. Hence the Seese's theorem implies that there is an algorithm of running time $f(k) \cdot n^{\mathcal{O}(1)}$ solving this problem. However these theories are not applicable with instantiations of $\mathcal{P}$, like planarity, that are not FOL-expressible.

Another island of tractability for graph modification problems is provided by Courcelle's theorem and similar theorems on graphs of bounded widths. For example, graph modification problems are fixed-parameter tractable in cases where the target property $\mathcal{P}$ is MSOL-expressible under the additional assumption that the graphs in $\mathcal{P}$ have fixed treewidth (or bounded rankwidth, for MSOL_1-properties, see e.g., [7]).

To conclude, according to the current state of the art, all known algorithmic meta-theorems concerning fixed-parameter tractability of graph modification problems are attainable either when the target property $\mathcal{P}$ is FOL-expressible and the structure is sparse or when $\mathcal{P}$ is MSOL/MSO_1-expressible and the structure has bounded tree/rank-width. Interestingly, *planarity* is the typical property that escapes the above pattern: it is *not* FOL-expressible and it has *unbounded* treewidth.

Graph planarization. The Planar Vertex Deletion problem is a generalization of planarity testing. For a given graph G the goal is to find a vertex set of size at most k whose removal makes the resulting graph planar. Planarity is a nontrivial and hereditary graph property, hence by the result of Lewis and Yannakakis [27], the decision version of Planar Vertex Deletion is NP-complete. The parameterized complexity of this problem has been extensively studied.

The non-uniform fixed-parameter tractability of Planar Vertex Deletion (parameterized by k) follows from the deep result of Robertson and Seymour in Graph Minors theory [32], that every minor-closed graph class can be recognized in polynomial time. Since the class of graphs that can be made planar by removing at most k vertices is minor-closed, the result of Robertson and Seymour implies that for Planar Vertex Deletion, for each k, there exists a (non-uniform) algorithm that in time $\mathcal{O}(n^3)$ solves Planar Vertex Deletion. Significant amount of work was involved to improve the enormous constants hidden in the big-Oh and the polynomial dependence in n. Marx and Schlotter [29] gave an algorithm that solves the problem in time $f(k)\cdot n^2$, where f is some function of k only. Kawarabayashi [24] obtained the first linear time algorithm of running time $f(k)\cdot n$ and Jansen, Lokshtanov, and Saurabh [23] obtained an algorithm of running time $\mathcal{O}(2^{\mathcal{O}(k\log k)}\cdot n)$. For the related problem of contracting at most k edges to obtain a planar graph, Planar Edge Contraction, an $f(k)\cdot n^{\mathcal{O}(1)}$ time algorithm was obtained by Golovach, van 't Hof and Paulusma [19]. Approximation algorithms for Planar Vertex Deletion and for Planar Edge Deletion were studied in [2–4].

Our results. Let $\boxtimes$ be one of the following operations on graphs: Vertex removal, edge removal, edge contraction, or edge addition. We are interested whether, for a given graph G and an FOL-sentence ϕ, it is possible to transform G by applying at most k $\boxtimes$-operations, into a planar graph with the property defined by ϕ. We refer to this problem as the Graph $\boxtimes$-Modification to Planarity and ϕ problem. For example, when $\boxtimes$ is the vertex removal operation and ϕ is a tautology, then the problem is Planar Vertex Deletion. Similarly, Graph $\boxtimes$-Modification to Planarity and ϕ generalizes Planar Edge Deletion and Planar Edge Contraction. On the other hand, for the special case of $k=0$ this is the problem of deciding FOL-expressible properties on planar graphs.

Examples of first-order expressible properties are deciding whether the input graph G contains a fixed graph H as a subgraph (H-Subgraph Isomorphism), deciding whether there is a homomorphism from a fixed graph H to G to (H-Homomorphism), satisfying degree constraints (the degree of every vertex of the graph should be between a and b for some constants a and b), excluding a subgraph of constant size or having a dominating set of constant size. Thus Graph $\boxtimes$-Modification to Planarity and ϕ encompasses the variety of graph modification problems to planar graphs with specific properties. For example, can we delete k vertices (or edges) such that the obtained graph is planar and each vertex belongs to a triangle? Reversely, can we delete at most k vertices (or edges) from a graph such that the resulting graph is a triangle-free planar graph? Can we add (or contract) at most k edges to such that the resulting graph is 4-regular and planar? Or can we delete at most k edges resulting in a square-free or claw-free planar graph?

Informally, our main result can be stated as follows.

► **Theorem** (Informal). *Graph $\boxtimes$-Modification to Planarity and ϕ is solvable in time $f(k, \phi) \cdot n^2$, for some function f depending on k and ϕ only. Thus the problem is fixed-parameter tractable, when parameterized by $k + |\phi|$.*

Our theorem not only implies that Planar Vertex Deletion is fixed-parameter tractable parameterized by k (proved in [23, 29]) and that deciding whether a planar graph has a first-order logic property ϕ is fixed-parameter tractable parameterized by $|\phi|$ (that follows from [9, 11, 16, 22]). It also implies a variety of new algorithmic results about graph modification problems to planar graphs with some specific properties that cannot be obtained by applying the known results directly. Of course, for some formulas ϕ, Graph $\boxtimes$-Modification to Planarity and ϕ can be solved by more simple techniques. For example, if ϕ defines a hereditary property characterized by a finite family of forbidden induced subgraphs $\mathcal{F}$, then deciding, whether it is possible to delete at most k vertices to obtain a planar $\mathcal{F}$-free graph, can be done by combining the straightforward branching algorithm and, say, the algorithm of Jansen, Lokshtanov, and Saurabh [23] for Planar Vertex Deletion. For this, we iteratively find a copy of each $F \in \mathcal{F}$ and if such a copy exists we branch on all the possibilities to destroy this copy of F by deleting a vertex. By this procedure, we obtain a search tree of depth at most k, whose leaves are all $\mathcal{F}$-free induced subgraphs of the input graph that could be obtained by at most k vertex deletions. Then for each leaf, we use the planarization algorithm limited by the remaining budget. However, this does not work for edge modifications, because deleting an edge in order to ensure planarity may result in creating a copy of a forbidden subgraph. For such type of problems, even for very "simple" ones, like deleting k edges to obtain a claw-free planar graph, or planar graph without induced cycles of length 4, our theorem establishes the first fixed-parameter algorithms. Also our theorem is applicable to the situation when ϕ defines a hereditary property that requires an infinite family of forbidden subgraphs for its characterization and for non-hereditary properties expressible in FOL.

In our paper, we show the result for Graph $\boxtimes$-Modification to Planarity and ϕ, but further we argue that it can be extended for modification problems to graphs embeddable to a surface of a given Euler genus.

The price we pay for such generality is the running time. While the polynomial factor in the running time of our algorithm is comparable with the running time of the algorithm of Marx and Schlotter [29] for Planar Vertex Deletion, it is worse than the more advanced algorithms of Kawarabayashi [24] and Jansen et al. [23]. Similarly, the algorithms for deciding first-order logic properties on graph classes [11, 16, 22] are faster than our algorithm.

The proof of the main theorem is based on a non-trivial combination of the *irrelevant vertex* technique of Robertson and Seymour [30, 31] with the *Gaifman's Locality Theorem* [17]. While both techniques were widely used, see [1, 8, 19, 21, 23, 28] and [9, 12, 16], the combination of the two techniques requires novel ideas. Following the popular trend in Theoretical Computer Science, an alternative title for our paper could be *"Robertson and Seymour meet Gaifman"*.

Organization of the paper. In Section 2 we give the formal definition of the general Graph $\boxtimes$-Modification to Planarity and ϕ problem, present the theoretical background around Gaifman's Locality Theorem, and give some preliminary definitions and results. In Section 3 we highlight the main ideas behind the proof explain how our arguments can be extended in cases where the target property is having bounded Euler genus and being a model of an FOL-sentence ϕ. Finally, in Section 4 we provide some directions for further research.

2 Problem definition and preliminaries

Before we explain our techniques, we give some necessary definitions. We denote by $\mathbb{N}$ the set of all non-negative integers. Given an $n \in \mathbb{N}$, we denote by $\mathbb{N}_{\geq n}$ the set containing all integers equal or greater than n. Given two integers x and y we define by $[x, y] = \{x, x+1, \ldots, y-1, y\}$. Given an $n \in \mathbb{N}_{\geq 1}$, we also define $[n] = [1, n]$.

All graphs in this paper are undirected, finite, and they do not have loops or multiple edges. Given a graph G, we denote by $V(G)$ and $E(G)$ the set of its vertices and edges, respectively. If $S \subseteq V(G)$, then we denote by $G \setminus S$ the graph obtained by G after removing from it all vertices in S, together with their incident edges. Also, we denote by $G \setminus v$ the graph $G \setminus \{v\}$, for some $v \in V(G)$. We also denote by $G[S]$ the graph $G \setminus (V(G) \setminus S)$.

2.1 Modifications on graphs

We define $\mathsf{OP} := \{\mathsf{vr}, \mathsf{er}, \mathsf{ec}, \mathsf{ea}\}$, that is the set of graph operations of removing a vertex, removing an edge, contracting an edge, and adding an edge, respectively. Given an operation $\boxtimes \in \mathsf{OP}$, a graph G, and a vertex set $R \subseteq V(G)$, we define the *application domain* of the operation $\boxtimes$ as

$$\boxtimes\langle G, R\rangle = \begin{cases} R, & \text{if } \boxtimes = \mathsf{vr}, \\ E(G) \cap \binom{R}{2}, & \text{if } \boxtimes = \mathsf{er}, \mathsf{ec}, \text{and} \\ \binom{R}{2} \setminus E(G), & \text{if } \boxtimes = \mathsf{ea}. \end{cases}$$

Notice that $\boxtimes\langle G, R\rangle$ is either a vertex set or a set of subsets of vertices each of size two.

Given a set $S \subseteq \boxtimes\langle G, R\rangle$, we define $G \boxtimes S$ as the graph obtained after applying the operation $\boxtimes$ on the elements of S. The vertices of G that are *affected* by the modification of G to $G \boxtimes S$, denoted by $A(S)$, are the vertices in S, in case $\boxtimes = \mathsf{vr}$ or the endpoints of the edges of S, in case $\boxtimes \in \{\mathsf{er}, \mathsf{ec}, \mathsf{ea}\}$.

Given an FOL-sentence ϕ and some $\boxtimes \in \mathsf{OP}$, we define the following meta-problem:

> Graph $\boxtimes$-Modification to Planarity and ϕ (In short: G$\boxtimes$MPϕ)
> **Input:** A graph G and a non-negative integer k.
> **Question:** Is there a set $S \subseteq \boxtimes\langle G, V(G)\rangle$ of size k such that $G \boxtimes S$ is a planar graph and $G \boxtimes S \models \phi$?

Let $(x_1, \ldots, x_\ell) \in \mathbb{N}^\ell$ and $f, g : \mathbb{N} \to \mathbb{N}$. We use notation $f(n) = \mathcal{O}_{x_1,\ldots,x_\ell}(g(n))$ to denote that there exists a computable function $h : \mathbb{N}^\ell \to \mathbb{N}$ such that $f(n) = h(x_1, \ldots, x_\ell) \cdot g(n)$. We are ready to give the formal statement of the main theorem of this paper.

► **Theorem 1.** *There exists a function $f : \mathbb{N}^2 \to \mathbb{N}$ such that, for every FOL-sentence ϕ and for every $\boxtimes \in \mathsf{OP}$, G$\boxtimes$MP$\phi$ is solvable in $\mathcal{O}_{k,|\phi|}(n^2)$ time.*

2.2 Gaifman's theorem

For vertices u, v of graph G, we use $d_G(u, v)$ to denote the distance between u and v in G. We also use $N_G^{(\leq r)}(v)$ to denote the set of vertices of G at distance at most r from v.

Formulas. In this paper we deal with logic formulas on graphs. In particular we deal with formulas of first-order logic (FOL) and monadic second-order logic (MSO_2). The syntax of FOL-formulas includes the logical connectives $\vee, \wedge, \neg$, a set of variables for vertices, the quantifiers $\forall, \exists$ that are applied to these variables, the predicate $u \sim v$, where u and v are vertex variables and whose interpretation is that u and v are adjacent, and the equality of variables representing vertices. An MSO_2-formula, in addition to the variables for vertices of FOL-formulas, may also contain variables for subsets of vertices or subsets of edges. The syntax of MSO_2-formulas is obtained after enhancing the syntax of FOL-formulas so to further allow quantification on subsets of vertices or subsets of edges and introduce the predicates $v \in S$ (resp. $e \in F$) whose interpretation is that the vertex v belongs in the vertex set S (resp. the edge e belongs in the edge set F).

An FOL-formula ϕ is in *prenex normal form* if it is written as $\phi = Q_1 x_1 \dots Q_n x_n \psi$ such that for every $i \in [n]$, $Q_i \in \{\forall, \exists\}$ and ψ is a quantifier-free formula on the variables $x_1, \dots, x_n$. Then $Q_1 x_1 \dots Q_n x_n$ is referred as the *prefix* of ϕ. For the rest of the paper, when we mention the term "FOL-formula", we mean an FOL-formula on graphs that is in prenex normal form. Given an FOL-formula ϕ, we say that a variable x is a *free variable* in ϕ if it does not occur in the prefix of ϕ. We write $\phi(x_1, \dots, x_r)$ to denote that ϕ is a formula with free variables $x_1, \dots, x_r$. We call a formula without free variables a *sentence*. For a sentence ϕ and a graph G, we write $G \models \phi$ to denote that ϕ evaluates to *true* on G. Also, for a sentence ϕ we denote its length by $|\phi|$.

Gaifman sentences. Given an FOL-formula $\psi(x)$ with one free variable x, we say that $\psi(x)$ is *r-local* if the validity of $\psi(x)$ depends only on the r-neighborhood of x, that is for every graph G and $v \in V(G)$ we have

$$G \models \psi(v) \iff N_G^{(\leq r)}(v) \models \psi(v).$$

Observe that there exists an FOL-formula $\delta_r(x, y)$ such that for every graph G and $v, u \in V(G)$, we have $d_G(u, v) \leq r \iff G \models \delta_r(v, u)$ (see [13, Lemma 12.26]).

We say that an FOL-sentence ϕ is a *Gaifman sentence* when it is a Boolean combination of sentences $\phi_1, \dots, \phi_m$ such that, for every $h \in [m]$,

$$\phi_h = \exists x_1 \dots \exists x_{\ell_h} \big(\bigwedge_{1 \leq i < j \leq \ell_h} d(x_i, x_j) > 2r_h \wedge \bigwedge_{i \in [\ell_h]} \psi_h(x_i) \big), \tag{1}$$

where $\ell_h, r_h \geq 1$ and ψ_h is an r_h-local formula with one free variable. We refer to the variables $x_1, \dots, x_{\ell_h}$ for each $h \in [m]$ as the *basic variables* of ϕ. Moreover, for every $h \in [m]$ we call ϕ_h a *basic local sentence* of ϕ and the formula ψ_h a *local formula* of ϕ.

▶ **Proposition 2** (Gaifman's Theorem [17]). *Every FOL-sentence ϕ is equivalent to a Gaifman sentence ϕ'. Furthermore, ϕ' can be computed effectively.*

2.3 Equivalent formulations

Given a Gaifman sentence ϕ combined from sentences $\phi_1, \ldots, \phi_m$ and a unary relation symbol R, we define $\phi\|_R$ as the sentence that is the same Boolean combination of sentences $\phi_1\|_R, \ldots, \phi_m\|_R$ such that, for every $h \in [m]$,

$$\phi_h\|_R = \exists x_1 \ldots \exists x_{\ell_h} \Big(\bigwedge_{i\in[\ell_h]} x_i \in R \wedge \bigwedge_{1\leq i<j\leq \ell_h} d(x_i, x_j) > 2r_h \wedge \bigwedge_{i\in[\ell_h]} \psi_h(x_i)\Big), \tag{2}$$

where $\ell_h, r_h \geq 1$ and ψ_h is an r_h-local formula with one free variable.

Let (G, k) be an instance of the G⊠MPϕ problem. We may assume, because of Proposition 2, that ϕ is a Gaifman sentence. We consider an enhanced version of the G⊠MPϕ problem as follows. Let (G, R, k) be a triple, where G is a graph, $R \subseteq V(G)$, and $k \in \mathbb{N}$. We say that (G, R, k) is a *$(\phi, \boxtimes)$-triple* if there exists a set $S \subseteq \boxtimes\langle G, R\rangle$ such that $|S| \leq k$, $G \boxtimes S$ is a planar graph, and $G \boxtimes S \models \phi\|_R$. Also, we say that a set $S \subseteq \boxtimes\langle G, V(G)\rangle$ is a *⊠-planarizer* of G if $G \boxtimes S$ is planar. It is easy to observe that the property that (G, R, k) is a $(\phi, \boxtimes)$-triple can be expressed in MSO_2. This is easy in case $\boxtimes \in \{\mathsf{vr}, \mathsf{er}, \mathsf{ec}\}$. In the case where $\boxtimes = \mathsf{ea}$, we observe the following:

▶ **Observation 3.** *Let $\boxtimes = \mathsf{ea}$, G be a graph, and $S \subseteq \langle G, V(G)\rangle$ where $S = \{\{v_1, u_1\}, \ldots, \{v_r, u_r\}\}$. Then there exists an MSO_2-formula $\phi_{\mathcal{P},S}$ on structures of the type $(G, v_1, u_1, \ldots, v_r, u_r)$ such that*

$$G \boxtimes S \text{ is a planar graph} \iff (G, v_1, u_1, \ldots, v_r, u_r) \models \phi_{\mathcal{P},S}.$$

Treewidth. A *tree decomposition* of a graph G is a pair (T, χ) where T is a tree and $\chi : V(T) \to 2^{V(G)}$ such that

1. $\bigcup_{t\in V(T)} \chi(t) = V(G)$;
2. for every edge e of G there is a $t \in V(T)$ such that $\chi(t)$ contains both endpoints of e and
3. for every $v \in V(G)$, the subgraph of T induced by $\{t \in V(T) \mid v \in \chi(t)\}$ is connected.

The *width* of (T, χ) is defined as $\mathbf{w}(T, \chi) := \max\big\{ |\chi(t)| - 1 \bigm| t \in V(T)\big\}$. The *treewidth of* G is defined as

$$\mathbf{tw}(G) := \min\big\{ \mathbf{w}(T, \chi) \bigm| (T, \chi) \text{ is a tree decomposition of } G\big\}.$$

Theorem 1 is a consequence of the following lemma.

▶ **Lemma 4.** *Given a Gaifman sentence ϕ and a $\boxtimes \in \mathsf{OP}$, there exists a function $f_1 : \mathbb{N}^2 \to \mathbb{N}$, and an algorithm with the following specifications:*

Reduce_Instance(k, G, S, R)

Input: an integer $k \in \mathbb{N}$, a graph G, a set $R \subseteq V(G)$, and a set $S \subseteq R$ that is a vr*-planarizer of G of size at most k.*

Output: One of the following:

1. - *if $\boxtimes \in \{\mathsf{er}, \mathsf{ec}, \mathsf{ea}\}$: a report that (G, k) is a* no*-instance of* G⊠MPϕ.
   - *if $\boxtimes = \mathsf{vr}$: a vertex $u \in S$ such that $S \setminus \{u\}$ is a* vr*-planarizer of $G \setminus u$ of size at most $k - 1$ and (G, k) and $(G \setminus u, k - 1)$ are equivalent instances of* G⊠MPϕ.
2. *a vertex set $X \subseteq V(G)$ and a vertex $v \in X$ such that $S \subseteq R \setminus X$ and (G, R, k) is a $(\phi, \boxtimes)$-triple iff $(G \setminus v, R \setminus X, k)$ is a $(\phi, \boxtimes)$-triple.*
3. *a tree decomposition of G of width at most $f_1(k, |\phi|)$.*

Moreover, this algorithm runs in $\mathcal{O}_{k,|\phi|}(n)$ steps.

Given Lemma 4, Theorem 1 can be proved as follows.

Proof of Theorem 1. Let ϕ be an FOL-formula. By Proposition 2, ϕ is equivalent to a Gaifman sentence ϕ'. Using the planarization algorithm from [23], we compute, in $\mathcal{O}_k(n)$ steps, a vr-planarizer S of G of size at most k. If $\boxtimes = \mathsf{ea}$, then $S := \emptyset$, while if $\boxtimes \in \{\mathsf{vr}, \mathsf{er}, \mathsf{ec}\}$, then if such a set does not exist the we safely return a negative answer (for the case of $\boxtimes = \mathsf{er}, \mathsf{ec}$, this is due to the fact that if there exists a ec- or an er-planarizer of G of size at most k then also a vr-planarizer of G of size at most k exists (see [19, Lemma 1])). We are now in position to apply recursively the algorithm **Reduce_Instance**(k, G, S, R) of Lemma 4 until either an answer or the third case appears. In the first case, we either return a negative answer, if $\boxtimes \in \{\mathsf{er}, \mathsf{ec}, \mathsf{ea}\}$, or set $(k, G, S, R) := (k-1, G \setminus v, S \setminus v, R)$ if $\boxtimes = \mathsf{vr}$, while in the second case we set $(k, G, S, R) := (k, G \setminus v, S, R \setminus X)$. In the third case we have that $\mathbf{tw}(G) \leq f_1(k, |\phi'|)$. Recall that the property that (G, R, k) is a $(\phi, \boxtimes)$-triple can be expressed in MSO_2, thus the status of the final equivalent instance (G, R, k) can be evaluated in $\mathcal{O}_{k,|\phi|}(n)$ steps by applying Courcelle's theorem. As the recursion takes at most n steps, we obtain the claimed running time. ◀

3 The algorithm

3.1 Two main lemmata

We now give two lemmata, whose combination gives the proof of Lemma 4. Before we state them, we give a series of definitions.

Let $\boxtimes \in \mathsf{OP}$, G be a graph, $k \in \mathbb{N}$, and let S be a $\boxtimes$-planarizer of G. We say that S is an *inclusion-minimal $\boxtimes$-planarizer* of G if none of its proper subsets is a $\boxtimes$-planarizer of G. Notice that, in the special case where $\boxtimes = \mathsf{ea}$, the unique inclusion-minimal $\boxtimes$-planarizer of G is the empty set of edges. We say that a set $Q \subseteq V(G)$ is *$\boxtimes$-planarization irrelevant* if for every inclusion-minimal $\boxtimes$-planarizer S of G that has size at most k, it holds that $A(S) \cap Q = \emptyset$.

Partially disk-embedded graphs. We define a *closed disk* Δ to be a subset of the plane homeomorphic to the set $\{(x, y) \mid x^2 + y^2 \leq 1\}$ and we use $\mathbf{bor}(\Delta)$ to denote its boundary. We say that a graph G is *partially disk-embedded in some closed disk* Δ, if there is some subgraph K of G that is embedded in Δ such that $\mathbf{bor}(\Delta)$ is a cycle of K and no vertex in $\Delta \setminus \mathbf{bor}(\Delta)$ is adjacent to a vertex not in Δ. We use the term *partially Δ-embedded graph G* to denote that a graph G is partially disk-embedded in some closed disk Δ. We also call the graph K *compass* of the partially Δ-embedded graph G and we always assume that we accompany a partially Δ-embedded graph G together with an embedding of its compass in Δ that is the set $G \cap \Delta$.

Grids and walls. Let $k, r \in \mathbb{N}$. The *$(k \times r)$-grid* is the Cartesian product of two paths on k and r vertices respectively. An *elementary r-wall*, for some odd $r \geq 3$, is the graph obtained from a $(2r \times r)$-grid with vertices (x, y), $x \in [2r] \times [r]$, after the removal of the "vertical" edges $\{(x, y), (x, y+1)\}$ for odd $x + y$, and then the removal of all vertices of degree one. Notice that, as $r \geq 3$, an elementary r-wall is a planar graph that has a unique (up to topological isomorphism) embedding in the plane such that all its finite faces are incident to exactly six edges. The *perimeter* of an elementary r-wall is the cycle bounding its infinite face, while the cycles bounding its finite faces are called *bricks*. Given an elementary wall $\overline{W}$, a *vertical path* of $\overline{W}$ is one whose vertices, in ordering of appearance, are $(i, 1), (i, 2), (i+1, 2), (i+1, 3), (i, 3), (i, 4), (i+1, 4), (i+1, 5), (i, 5), \ldots, (i, r-2), (i, r-1), (i+1, r-1), (i+1, r)$, for some $i \in \{1, 3, \ldots, 2r-1\}$. Also an *horizontal path* of $\overline{W}$ is the one whose vertices, in ordering of appearance, are $(1, j), (2, j), \ldots, (2r, j)$, for some $j \in [2, r-1]$, or $(1, 1), (2, 1), \ldots, (2r-1, 1)$ or $(2, r), (2, r), \ldots, (2r, r)$.

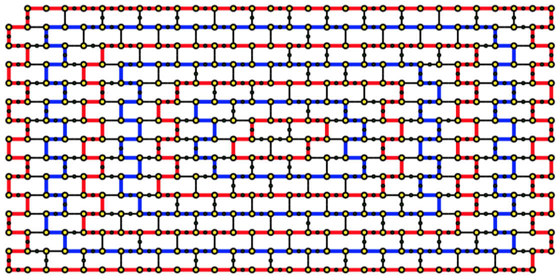

Figure 1 An 15-wall and its 7 layers.

An *r-wall* is any graph W obtained from an elementary r-wall $\overline{W}$ after subdividing edges (see Figure 1). We call the vertices that where added after the subdivision operations *subdivision vertices*, while we call the rest of the vertices (i.e., those of $\overline{W}$) *branch vertices*. The *perimeter* of W, denoted by $\mathsf{perim}(W)$, is the cycle of W whose non-subdivision vertices are the vertices of the perimeter of $\overline{W}$. Also, a vertical (resp. horizontal) path of W is a subdivided vertical (resp. horizontal) path of $\overline{W}$.

A subgraph W of a graph G is called a *wall* of G if W is an r-wall for some odd $r \geq 3$ and we refer to r as the *height* of the wall W.

Let W be a wall of a graph G and K' be the connected component of $G \setminus \mathsf{perim}(W)$ that contains $W \setminus \mathsf{perim}(W)$. The *compass* of W, denoted by $\mathsf{comp}(W)$, is the graph $G[V(K') \cup V(\mathsf{perim}(W))]$. Observe that W is a subgraph of $\mathsf{comp}(W)$ and $\mathsf{comp}(W)$ is connected.

The *layers* of an r-wall W are recursively defined as follows. The first layer of W is its perimeter. For $i = 2, \ldots, (r-1)/2$, the i-th layer of W is the $(i-1)$-th layer of the subwall W' obtained from W after removing from W its perimeter and all occurring vertices of degree one. Notice that each $(2r+1)$-wall has r layers (see Figure 1). The *central vertices* of W, denoted by $\mathsf{center}(W)$, are the two branch vertices of W that do not belong to any of its layers.

We are now in position to state the following two lemmata.

▶ **Lemma 5.** *Given a $\boxtimes \in \mathsf{OP}$, there exist two functions $f_1, f_2 : \mathbb{N}^2 \to \mathbb{N}$, and an algorithm with the following specifications:*

Find_Area(k, q, G, S)

Input: a $k \in \mathbb{N}$, an odd $q \in \mathbb{N}_{\geq 1}$, a graph G, and a set $S \subseteq V(G)$ that is a vr-planarizer of G of size at most k.

Output: One of the following:

1. - *if $\boxtimes \in \{\mathsf{er}, \mathsf{ec}, \mathsf{ea}\}$: a report that (G, k) is a no-instance of* G$\boxtimes$MPϕ.
   - *if $\boxtimes = \mathsf{vr}$: a vertex $u \in S$ such that $S \setminus u$ is a vr-planarizer of $G \setminus u$ of size at most $k-1$ and (G, k) and $(G \setminus u, k-1)$ are equivalent instances of* G$\boxtimes$MPϕ.
2. *a q-wall W of G and a closed disk Δ such that*
   - *the compass of W has treewidth at most $f_2(k, q)$,*
   - *G is partially Δ-embedded, where $G \cap \Delta = \mathsf{comp}(W)$, $\mathbf{bor}(\Delta) = \mathsf{perim}(W)$,*

- $V(\textsf{comp}(W))$ *is* $\boxtimes$*-planarization irrelevant, and*
- $N_G(S) \cap V(\textsf{comp}(W)) = \emptyset$*, or*

3. *a tree decomposition of* G *of width at most* $f_1(k, q)$*.*

Moreover, this algorithm runs in $\mathcal{O}_{k,q}(n)$ *steps.*

By $N_G(S)$ we denote the vertices not in S adjacent in G with vertices in S. In the first possible output of the algorithm of Lemma 5 we have either a negative answer to the G$\boxtimes$MPϕ problem or an equivalent instance of G$\boxtimes$MPϕ with reduced value of k.

The main steps of the proof of Lemma 5 are the following. In case, $\boxtimes = \textsf{ea}$ we first check whether G is planar. If not, we report a negative answer, otherwise we find a wall W in G whose size is a "big-enough" function of k and whose compass has "small-enough" treewidth using [18, Lemma 4.2]. This wall contains an (also "big-enough") subwall of W whose compass is not affected by S. In case $\boxtimes \in \{\textsf{vr}, \textsf{er}, \textsf{ec}\}$, we consider the neighbors of S in the planar graph G', that is the set $N := N_G(S)$. Moreover, we consider a big enough triangulated grid Γ as a contraction of G' (using [14, Theorem 3]) and the set N_Γ of the "contraction-heirs" of the vertices of N in Γ. If $|N_\Gamma|$ is "big-enough", then we prove, using the main technical result of [10], that some of the vertices of S should be affected by every possible solution, in case $\boxtimes = \textsf{vr}$, or that we have a $\textsf{no}$-instance, in case $\boxtimes \in \{\textsf{er}, \textsf{ec}\}$. If $|N_\Gamma|$ is "small-enough", then we can find a "big-enough" wall W in G whose compass is not affected by S (again using the previously mentioned result of [18]). The proof is completed by proving that this wall contains some "big-enough" subwall that is not affected by any inclusion-minimal $\boxtimes$-planarizer.

The next lemma deals with the second possible output of the algorithm of Lemma 5 and contains the "core arguments" of this paper.

▶ **Lemma 6.** *Given a Gaifman sentence* ϕ *and a* $\boxtimes \in \textsf{OP}$*, there exists a function* $f_3 : \mathbb{N}^2 \to \mathbb{N}$ *and an algorithm with the following specifications:*

Find_Vertex$(k, \Delta, G, R, \tilde{W})$

Input: a $k \in \mathbb{N}$*, a partially* Δ*-embedded graph* G*, a set of annotated vertices* $R \subseteq V(G)$*, and a* q*-wall* $\tilde{W}$ *of* G *such that*

- $q = f_3(k, |\phi|)$*.*
- *the compass of* $\tilde{W}$ *has treewidth at most* $f_2(k, q)$*,*
- $G \cap \Delta = \textsf{comp}(\tilde{W})$*,* $\textbf{bor}(\Delta) = \textsf{perim}(\tilde{W})$*,*
- $V(\textsf{comp}(\tilde{W}))$ *is* $\boxtimes$*-planarization irrelevant, and*

Output: a vertex set $X \subsetneq V(\textsf{comp}(\tilde{W}))$ *and a vertex* $v \in X$ *such that* (G, R, k) *is a* $(\phi, \boxtimes)$*-triple iff* $(G \setminus v, R \setminus X, k)$ *is a* $(\phi, \boxtimes)$*-triple.*

Moreover, this algorithm runs in $\mathcal{O}_{k,|\phi|}(n)$ *steps.*

Notice that the above algorithm produces a $(\phi, \boxtimes)$-triple where both R and G are reduced. To see why Lemma 4 follows from Lemma 5 and Lemma 6, observe that in the second possible output of the algorithm **Find_Area**(k, q, G, S) we can call the algorithm **Find_Vertex**$(k, \Delta, G, R, \tilde{W})$, where $\tilde{W} := W$, which outputs a vertex set $X \subsetneq V(\textsf{comp}(\tilde{W}))$ and a vertex $v \in X$ such that (G, R, k) is a $(\phi, \boxtimes)$-triple iff $(G \setminus v, R \setminus X, k)$ is a $(\phi, \boxtimes)$-triple. Observe that since $N_G(S) \cap V(\textsf{comp}(\tilde{W})) = \emptyset$, then $S \subseteq R \setminus X$. We insist that the algorithm **Find_Vertex**$(k, \Delta, G, R, \tilde{W})$ does not use the fact that $N_G(S) \cap V(\textsf{comp}(\tilde{W})) = \emptyset$ but we use the latter to guarantee that $S \subseteq R \setminus X$. For the running time of Lemma 4, recall that the two algorithms of Lemma 5 and Lemma 6 run in $\mathcal{O}_{k,|\phi|}(n)$ steps.

3.2 Sketch of the proof of Lemma 6

In order to prove Lemma 6, we first find a collection $\mathcal{W}$ of "sufficiently many" subwalls of $\tilde{W}$ each with ρ layers (where ρ is "big-enough"), whose compasses are pairwise vertex-disjoint.

The key idea is to define a "characteristic" of each wall $W \in \mathcal{W}$ that encodes all possible ways that a $\boxtimes$-planarizer S of G affects $\mathsf{comp}(W)$ along with the ways that the fact that $G \boxtimes S \models \phi$ is certified by a vertex assignment to the basic variables of the Gaifman formula ϕ in $\mathsf{comp}(W)$. Recall that $\phi\|_R$ is a Boolean combination of sentences $\phi_1\|_R, \ldots, \phi_m\|_R$ so that for every $h \in [m]$,

$$\phi_h\|_R = \exists x_1 \ldots \exists x_{\ell_h} \big(\bigwedge_{i\in[\ell_h]} x_i \in R \wedge \bigwedge_{1\leq i<j\leq \ell_h} d(x_i,x_j) > 2r_h \wedge \bigwedge_{i\in[\ell_h]} \psi_h(x_i)\big),$$

where $\ell_h, r_h \geq 1$ and ψ_h is an r_h-local formula with one free variable. Notice that $\phi\|_R$ is evaluated on annotated graphs of the form (G,R). Clearly, $\phi\|_R$ is a sentence in Monadic Second Order Logic, in short, a MSO_2-sentence.

As a first step, for every $h \in [m]$, $W \in \mathcal{W}$, $S \subseteq \boxtimes\langle G,R\rangle$ of size at most k, $I_h \subseteq [\ell_h]$, and $t \in [\rho]$, we define:

$$\mathsf{sig}_{\phi_h,\boxtimes}^{(S,I_h,t)}(W) := \begin{cases} 1, & \text{if } \exists \tilde{X} = \{x_i \mid i \in I_h\} \subseteq V(\mathsf{comp}(W^{(t)}) \boxtimes S) \cap R \text{ such that } \tilde{X} \\ & \text{is } (|I_h|, r_h)\text{-scattered in } \mathsf{comp}(W^{(t)}) \boxtimes S \text{ and } G \boxtimes S \models \bigwedge_{x\in\tilde{X}} \psi_h(x), \\ 0, & \text{otherwise.} \end{cases}$$

In the above definition, $W^{(t)}$ is the subwall of W that has t layers (which are the last t layers of W) and the same center as W. Also, a set X of vertices is (α,β)-*scattered*, if $|X| = \alpha$ and there are no two vertices in X within distance $\leq 2\beta$. Intuitively, $\mathsf{sig}_{\phi_h,\boxtimes}^{(S,I_h,t)}(W) = 1$ if the application of the operation $\boxtimes$ on G as defined by S gives rise to the existence of a scattered set $\tilde{X}$ in the compass of $W^{(t)}$ so that when the vertices of $\tilde{X}$ are assigned to the basic variables of ϕ_h corresponding to I_h, the local formula ψ_h is satisfied for each $x_i \in \tilde{X}$ in the modified graph.

Next, for every $W \in \mathcal{W}$ and every $S \subseteq \boxtimes\langle G,R\rangle$ of size at most k we define:

$$\mathsf{msig}_{\phi,\boxtimes}^{(S)}(W) = \Big((\mathsf{sig}_{\phi_1,\boxtimes}^{(S,I_1,t)}(W), \ldots, \mathsf{sig}_{\phi_m,\boxtimes}^{(S,I_m,t)}(W)) \mid (I_1,\ldots,I_m,t) \in 2^{[\ell_1]} \times \cdots \times 2^{[\ell_m]} \times [\rho] \Big).$$

Clearly, $\mathsf{msig}_{\phi,\boxtimes}^{(S)}(W)$ can be seen as a $(2^\ell \cdot \rho)$-tuple of binary m-tuples, given that $\ell := \sum_{h\in[m]} \ell_h$. Let SIG be the set of all such tuples and notice that $|\mathsf{SIG}|$ is bounded by some function of k and $|\phi|$ and $\Big\{\mathsf{msig}_{\phi,\boxtimes}^{(S)}(W) \;\Big|\; W \in \mathcal{W}, S \subseteq \boxtimes\langle G,R\rangle \text{ of size at most } k\Big\} \subseteq \mathsf{SIG}$.

It is now time to define the characteristic of a wall $W \in \mathcal{W}$. We set $r := \max_{h\in[m]}\{r_h\}$ and $d := 2(r + (\ell+1)r + r)$. We define the $(\phi,\boxtimes)$-*characteristic* of W as follows:

$$\begin{aligned}(\phi,\boxtimes)\text{-}\mathsf{char}(W) \quad = \quad \{(s,\sigma,t) \in [0,k] \times \mathsf{SIG} \times [d+1,\rho] \mid \exists S \subseteq \boxtimes\langle G,R\rangle, \\ |S| = s, \\ A(S) \subseteq V(\mathsf{comp}(W^{(t-d)})) \cap R, \\ \mathsf{comp}(W) \boxtimes S \text{ is planar, and} \\ \mathsf{msig}_{\phi,\boxtimes}^{(S)}(W) = \sigma\}.\end{aligned}$$

Notice that all queries in the definition of $(\phi,\boxtimes)$-$\mathsf{char}(W)$ can be expressed in MSO_2. Indeed, this is easy to see when $\boxtimes \in \{\mathsf{vr},\mathsf{er},\mathsf{ec}\}$, as in this case the query "$\mathsf{comp}(W) \boxtimes S$ is planar" is trivially true, since $V(\mathsf{comp}(\tilde{W}))$ is $\boxtimes$-planarization irrelevant. In the case where $\boxtimes = \mathsf{ea}$, the MSO_2 expressibility follows from Theorem 3. As each $W \in \mathcal{W}$ has treewidth bounded by a function of k and $|\phi|$, it follows by the theorem of Courcelle that $(\phi,\boxtimes)$-$\mathsf{char}(W)$ can be computed in $\mathcal{O}_{k,|\phi|}(n)$ time.

We say that two walls are $(\phi, \boxtimes)$-*equivalent* if they have the same $(\phi, \boxtimes)$-characteristic. Since the collection $\mathcal{W}$ contains "sufficiently many" walls, then we can find a collection $\mathcal{W}' \subseteq \mathcal{W}$ of also "sufficiently many" walls that are pairwise equivalent. We fix a wall $W_1 \in \mathcal{W}'$ and we set $X := \mathsf{comp}(W_1^{(r)})$, where $r = \max_{h\in[m]}\{r_h\}$, and $v \in \mathsf{center}(W_1)$.

In what follows, we highlight the ideas of the proof of the fact that if (G, R, k) is a $(\phi, \boxtimes)$-triple, then $(G \setminus v, R \setminus X, k)$ is a $(\phi, \boxtimes)$-triple. We first consider a set $S \subseteq \boxtimes\langle G, R\rangle$ of size at most k that certifies that (G, R, k) is a $(\phi, \boxtimes)$-triple. Then, we pick a wall $W_2 \in \mathcal{W}' \setminus \{W_1\}$ whose compass is not affected by S. We are allowed to pick this wall since there are "sufficiently many" walls equivalent to W_1 in $\mathcal{W}'$. Our strategy is to use the fact that W_1 and W_2 are $(\phi, \boxtimes)$-equivalent in order to state a "replacement argument": we can find a $t \in [\rho]$, such that the subset S_{in} of S that affects $\mathsf{comp}(W_1^{(t)})$ and the set X of vertices of $\mathsf{comp}(W_1^{(t)})$ that are assigned to the basic variables of ϕ in order to certify that $G \boxtimes S \models \phi$, can be replaced by their "equivalent" sets $\tilde{S}$ and $\tilde{X}$ in $\mathsf{comp}(W_2^{(t)})$. As a consequence of this, for every possible solution S and vertex assignment to the basic variables of ϕ, we can find both a new solution and a new vertex assignment that "avoid" the "inner part" of W_1. This implies that the validity of any local formula of ϕ does not depend on the central vertices of W_1. Thus, we can declare one of them "irrelevant" and safely remove it from G, while storing (by reducing R to $R \setminus X$) the fact that every possible solution S and vertex assignment to the basic variables of ϕ can "avoid" the "inner part" of W_1.

To further inspect how this "replacement" is achieved, we need to dive deeper into the technicalities of the proof (through an intuitive perspective). Given a wall W, we refer to a *wall-annulus* of W as the subgraph of W that is obtained from W after removing from W all its layers, except a fixed number of consecutive layers. We think of every wall $W \in \mathcal{W}$ as divided in consecutive wall-annuli of fixed size. Since ρ is "big-enough", then we can find also "many enough" such wall-annuli. We denote each one of them by $A_i(W)$. Given a $W \in \mathcal{W}$, every wall-annulus $A_i(W)$ is divided in some regions as depicted in Figure 2.

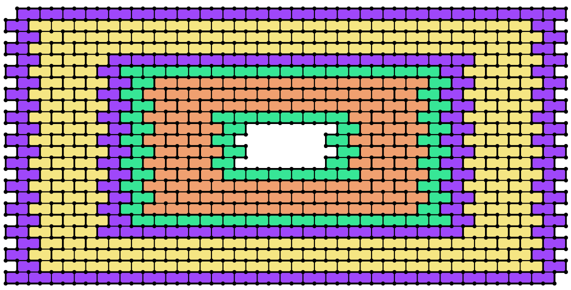

Figure 2 An example of a wall-annulus $A_i(W)$ of a wall $W \in \mathcal{W}$, together with its regions refered in the proof of Lemma 6.

The regions depicted in purple and green are consisting of r layers of the wall W (recall that $r = \max_{h\in[m]}\{r_h\}$). The regions depicted in yellow and orange are "big-enough" so as to be able to find an also "big-enough" wall-annulus that "avoids" a given vertex assignment to the basic variables of ϕ.

Since ρ is "big-enough", then we can find a wall-annulus $A_i(W_1)$ that is not affected by S. This allows us to partition S in two sets, S_{in} and S_{out} in the obvious way. The fact that W_1 and W_2 are $(\phi, \boxtimes)$-equivalent implies the existence of a set $\tilde{S}$ in W_2 certifying that $(\phi, \boxtimes)\text{-}\mathsf{char}(W_2) = (\phi, \boxtimes)\text{-}\mathsf{char}(W_1)$. Thus, by setting $S' := \tilde{S} \cup S_{\mathsf{out}}$, we have that $S' \subseteq \boxtimes\langle G, R'\rangle$, $|S'| = |S|$, and $G \boxtimes S'$ is planar. The latter is guaranteed by the fact that $V(\mathsf{comp}(\tilde{W}))$ is $\boxtimes$-planarization irrelevant, in the case $\boxtimes \in \{\mathsf{vr}\ , \mathsf{er}, \mathsf{ec}\}$, while in the case that

$\boxtimes = \mathsf{ea}$, the existence of the outer purple buffer of $A_i(W_1)$ (resp. $A_i(W_2)$) allows us to treat S_{in} (resp. $\tilde{S}$) and S_{out} separately, while not spoiling planarity. The last part of the proof requires to prove that $G \boxtimes S \models \phi\|_R \iff G \boxtimes S' \models \phi\|_{R'}$.

For simplicity, here we only argue why $G \boxtimes S \models \phi_h\|_R \implies G \boxtimes S' \models \phi_h\|_{R'}$ holds, as the arguments in the proof of the inverse direction are completely symmetrical. Therefore, given an (ℓ_h, r_h)-scattered set X such that ϕ_h is satisfied if the vertices of X are assigned to the basic variables of ϕ_h, we aim to find a $t \in [\rho]$ in order to "replace" the vertices in $X \cap V(\mathsf{comp}(W_1^{(t)}))$ with a set $\tilde{X}$ of vertices in $\mathsf{comp}(W_2^{(t)})$ such that the resulting vertex set X' is (ℓ_h, r_h)-scattered and ϕ_h is satisfied if the vertices of X' are assigned to the basic variables of ϕ_h. Notice that for every $h \in [m]$ such that $G \boxtimes S \models \phi_h\|_R$, these "replacement arguments" are pairwise independent.

We first deal with the possibility that the given scattered set X intersects some "inner part" of $\mathsf{comp}(W_2)$. Thus, in order to "clean" the "inner part" of $\mathsf{comp}(W_2)$, we find a wall $W_3 \in \mathcal{W}' \setminus \{W_1, W_2\}$ that "avoids" both S and X (for different $h \in [m]$, the choice of W_3 may coincide). Also, we consider a $\tilde{t} \in [\rho]$ corresponding to a layer in the yellow region of the wall-annulus $A_i(W_2)$ such that the annulus of the wall-annulus of $A_i(W_2)$ bounded by the $(t - r + 1)$-th and t-th layer of W_2 is not intersected by X. Then, we "replace" the vertices of X in $\mathsf{comp}(W_2^{(\tilde{t})})$, call it X_{in} with an "equivalent" vertex set $\tilde{X}$ in $\mathsf{comp}(W_3^{(\tilde{t})})$ (notice that this is achieved by arguing for $S := \emptyset$ in the notion of $(\phi, \boxtimes)$-characteristic). This results to an (ℓ_h, r_h)-scattered set Y such that Y does not intersect $\mathsf{comp}(W_2^{(\tilde{t})})$ and $G \boxtimes S \models \bigwedge_{x \in Y} \psi_h(x)$ (see Figure 3).

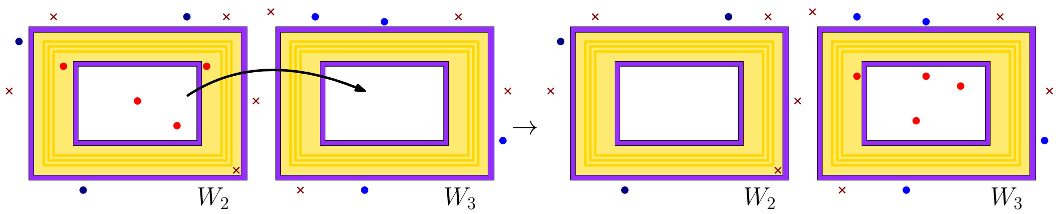

Figure 3 The "cleaning" of the "inner part" of $\mathsf{comp}(W_2)$. Left: The set $A(S)$ is depicted in cross vertices, the set $X \setminus X_{\mathsf{in}}$ is depicted in blue, and the set X_{in} is depicted in red. Right: The set $A(S)$ is depicted in cross vertices, the set $Y \setminus X_{\mathsf{in}}$ is depicted in blue, and the set $\tilde{X}$ is depicted in red.

Now, we are allowed to pick a $t \in [\rho]$ corresponding to an "orange" layer of $A_i(W_1)$ such that the annulus of the wall-annulus of $A_i(W_1)$ bounded by the $(t' - r)$-th and t'-th layer of W_1 is not intersected by X. If we set Y_{in} to be the set of vertices of Y in $\mathsf{comp}(W_1^{(t')})$, then since $\mathsf{msig}_{\phi,\boxtimes}^{(S_{\mathsf{in}})}(W_1) = \mathsf{msig}_{\phi,\boxtimes}^{(\tilde{S})}(W_2)$, then there exists a set $\tilde{Y}$ in $\mathsf{comp}(W_2^{(t')})$ that is "equivalent" to Y_{in} (see Figure 4).

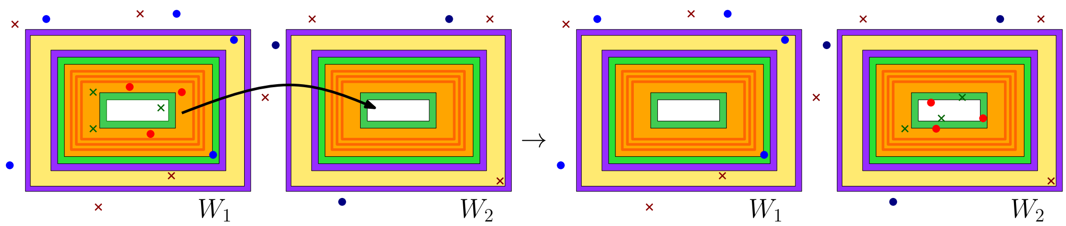

Figure 4 The last part of the proof. Left: The set $A(S_{\mathsf{out}})$ is depicted in red cross vertices, the set $A(S_{\mathsf{in}})$ is depicted in green cross vertices, the set $Y \setminus Y_{\mathsf{in}}$ is depicted in blue, and the set Y_{in} is depicted in red. Right: The set $A(S_{\mathsf{out}})$ is depicted in red cross vertices, the set $A(\tilde{S})$ is depicted in green cross vertices, the set $Y \setminus Y_{\mathsf{in}}$ is depicted in blue, and the set $\tilde{Y}$ is depicted in red.

Therefore, since $\tilde{Y}$ is in the orange region of $\mathsf{comp}(W_2)$ and Y is "avoiding" $\mathsf{comp}(W_2^{(\tilde{t})})$, then we can derive that Y and $\tilde{Y}$ are "separated" by a green and a purple region of $A_i(W_2)$. Thus, $X' := (Y \setminus Y_{\mathsf{in}}) \cup \tilde{Y}$ is an (ℓ_h, r_h)-scattered set of $G \boxtimes S'$ that "avoids" $\mathsf{comp}(W_1^{(r)})$. Moreover, ϕ_h is satisfied given that the vertices of X' of $G \boxtimes S'$ are assigned to the basic variables of ϕ_h. The proof is concluded.

3.3 Extension on graphs of bounded genus

The immediate question is whether our results can be extended to target properties that are more general than planarity (and still not FOL-expressible). The first candidate is the $\boxtimes$-Modification to g-Euler Genus and ϕ, where we ask for a set $S \subseteq \boxtimes\langle G, V(G)\rangle$ of size k such that $G \boxtimes S$ has Euler genus at most g. Notice that the property of having Euler genus at most g is not FOL-expressible. On the positive side, this property is MSO_2-expressible as there is a set $\mathcal{B}_g$ of graphs such that G has Euler genus at most g iff none of the graphs in $\mathcal{B}_g$ is a minor of G and minor containment is MSO_2-expressible. We next argue about how to adapt the techniques of this paper in order to prove that this problem can be solved in $\mathcal{O}_{k,|\phi|,g}(n^2)$ when $\boxtimes \in \{\mathsf{vr}, \mathsf{er}, \mathsf{ec}\}$. For this we first straightforwardly extend the notions of $\boxtimes$-planarization irrelevant vertex set and $\boxtimes$-planarizer to the respective notions of $\boxtimes$-*g-Euler Genus irrelevant* vertex set $\boxtimes$-*g-euler genus enforcer*. Our aim is to prove a more general version of Lemma 4 where $\boxtimes$-planarizer is replaced by $\boxtimes$-g-euler genus enforcer. The $\mathcal{O}_{k,|\phi|,g}(n^2)$ time algorithm for $\boxtimes$-Modification to g-Euler Genus and ϕ follows directly from this extended version of Lemma 4 with the same arguments as its planarization counterpart. The extended version of Lemma 4 in turn is a consequence of the generalized versions of Lemma 5 and Lemma 6 where $\boxtimes$-planarizer is replaced by $\boxtimes$-g-euler genus enforcer and $\boxtimes$-planarization irrelevant is replaced by $\boxtimes$-g-Euler Genus irrelevant. The generalized version of Lemma 5 follows as the same arguments also hold on bounded-genus graphs: the result we use from [18] has a bounded-genus analogue, the results from [14] and [10] hold for the more general graph class of apex-minor-free graphs. Also the fact that the "big-enough" q-wall that we find is $\boxtimes$-g-Euler Genus irrelevant can be proven using arguments from [25]. Having the extended version of Lemma 5, the proof of the extended version of Lemma 6 is almost identical as we still work inside a disk Δ where G is partially embedded, so that local modifications should locally respect planarity. To be precise, the main difference is that in the definition of d, we now demand that d is also lower bounded by some big-enough function of the genus which guarantees that local modifications in the disk Δ do not alter the genus of the whole graph.

4 Further research directions

In this paper we provide an algorithmic-meta theorem for the graph modifiction problem where the modification operation is in $\{\mathsf{vr}, \mathsf{er}, \mathsf{ec}, \mathsf{ea}\}$ and the target property is planarity plus being a model of some FOL-sentence ϕ. We also argued how to extend this result for modification operations in $\{\mathsf{vr}, \mathsf{er}, \mathsf{ec}\}$ for the case where instead of planarity we consider the class of graphs embeddable in a surface of Euler genus g, for fixed g. The two general challenges that we distinguish are the following.

- Pick a (non-empty) subset $\mathfrak{O}$ of $\{\mathsf{vr}, \mathsf{er}, \mathsf{ec}, \mathsf{ea}\}$ and define Graph $\mathfrak{O}$-Modification to Planarity and ϕ in the obvious way, by permitting any modification operation from $\mathfrak{O}$. It is possible (however more technical) to adapt our results for this problem in the case where $\mathsf{ea} \notin \mathfrak{O}$. However, in the case where $\mathsf{ea} \in \mathfrak{O}$ (while $|\mathfrak{O}| > 1$) the problem becomes considerably more complicated as parts of the graph may be relocated during the modification operation (in fact, from a more general perspective, the same issue appears for the EA-MODIFICATION TO g-EULER GENUS AND ϕ problem that we avoided to consider in Subsection 3.3). We believe that this issue can be tackled using the techniques of [15]. However, the technical details of such an enterprise seem to be quite involved.

- Consider other target properties, alternative to planarity, that are not FOL-expressible. A natural challenge in this direction is to consider some finite set of graphs $\mathcal{H}$ and define the $\boxtimes$-MODIFICATION TO EXCLUDING $\mathcal{H}$-MINORS AND ϕ problem where the target property, apart from being a model of ϕ, is to exclude every graph in $\mathcal{H}$ as a minor. Notice that if $\mathcal{H}$ contains some planar graph, then the yes-instance of the problem has bounded treewidth, therefore the problem is fixed-parameter tractable due to Courcelle's Theorem. The result of this paper can be seen as $\boxtimes$-MODIFICATION TO EXCLUDING $\{K_5, K_{3,3}\}$-MINORS AND ϕ that is the simplest, however essential, version of the general problem. We conjecture that the same results can be achieved for every $\mathcal{H}$ and we believe that the techniques introduced in this paper can be the starting point of such a project.

References

1 Isolde Adler, Stavros G. Kolliopoulos, Philipp Klaus Krause, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Tight bounds for linkages in planar graphs. In Luca Aceto, Monika Henzinger, and Jirí Sgall, editors, *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part I*, volume 6755 of *Lecture Notes in Computer Science*, pages 110–121. Springer, 2011. doi:10.1007/978-3-642-22006-7_10.

2 Chandra Chekuri and Anastasios Sidiropoulos. Approximation algorithms for euler genus and related problems. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 167–176. IEEE Computer Society, 2013. doi:10.1109/FOCS.2013.26.

3 Julia Chuzhoy. An algorithm for the graph crossing number problem. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 303–312. ACM, 2011. doi:10.1145/1993636.1993678.

4 Julia Chuzhoy, Yury Makarychev, and Anastasios Sidiropoulos. On graph crossing number and edge planarization. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 1050–1069. SIAM, 2011. doi:10.1137/1.9781611973082.80.

5 Bruno Courcelle. The monadic second-order logic of graphs. i. recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990. doi:10.1016/0890-5401(90)90043-H.

6 Bruno Courcelle. The monadic second-order logic of graphs III: tree-decompositions, minor and complexity issues. *RAIRO Theor. Informatics Appl.*, 26:257–286, 1992. doi:10.1051/ita/1992260302571.

7 Bruno Courcelle and Sang-il Oum. Vertex-minors, monadic second-order logic, and a conjecture by seese. *J. Comb. Theory, Ser. B*, 97(1):91–126, 2007. doi:10.1016/j.jctb.2006.04.003.

8 Marek Cygan, Dániel Marx, Marcin Pilipczuk, and Michal Pilipczuk. The planar directed k-vertex-disjoint paths problem is fixed-parameter tractable. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 197–206. IEEE Computer Society, 2013. doi:10.1109/FOCS.2013.29.

9 Anuj Dawar, Martin Grohe, and Stephan Kreutzer. Locally excluding a minor. In *22nd IEEE Symposium on Logic in Computer Science (LICS 2007), 10-12 July 2007, Wroclaw, Poland, Proceedings*, pages 270–279. IEEE Computer Society, 2007. doi:10.1109/LICS.2007.31.

10 Erik D. Demaine, Fedor V. Fomin, Mohammad Taghi Hajiaghayi, and Dimitrios M. Thilikos. Bidimensional parameters and local treewidth. *SIAM J. Discret. Math.*, 18(3):501–511, 2004. doi:10.1137/S0895480103433410.

11 Zdenek Dvorák, Daniel Král, and Robin Thomas. Testing first-order properties for subclasses of sparse graphs. *J. ACM*, 60(5):36:1–36:24, 2013. doi:10.1145/2499483.

12 Jörg Flum and Martin Grohe. Fixed-parameter tractability, definability, and model-checking. *SIAM J. Comput.*, 31(1):113–145, 2001. doi:10.1137/S0097539799360768.

13 Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006. doi:10.1007/3-540-29953-X.

14 Fedor V. Fomin, Petr A. Golovach, and Dimitrios M. Thilikos. Contraction obstructions for treewidth. *J. Comb. Theory, Ser. B*, 101(5):302–314, 2011. doi:10.1016/j.jctb.2011.02.008.

15 Fedor V. Fomin, Petr A. Golovach, and Dimitrios M. Thilikos. Modification to planarity is fixed parameter tractable. In Rolf Niedermeier and Christophe Paul, editors, *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13-16, 2019, Berlin, Germany*, volume 126 of *LIPIcs*, pages 28:1–28:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.STACS.2019.28.

16 Markus Frick and Martin Grohe. Deciding first-order properties of locally tree-decomposable structures. *J. ACM*, 48(6):1184–1206, 2001. doi:10.1145/504794.504798.

17 Haim Gaifman. On local and non-local properties. In J. Stern, editor, *Proceedings of the Herbrand Symposium*, volume 107 of *Studies in Logic and the Foundations of Mathematics*, pages 105–135. Elsevier, 1982. doi:10.1016/S0049-237X(08)71879-2.

18 Petr A. Golovach, Marcin Kaminski, Spyridon Maniatis, and Dimitrios M. Thilikos. The parameterized complexity of graph cyclability. *SIAM J. Discret. Math.*, 31(1):511–541, 2017. doi:10.1137/141000014.

19 Petr A. Golovach, Pim van 't Hof, and Daniël Paulusma. Obtaining planarity by contracting few edges. *Theor. Comput. Sci.*, 476:38–46, 2013. doi:10.1016/j.tcs.2012.12.041.

20 Martin Grohe. Logic, graphs, and algorithms. In Jörg Flum, Erich Grädel, and Thomas Wilke, editors, *Logic and Automata: History and Perspectives [in Honor of Wolfgang Thomas]*, volume 2 of *Texts in Logic and Games*, pages 357–422. Amsterdam University Press, 2008.

21 Martin Grohe, Ken-ichi Kawarabayashi, Dániel Marx, and Paul Wollan. Finding topological subgraphs is fixed-parameter tractable. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 479–488. ACM, 2011. doi:10.1145/1993636.1993700.

22 Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3):17:1–17:32, 2017. doi:10.1145/3051095.

23 Bart M. P. Jansen, Daniel Lokshtanov, and Saket Saurabh. A near-optimal planarization algorithm. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1802–1811. SIAM, 2014. doi:10.1137/1.9781611973402.130.

24 Ken-ichi Kawarabayashi. Planarity allowing few error vertices in linear time. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*, pages 639–648. IEEE Computer Society, 2009. doi:10.1109/FOCS.2009.45.

25 Tomasz Kociumaka and Marcin Pilipczuk. Deleting vertices to graphs of bounded genus. *Algorithmica*, 81(9):3655–3691, 2019. doi:10.1007/s00453-019-00592-7.

26 Stephan Kreutzer. Algorithmic meta-theorems. In Javier Esparza, Christian Michaux, and Charles Steinhorn, editors, *Finite and Algorithmic Model Theory*, volume 379 of *London Mathematical Society Lecture Note Series*, pages 177–270. Cambridge University Press, 2011.

27 John M. Lewis and Mihalis Yannakakis. The node-deletion problem for hereditary properties is np-complete. *J. Comput. Syst. Sci.*, 20(2):219–230, 1980. doi:10.1016/0022-0000(80)90060-4.

28 Dániel Marx. Can you beat treewidth? *Theory Comput.*, 6(1):85–112, 2010. doi:10.4086/toc.2010.v006a005.

29 Dániel Marx and Ildikó Schlotter. Obtaining a planar graph by vertex deletion. *Algorithmica*, 62(3-4):807–822, 2012. doi:10.1007/s00453-010-9484-z.

30 Neil Robertson and Paul D. Seymour. Graph minors. II. algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986. doi:10.1016/0196-6774(86)90023-4.

31 Neil Robertson and Paul D. Seymour. Graph minors .xiii. the disjoint paths problem. *J. Comb. Theory, Ser. B*, 63(1):65–110, 1995. doi:10.1006/jctb.1995.1006.

32 Neil Robertson and Paul D. Seymour. Graph minors. XX. wagner's conjecture. *J. Comb. Theory, Ser. B*, 92(2):325–357, 2004. doi:10.1016/j.jctb.2004.08.001.

33 Detlef Seese. Linear time computable problems and first-order descriptions. *Math. Struct. Comput. Sci.*, 6(6):505–526, 1996.