Compound Logics for Modification Problems

FEDOR V. FOMIN and PETR A. GOLOVACH, Department of Informatics, University of Bergen, Bergen, Norway
IGNASI SAU, LIRMM, Univ Montpellier, CNRS, Montpellier, France
GIANNOS STAMOULIS, Institute of Informatics, University of Warsaw, Warsaw, Poland
DIMITRIOS M. THILIKOS, LIRMM, Univ Montpellier, CNRS, Montpellier, France

We introduce a novel model-theoretic framework inspired from graph modification and based on the interplay between model theory and algorithmic graph minors. The core of our framework is a new compound logic operating with two types of sentences, expressing graph modification: the *modulator sentence*, defining some property of the modified part of the graph, and the *target sentence*, defining some property of the resulting graph. In our framework, modulator sentences are in counting monadic second-order logic (CMSO) and have models of bounded treewidth, while target sentences express first-order logic (FO) properties. Our logic captures problems that are not definable in FO and, moreover, may have instances of unbounded treewidth. Our main result is that, for this compound logic, model-checking can be done in quadratic time on minor-free graphs. The proposed logic can be seen as a general framework to capitalize on the potential of the *irrelevant vertex technique*. It gives a way to deal with problem instances of unbounded treewidth, for which Courcelle's theorem does not apply. The proof of our meta-theorem combines novel combinatorial results related to the Flat Wall theorem along with elements of the proof of Courcelle's theorem and Gaifman's theorem. Our algorithmic meta-theorem encompasses, unifies, and extends the known meta-algorithmic results for CMSO and FO on minor-closed graph classes.

CCS Concepts: • **Mathematics of computing → Graph theory**; **Graph algorithms**; • **Theory of computation → Logic**;

Additional Key Words and Phrases: Algorithmic meta-theorems, Graph modification problems, Model-checking, Graph minors, First-order logic, Monadic second-order logic, Flat Wall theorem, Irrelevant vertex technique

The results of this article appeared in the *Proceedings of the 50th International Colloquium on Automata, Languages and Programming (ICALP 2023)* [44].
Fedor V. Fomin, Petr A. Golovach, and Dimitrios M. Thilikos were supported by the Franco-Norwegian project PHC AURORA 2024. Fedor V. Fomin and Petr A. Golovach were supported by the Research Council of Norway via the project BWCA (314528). Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos where supported by the ANR projects DEMOGRAPH (ANR-16-CE40-0028), ESIGMA (ANR-17-CE23-0010), and the French-German Collaboration ANR/DFG Project UTMA (ANR-20-CE92-0027). Ignasi Sau was also supported by the ANR project ELIT (ANR-20-CE48-0008-01). Most of the research work for this article was conducted when Giannos Stamoulis was affiliated with LIRMM, Univ Montpellier, CNRS, Montpellier, France.
Authors' Contact Information: Fedor V. Fomin, Department of Informatics, University of Bergen, Bergen, Norway; e-mail: fedor.fomin@uib.no; Petr A. Golovach, Department of Informatics, University of Bergen, Bergen, Norway; e-mail: petr.golovach@uib.no; Ignasi Sau, LIRMM, Univ Montpellier, CNRS, Montpellier, France; e-mail: ignasi.sau@lirmm.fr; Giannos Stamoulis (corresponding author), Institute of Informatics, University of Warsaw, Warsaw, Poland; e-mail: giannos.stamoulis@mimuw.edu.pl; Dimitrios M. Thilikos, LIRMM, Univ Montpellier, CNRS, Montpellier, France; e-mail: sedthilk@thilikos.info.

ACM 1557-945X/2024/12-ART2
https://doi.org/10.1145/3696451

ACM Reference format:
Fedor V. Fomin, Petr A. Golovach, Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. 2024. Compound Logics for Modification Problems. *ACM Trans. Comput. Logic* 26, 1, Article 2 (December 2024), 57 pages.
https://doi.org/10.1145/3696451

1 Introduction

Our work is kindled by the current algorithmic advances in graph modification. The core of our approach is a novel model-theoretic framework that is based on the interplay between model theory and algorithmic graph minors. Departing from this new perspective, we obtain an algorithmic meta-theorem that encompasses, unifies, and extends all known meta-algorithmic results for **counting monadic second-order logic (CMSO)** and **first-order logic (FO)** on minor-closed graph classes.

1.1 State of the Art and Our Contribution

Modification Problems. A *graph modification problem* asks whether it is possible to apply a series of modifications to a graph in order to transform it to a graph with some desired target property. Such problems have been the driving force of Parameterized Complexity where parameterization quantifies the concept of "distance from triviality" [66] and measures the amount of the applied modification. Classically, modification operations may be vertex or edge deletions, edge additions/contractions, or combinations of them like taking a minor. In their generality, such problems are NP-complete [82, 111] and much research in Parameterized Complexity is on the design of algorithms in time $f(k) \cdot n^{O(1)}$, where the parameter k is some measure of the modification operation [32]. The target property may express desired structural properties that respond to certain algorithmic or combinatorial demands. A widely studied family of target properties are minor-closed graph classes such as edgeless graphs [25], forests [24, 76], bounded treewidth graphs [48, 49, 74], planar graphs [69, 70, 87], bounded genus graphs [77], or, most generally, minor-excluding graphs [1, 88, 101, 103]. However, other families of target properties have also been considered, such as those that exclude an odd cycle [41], a topological minor [50], an (induced) subgraph [23, 34, 100], an immersion [55], or an induced minor [57]. A broad class of graph modification problems concerns cuts. In a typical cut problem, one wants to find a minimum-size set of edges or vertices X in a graph G such that in the new graph $G \setminus X$, obtained by deleting X from G, some terminal-connectivity conditions are satisfied. For example, the condition can be that a set of specific terminals becomes separated or that at least one connected component in the new graph is of a specific size. The development of parameterized algorithms for cut problems is a popular trend in parameterized algorithms [20, 33, 67, 73, 75, 85, 86]. More involved modification measures of vertex set removals, related to treewidth or treedepth, have been considered very recently [3, 21, 22, 39, 68, 88].

Algorithmic Meta-Theorems. A vibrant line of research in Logic and Algorithms is the development of *algorithmic meta-theorems*. According to Grohe and Kreutzer [62], algorithmic meta-theorems state that certain families of algorithmic problems, typically defined by some logical and some combinatorial condition, can be solved "efficiently," under some suitable definition of this term. Algorithmic meta-theorems play an important role in the theory of algorithms as they reveal deep interplays between Algorithms, Logic, and Combinatorics. One of the most celebrated meta-theorems is Courcelle's theorem asserting that graph properties definable in CMSO are decidable in linear time on graphs of bounded treewidth [27–29]; see also [5, 18]. Another stream of research concerns identifying wide combinatorial structures where model-checking for FO can be done in polynomial time. This includes graph classes of bounded degree [107], graph classes of bounded local treewidth [51], *minor-closed graph classes* [42], graph classes locally excluding a minor [35], and more powerful concepts of sparsity, such as having bounded expansion [38, 89–92], nowhere denseness [63], or having bounded twin-width [16]. (See [61, 79] for surveys. Also for results on the combinatorial horizon of FO and CMSO (and its variants) see [15, 16, 63] and [11, 12] respectively.)

Another line of research, already mentioned in [61], is to prove algorithmic meta-theorems for extensions of FO of greater expressibility. Two such extensions have been recently presented. The first one consists in enhancing FO with predicates that can express k-connectivity for every $k \geq 1$. This extension of FO, was introduced independently by Schirrmacher, Siebertz, and Vigny in [106] (under the name FO+conn) and by Bojańczyk in [13] (under the name *separator logic*). The second and more expressive extension, also introduced by Schirrmacher, Siebertz, and Vigny in [106], is FO+DP, that enhances FO with predicates expressing the existence of disjoint paths between certain pairs of vertices. For FO+conn, an algorithmic meta-theorem for model-checking on graphs excluding a topological minor has been very recently given by Pilipczuk, Schirrmacher, Siebertz, Toruńczyk, and Vigny [97]. For the more expressive FO+DP, an algorithmic meta-theorem for model-checking on graphs excluding a minor has been very recently given by Golovach, Stamoulis, and Thilikos in [59] (see [58] for the full version), see also [105].

Research on the meta-algorithmics of FO is quite active and has moved to several directions such as the study of FO-interpretability [14, 53, 93–96] or the enhancement of FO with counting/numerical predicates [37, 64, 80, 81] (see also [40, 60, 65, 110] for other extensions).

Our Contribution. In this article, we initiate an alternative approach consisting in combining the expressive power of FO and **monadic second-order logic (MSO)**. We introduce a *compound logic* that models computational problems through the lens of the "modulator vs target" duality of graph modification problems. Each sentence of this logic is a composition of two types of sentences. The first one, called the *modulator sentence*, models a modification operation, while the second one, called the *target sentence*, models a target property. Informally, our result asserts that if some appropriate version of the modulator sentence meets the meta-algorithmic assumptions of Courcelle's theorem [27] (i.e., CMSO-definability and assuming bounded treewidth) and the target sentence meets the meta-algorithmic assumptions of the theorem of Flum and Grohe [42] (i.e., FO-definability and assuming minor-exclusion), then model-checking for the composed compound sentence can be done, constructively, in quadratic time on graphs excluding some graph as a minor. Our main result (Theorem 3) can be seen as a "two-dimensional product" of the two aforementioned meta-algorithmic results, contains both of them as special cases, and automatically implies the tractability of wide families of problems that *neither* are FO-definable *nor* have instances of bounded treewidth.

1.2 Our Results

In this subsection we give formal statements of our results. We need first some definitions.

Preliminaries on Graphs. All graphs in this article are assumed to be finite and most of our graph definitions are compatible with Diestel's book [36]. Given a graph G, we denote by $\mathrm{cc}(G)$ the set of all connected components of G. For a graph G and a set $X \subseteq V(G)$, the *stellation* of X in G is the graph $\mathsf{stell}(G, X)$ obtained from G if, for every $C \in \mathrm{cc}(G \setminus X)$, we contract all the edges of C to a single vertex v_C. The *torso* of X in G is the graph *torso*(G, X) obtained from $\mathsf{stell}(G, X)$ if, for every v_C where $C \in \mathrm{cc}(G \setminus X)$, we add all edges between neighbors of v_C and finally remove all v_C's from the resulting graph. See Figure 1 for an example of a pair (G, X), the graph $\mathsf{stell}(G, X)$, and the graph $\mathsf{torso}(G, X)$.

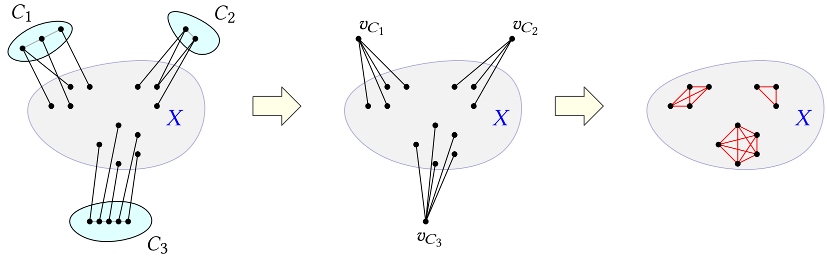

Fig. 1. Left: A graph G, a set X, and the vertex sets C_1, C_2, and C_3 of the connected components of $G \setminus X$. Middle: The graph $\mathsf{stell}(G, X)$. Right: The graph $\mathsf{torso}(G, X)$.

Given a family of graphs $\mathcal{H}$, we define $\mathsf{excl}(\mathcal{H})$ as the class of all graphs minor-excluding the graphs in $\mathcal{H}$ and note that $\mathsf{excl}(\mathcal{H})$ is a minor-closed class (see Section 2 for the definition of minor relation, minor closeness, and minor-exclusion). The *Hadwiger number* of a graph G, denoted by $\mathbf{hw}(G)$, is the minimum k such that G does not contain K_k, i.e., the complete graph on k vertices, as a minor. We also use the well-known parameter of *treewidth* of a graph G, denoted by $\mathbf{tw}(G)$, that is defined in Section 2. Given a class of graphs $\mathcal{G}$ where the treewidth of every graph in $\mathcal{G}$ is bounded by some fixed value, we define $\mathbf{tw}(\mathcal{G}) = \max\{\mathbf{tw}(G) \mid G \in \mathcal{G}\}$. We define $\mathbf{hw}(\mathcal{G})$ analogously. We use $\mathcal{G}_{\text{all}}$ for the class of all finite graphs.

Preliminaries on Logic. We use CMSO (resp. FO) for the set of sentences in CMSO (resp. FO)—see Section 2.3 for the definitions. Given some vocabulary τ and a sentence $\varphi \in \mathsf{CMSO}[\tau]$, we denote by $\mathrm{Mod}(\varphi)$ the class of all finite models of φ, i.e., all finite structures that are models of φ. In this introduction, in order to simplify our presentation, all structures that we consider are either graphs or annotated graphs, i.e., pairs (G, X) where G is a graph and $X \subseteq V(G)$. In the first case $\tau = \{\mathsf{E}\}$, and in the second $\tau = \{\mathsf{E}, \mathsf{X}\}$.

We define the set $\mathsf{CMSO}^{\mathsf{tw}}[\{\mathsf{E}, \mathsf{X}\}]$ as the set that contains every sentence $\beta \in \mathsf{CMSO}[\{\mathsf{E}, \mathsf{X}\}]$ for which there exists some c_β such that the torsos of all the models of β have treewidth at most c_β. Formally,

$$\begin{aligned}\mathsf{CMSO}^{\mathsf{tw}}[\{\mathsf{E}, \mathsf{X}\}] = \{\beta \in \mathsf{CMSO}[\{\mathsf{E}, \mathsf{X}\}] \mid\ & \text{there exists some } c_\beta \text{ such that} \\ & \text{for every graph } G \text{ and every } X \subseteq V(G) \\ & \text{if } (G, X) \models \beta \text{ then } \mathbf{tw}(\mathsf{torso}(G, X)) \leq c_\beta\}.\end{aligned}$$

For simplicity, we use $\mathsf{CMSO}^{\mathsf{tw}}$ and FO as shortcuts for $\mathsf{CMSO}^{\mathsf{tw}}[\{\mathsf{E}, \mathsf{X}\}]$ and $\mathsf{FO}[\{\mathsf{E}\}]$, respectively.

Algorithmic Meta-Theorems. We are now in position to restate three major meta-algorithmic results that were mentioned in the previous subsection.

Proposition 1 (Courcelle [27]). *For every $\beta \in \mathsf{CMSO}^{\mathsf{tw}}$, there is an algorithm deciding membership in $\mathrm{Mod}(\beta)$ in linear time.*

Proposition 2 (Flum and Grohe [42]). *For every $\sigma \in \mathsf{FO}$, there is an algorithm deciding membership in $\mathrm{Mod}(\sigma)$ in quadratic time on graphs of bounded Hadwiger number.*

Some comments are in order. The statements of Proposition 1 and Proposition 2 have been adapted so to incorporate the combinatorial demands in the logical condition. While they can both be stated for structures, we state Proposition 1 for annotated graphs and Proposition 2 for graphs in order to facilitate our presentation. In the classic formulation of Courcelle's theorem, we are given a sentence $\beta \in \mathsf{CMSO}$ and a tree decomposition of bounded treewidth. As such a decomposition can be found in linear time, using e.g., [8, 9, 78], the linearity in the running time of Courcelle's theorem is preserved when it is stated in the form of Proposition 1. For the theorem of Flum and Grohe, the situation is different as the combinatorial demand is minor-exclusion of a clique, which is not definable is FO. As we already mentioned, Proposition 1 and Proposition 2 cannot deal, in general, with modification problems to properties of unbounded treewidth.

We stress that Proposition 1 and Proposition 2 are non-constructive. In order to construct the algorithms promised by Proposition 1, one should also know the bound c_β on the treewidth of the models of $\beta \in \mathsf{CMSO}^{\mathsf{tw}}$. Similarly, for Proposition 2, one should have an upper bound on the Hadwiger number of the input graphs.

A Logic for Modification Problems. As a key ingredient of our result, we define the following operation between sentences. Let $\beta \in \mathsf{CMSO}[\{\mathsf{E},\mathsf{X}\}]$ and $\sigma \in \mathsf{CMSO}[\{\mathsf{E}\}]$. We refer to β as the *modulator* sentence on annotated graphs and to σ as the *target* sentence on graphs. We define $\beta \triangleright \sigma$ so that for every non-empty graph G,

$$G \models \beta \triangleright \sigma \text{ if there is } X \subseteq V(G) \text{ such that } (\mathsf{stell}(G,X),X) \models \beta \text{ and } G \setminus X \models \sigma. \tag{1}$$

In other words, $G \models \beta \triangleright \sigma$ means that the stellation of X in G, along with X, is a model of the modulator sentence β and the $G \setminus X$ is a model of the target sentence σ. That way, β implies the modification operation and σ expresses the target graph property. It is easy to see and we prove formally in Corollary 6 that $\beta \triangleright \sigma \in \mathsf{CMSO}[\{\mathsf{E}\}]$. Given two sets of formulas $\mathcal{L}_1, \mathcal{L}_2$, we set

$$\mathcal{L}_1 \triangleright \mathcal{L}_2 = \{\beta \triangleright \sigma \mid \beta \in \mathcal{L}_1 \text{ and } \sigma \in \mathcal{L}_2\}$$

Our main result is the following.

Theorem 3. *For every $\varphi \in \mathsf{CMSO}^{\mathsf{tw}} \triangleright \mathsf{FO}$, there is an algorithm deciding membership in $\mathrm{Mod}(\varphi)$ in quadratic time on graphs of bounded Hadwiger number.*

Note that Theorem 3 expresses the conditions of Proposition 1 and Proposition 2. Indeed, Proposition 1 follows[1] if β expresses that $X = V(G)$ and σ demands that $G \setminus X$ is the empty graph and Proposition 2 follows if β demands that $X = \emptyset$. In other words, Proposition 1 follows if the target sentence becomes void while Proposition 2 follows if the modulator sentence is void. In this sense, Theorem 3 provides an alternative meta-algorithmic set up between the logical and the combinatorial condition (see Figure 2).

At this point, we would like to stress that our results become also relevant in the context of parameterized algorithm design for graph modification problems. For example, our results imply the existence of fixed-parameter tractable algorithms for deciding the $\mathcal{G}$*-treewidth* of a graph, a notion of modulator measure recently introduced by Eiben, Ganian, Hamm, and Kwon [39] (see also [3, 68]). Having $\mathcal{G}$-treewidth at most k is equivalent to asking that the $\mathsf{torso}(G,X)$ has treewidth at most k and the target property is containment in $\mathcal{G}$. In our case, the class $\mathcal{G}$ can be defined as the models of some given FO-sentence. We refer the reader to [2–4, 19, 21, 22, 39, 47, 68, 84] for an illustrative and non-exhaustive list of recent results, where alternative quality measures of the "modulator" to some graph class were considered, other than just its size.

Compound Logics: Generalizations of $\mathsf{CMSO}^{\mathsf{tw}} \triangleright \mathsf{FO}$. Sentences in $\mathsf{CMSO}^{\mathsf{tw}} \triangleright \mathsf{FO}$ provide an illustrative example of expressing modification problems. One can consider multiple levels of generalizations of this approach: First of all, one can remove minor-exclusion as a combinatorial condition for Theorem 3 and add it as a conjunct in the target formula (minor-exclusion is MSO-definable). Another way to obtain more general target sentences is to ask that the target FO-sentence is satisfied in each connected component of the graph after the removal of the modulator X. These

[1] In particular, Theorem 3 gives a quadratic-time algorithm that contains Proposition 1 as a linear-time black-box procedure for deciding models of bounded treewidth.

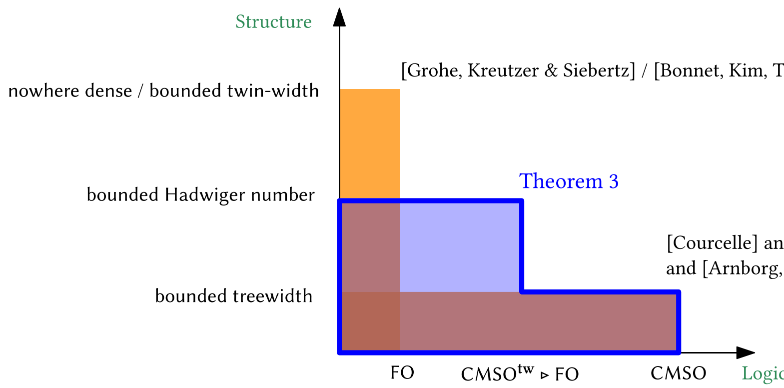

Fig. 2. Theorem 3 in the current meta-algorithmic landscape. The vertical axis is the combinatorial one and is marked by four different types of (structural) sparsity, while the horizontal one is the logical one and is marked with FO, CMSO$^{\mathrm{tw}} \triangleright$ FO, and CMSO.

two approaches can be combined in order to obtain richer, in terms of expressivity, target sentences, and, going even further, one can consider taking positive boolean combinations of the above target sentences. Let us denote by Θ_0 this general target logic that involves all the above ingredients, namely FO, minor-exclusion, connectivity, and positive boolean combinations. Aiming to model more modification problems, one can furthermore consider applying the operator $\triangleright$ recursively, i.e., defining the logic $\Theta = \bigcup_{i\in\mathbb{N}}$, where Θ_{i+1} is defined as CMSO$^{\mathrm{tw}} \triangleright \mathcal{L}_i$, where $\mathcal{L}_i$ is obtained from closing Θ_i under positive boolean combinations and/or connected components. We stress that in all the above considered logics one can include minor-exclusion either as a combinatorial condition for model-checking or as a target demand expressed in the (intermediate or final) target sentence(s). This versatile framework is presented in [43, 44], where the analog of Theorem 3 is proven for all the above compound logics (see [44] for an extended abstract of all these results and see [43] for a more detailed and in-depth presentation and formal proofs). In fact, the results of [43, 44] go even further, by considering generalizations of Θ by replacing FO by the *(Scattered) Disjoint-paths logic* FO+(S)DP. In [43, 44], it is shown that the variant $\Theta^{(\mathrm{s})\mathrm{dp}}$ of Θ obtained by replacing FO by FO+(S)DP also admits a quadratic-time model-checking algorithm on graphs of bounded Hadwiger number (resp. graphs embeddable in some fixed surface).

In this article, we present the simplest and more comprehensive form of the results of the framework introduced in [43, 44], namely the definition of CMSO$^{\mathrm{tw}} \triangleright$ FO and the proof of the model-checking algorithm of Theorem 3. The reason we opt for this is two-fold. First, we wish to illustrate in a more concise and visible way the principal ideas behind the definition of the "modulator vs target" duality of our formulas, already present in the (very special) case of CMSO$^{\mathrm{tw}} \triangleright$ FO. Once familiar with the definition of CMSO$^{\mathrm{tw}} \triangleright$ FO, one can then attempt to consider all the enhancements of CMSO$^{\mathrm{tw}} \triangleright$ FO mentioned above, up to Θ or $\Theta^{(\mathrm{s})\mathrm{dp}}$, to express more intricate modification problems. Also, the algorithm of Theorem 3 can be considered as a functional prototype of the results in [43, 44]. All crucial elements of the proofs in [43, 44] already appear in the proof of Theorem 3 and here are presented in a more streamlined way, avoiding additional technical overload coming from the general framework of [43, 44]. Furthermore, for the proof of Theorem 3, many ideas from [43, 44] are redesigned, improved, and simplified, providing clearer definitions and arguments.

Techniques. The algorithm and the proofs of Theorem 3 use as departure point core techniques from the proofs of Propositions 1 and 2, such as Courcelle's theorem for dealing with CMSO-sentences, the use of Gaifman's theorem for dealing with FO-sentences, and an extended version of the irrelevant vertex technique, introduced by Robertson and Seymour in [98], along with some suitable version of the Flat Wall theorem which appeared recently in [72, 104] (see also [6, 101–103]). The algorithm produces equivalent and gradually "strictly simpler" instances of an annotated version of the problem. Each equivalent instance is produced in linear time and this simplification is repeated until the graph has bounded treewidth (here we may apply Courcelle's theorem, that is Proposition 1). This yields a (constructive) quadratic-time algorithm. We stress that our approach avoids techniques that have been recently used for this type of problems such as *recursive understanding* (in [3]) or the use of *important separators* (in [68]) that give worst running times in n.

Organization of the Article. In Section 2 we provide some basic definitions that will be used throughout the article. In Section 3 we provide an overview of our proof. To describe the algorithm for Theorem 3, we first introduce an annotated version of the problem; this is done in Section 4. Then, in Section 5, we give some preliminary concepts and results and, in Section 6, we present the general scheme of the algorithm for Theorem 3. Section 7 is devoted to the gradual presentation of the main subroutine of the algorithm of Theorem 3 and its correctness. In Appendix A, we present some additional logical background and in particular the framework of logical transductions. We conclude the article with Section 8 by mentioning the limitations of our approach, possible extensions, and open research directions. In Appendix B we present the flat wall framework that we use in this article, which was introduced in [104].

2 Basic Definitions

Next sections are devoted to the formal statement and proof of our results. We present here some basic definitions.

2.1 Integers, Sets, and Tuples

We denote by $\mathbb{N}$ the set of non-negative integers. Given two integers p and q, the set $[p,q]$ refers to the set of every integer r such that $p \leq r \leq q$. For an integer $p \geq 1$, we set $[p] = [1,p]$ and $\mathbb{N}_{\geq p} = \mathbb{N} \setminus [0, p-1]$. Given a non-negative integer x, we denote by $\mathrm{odd}(x)$ the minimum odd number that is not smaller than x. For a set S, we denote by 2^S the collection of all subsets of S and, given an integer $r \in [|S|]$, we denote by $\binom{S}{r}$ the collection of all subsets of S of size r. Given two sets A, B and a function $f : A \to B$, for a subset $X \subseteq A$ we use $f(X)$ to denote the set $\{f(x) \mid x \in X\}$. Let $\mathcal{S}$ be a collection of objects where the operations $\cup$ and $\cap$ are defined. We denote $\bigcup \mathcal{S} = \bigcup_{X \in \mathcal{S}} X$.

2.2 Graphs

Basic Concepts on Graphs. All graphs considered in this article are undirected, finite, and without loops or multiple edges. We use standard graph-theoretic notation and we refer the reader to [36] for any undefined terminology. Let G be a graph. We say that a pair $(L,R) \in 2^{V(G)} \times 2^{V(G)}$ is a *separation* of G if $L \cup R = V(G)$ and there is no edge in G between a vertex in $L \setminus R$ and a vertex in $R \setminus L$. Given a vertex $v \in V(G)$, we denote by $N_G(v)$ the set of vertices of G that are adjacent to v in G. Also, given a set $S \subseteq V(G)$, we set $N_G(S) = \bigcup_{v \in S} N_G(v)$. For $S \subseteq V(G)$, we set $G[S] = (S, E \cap \binom{S}{2})$ and use the shortcut $G \setminus S$ to denote $G[V(G) \setminus S]$. Given a graph G and a set $X \subseteq V(G)$, we denote by $\partial_G(X)$ the set of vertices in X that are adjacent to vertices of $G \setminus X$.

Given a graph G and a set $S \subseteq V(G)$, we define $\mathrm{cc}(G,S)$ to be the collection of the vertex sets of the connected components of $G \setminus S$.

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 $\bigcup_{t \in V(T)} \chi(t) = V(G)$, for every edge e of G there is a $t \in V(T)$ such that $\chi(t)$

contains both endpoints of e, and 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 equal to $\max\left\{ |\chi(t)| - 1 \;\middle|\; t \in V(T)\right\}$ and the *treewidth* of G is the minimum width over all tree decompositions of G.

Contractions and Minors. The *contraction* of an edge $e = \{u, v\}$ of a simple graph G results in a simple graph G' obtained from $G \setminus \{u, v\}$ by adding a new vertex uv adjacent to all the vertices in the set $N_G(u) \cup N_G(v) \setminus \{u, v\}$. A graph G' is a *minor* of a graph G, denoted by $G' \preceq_{\mathsf{m}} G$, if G' can be obtained from G by a sequence of vertex removals, edge removals, and edge contractions. Given a finite collection of graphs $\mathcal{F}$ and a graph G, we use the notation $\mathcal{F} \preceq_{\mathsf{m}} G$ to denote that some graph in $\mathcal{F}$ is a minor of G. Given a family of graphs $\mathcal{F}$, we denote by $\mathsf{excl}(\mathcal{F})$ the graph class containing every graph that excludes all graphs in $\mathcal{F}$ as minors. A graph class $\mathcal{G}$ is *minor-closed* if every minor of a graph in $\mathcal{G}$ is also a member of $\mathcal{G}$.

2.3 FO and MSO

In this subsection, we present some basic notions on logical structures, we define FO and CMSO on structures, and present Gaifman's locality theorem. We refer the reader to [30] for a broader discussion on logical structures and MSO, from the viewpoint of graphs (see also [83]).

Structures. A *vocabulary* is a finite set of relation and constant symbols (we do not use function symbols). Every relation symbol R is associated with a positive integer that is called the *arity* of R, which we denote $\mathsf{ar}(\mathsf{R})$. A *structure* $\mathfrak{A}$ *of vocabulary* τ, in short a τ*-structure*, consists of a non-empty set $V(\mathfrak{A})$, called the *universe* of $\mathfrak{A}$, an r-ary relation $\mathsf{R}^{\mathfrak{A}} \subseteq V(\mathfrak{A})^r$ for each relation symbol $\mathsf{R} \in \tau$ of arity $r \geq 1$, and an element[2] $\mathsf{c}^{\mathfrak{A}} \in \{\emptyset\} \cup V(\mathfrak{A})$ for each constant symbol $\mathsf{c} \in \tau$. We refer to $\mathsf{R}^{\mathfrak{A}}$ (resp. $\mathsf{c}^{\mathfrak{A}}$) as the *interpretation of the symbol R (resp. c) in the structure* $\mathfrak{A}$. A structure $\mathfrak{A}$ is *finite* if its universe $V(\mathfrak{A})$ is a finite set. We denote by $\mathbb{STR}[\tau]$ the family of all finite τ-structures.

Let $\mathfrak{A}$ and $\mathfrak{B}$ be τ-structures (both containing $\emptyset$ to their universe). We say that $\mathfrak{A}$ is a *substructure* of $\mathfrak{B}$, and we write $\mathfrak{A} \subseteq \mathfrak{B}$, if $V(\mathfrak{A}) \subseteq V(\mathfrak{B})$, for every constant symbol $\mathsf{c} \in \tau$, $\mathsf{c}^{\mathfrak{A}} = \mathsf{c}^{\mathfrak{B}}$ if $\mathsf{c}^{\mathfrak{B}} \in V(\mathfrak{A})$ and $\mathsf{c}^{\mathfrak{A}} = \emptyset$ otherwise, and for every relation symbol $\mathsf{R} \in \tau$ of arity $r \geq 1$ we have $\mathsf{R}^{\mathfrak{A}} \subseteq \mathsf{R}^{\mathfrak{B}} \cap V(\mathfrak{A})^r$. We also say that $\mathfrak{A}$ is an *induced substructure* of $\mathfrak{B}$, if $\mathfrak{A} \subseteq \mathfrak{B}$ and for every relation symbol $\mathsf{R} \in \tau$ of arity $r \geq 1$ we have $\mathsf{R}^{\mathfrak{A}} = \mathsf{R}^{\mathfrak{B}} \cap V(\mathfrak{A})^r$. Given a set $S \subseteq V(\mathfrak{A})$, we use $\mathfrak{A}[S]$ to denote the τ-structure with universe S, where $\mathsf{R}^{\mathfrak{A}[S]} = \mathsf{R}^{\mathfrak{A}} \cap S^r$ for each relation symbol $\mathsf{R} \in \tau$ of arity $r \geq 1$ and for each constant symbol $\mathsf{c} \in \sigma$, $\mathsf{c}^{\mathfrak{A}[S]} = \mathsf{c}^{\mathfrak{A}}$, if $\mathsf{c}^{\mathfrak{A}} \in S$, while $\mathsf{c}^{\mathfrak{A}} = \emptyset$, if otherwise.

Let τ be a vocabulary without constant symbols. Given two τ-structures $\mathfrak{A}$ and $\mathfrak{B}$, we define the *disjoint union* of $\mathfrak{A}$ and $\mathfrak{B}$, and we denote it by $\mathfrak{A}\dot{\cup}\mathfrak{B}$, as the τ-structure where $V(\mathfrak{A}\dot{\cup}\mathfrak{B})$ is the disjoint union of $V(\mathfrak{A}) \setminus \{\emptyset\}$, $V(\mathfrak{B}) \setminus \{\emptyset\}$ and $\{\emptyset\}$ and for every relation symbol $\mathsf{R} \in \tau$, $\mathsf{R}^{\mathfrak{A}\dot{\cup}\mathfrak{B}} = \mathsf{R}^{\mathfrak{A}} \cup \mathsf{R}^{\mathfrak{B}}$.

An undirected graph without loops can be seen as an $\{\mathsf{E}\}$-structure $\mathfrak{G} = (V(\mathfrak{G}), \mathsf{E}^{\mathfrak{G}})$, where $\mathsf{E}^{\mathfrak{G}}$ is a binary relation that is symmetric and anti-reflexive. A *vocabulary of annotated graphs* is a vocabulary that contains the binary relation E that is interpreted as a symmetric and anti-reflexive relation and a collection $\mathsf{R}_1, \ldots, \mathsf{R}_h$ of h unary relations, for some $h \in \mathbb{N}$.

First-Order and Monadic Second-Order Logic. We now define the syntax and the semantics of FO and MSO of a vocabulary τ. We assume the existence of a countably infinite set of *first-order variables*, usually denoted by lowercase symbols $\mathsf{x}_1, \mathsf{x}_2, \ldots$, and of a countably infinite set of *set*

[2]We stress that we allow constant symbols to be interpreted as the element $\emptyset$, where $\emptyset$ is an element that is not in $V(\mathfrak{A})$. Throughout this article, we assume that the universe of every given structure is extended by adding the extra element $\emptyset$, while all relation symbols are interpreted as tuples of elements of $V(\mathfrak{A})$, not containing $\emptyset$. Moreover, we assume that for every formula that we consider, quantified first order variables are interpreted as elements of the original universe of the structure (and not $\emptyset$).

variables, usually denoted by uppercase symbols $\mathsf{X}_1, \mathsf{X}_2, \ldots$. A *first-order term* is either a first-order variable or a constant symbol. A *FO formula* of vocabulary τ is built from atomic formulas $\mathsf{x} = \mathsf{y}$ and $(\mathsf{x}_1, \ldots, \mathsf{x}_r) \in \mathsf{R}$, where $\mathsf{R} \in \tau$ and has arity $r \geq 1$, on first-order terms $\mathsf{x}, \mathsf{y}, \mathsf{x}_1, \ldots, \mathsf{x}_r$, by using the logical connectives $\vee, \wedge, \neg$ and the quantifiers $\forall, \exists$ on first-order variables. We denote by $\mathsf{FO}[\tau]$ the set of all FO-formulas of vocabulary τ.

A *MSO formula* of vocabulary τ is obtained by enhancing the syntax of FO-formulas by allowing the atomic formulas $\mathsf{x} \in \mathsf{X}$, for some first-order term x and some set variable X, and allowing quantification both on first-order and set variables. We denote by $\mathsf{MSO}[\tau]$ the set of all MSO-formulas of vocabulary τ. We make clear that what we call here MSO is what is commonly referred in the literature as MSO_1, in which, for the vocabulary of graphs, first-order variables are interpreted as vertices and set variables are interpreted as sets of vertices. Our approach uses Courcelle's theorem for bounded treewidth structures (Proposition 1) as a black-box, which applies for a more general logic than MSO_1, that is MSO_2. For the vocabulary of graphs, MSO_2 extends MSO_1 by also allowing quantification over edges and edge sets (see [30, Subsection 9.2] for formal definition of MSO_2 for general relational vocabularies). Using this fact, our results hold also in the case we define MSO to be MSO_2.

A *CMSO formula* of vocabulary τ is obtained by enhancing the syntax of MSO-formulas by allowing predicates of the form $\mathsf{Card}_p(\mathsf{X})$, expressing that $|\mathsf{X}|$ is a multiple of an integer $p > 1$. We denote by $\mathsf{CMSO}[\tau]$ the set of all CMSO-formulas of vocabulary τ.

The formulas in $\mathsf{FO}[\tau]$ and $\mathsf{CMSO}[\tau]$ are evaluated on τ-structures by interpreting every symbol in τ as its interpretation in the structure and every first-order (resp. set) variable as an element (resp. set of elements) of the universe of the structure. Given a formula φ, the *free variables* of φ are its variables that are not in the scope of any quantifier. We write $\varphi(\mathsf{x}_1, \ldots, \mathsf{x}_k, \mathsf{X}_1, \ldots, \mathsf{X}_\ell)$ to indicate that the free variables of the formula φ are $\mathsf{x}_1, \ldots, \mathsf{x}_k$ (first-order variables) and $\mathsf{X}_1, \ldots, \mathsf{X}_\ell$ (set variables). A *sentence* is a formula without free variables.

Given a τ-structure $\mathfrak{A}$, a formula $\varphi(\mathsf{x}_1, \ldots, \mathsf{x}_k, \mathsf{X}_1, \ldots, \mathsf{X}_\ell) \in \mathsf{CMSO}[\tau]$, $a_1, \ldots, a_k \in V(\mathfrak{A})$, and $A_1, \ldots, A_\ell \subseteq V(\mathfrak{A})$, we write $\mathfrak{A} \models \varphi(a_1, \ldots, a_k, A_1, \ldots, A_\ell)$ to denote that $\varphi(\mathsf{x}_1, \ldots, \mathsf{x}_k, \mathsf{X}_1, \ldots, \mathsf{X}_\ell)$ holds in $\mathfrak{A}$ if, for every $i \in [k]$, the variable x_i is interpreted as a_i and, for every $j \in [\ell]$, the variable X_j is interpreted as A_j. Two formulas $\varphi(\mathsf{x}_1, \ldots, \mathsf{x}_k, \mathsf{X}_1, \ldots, \mathsf{X}_\ell), \psi(\mathsf{x}_1, \ldots, \mathsf{x}_k, \mathsf{X}_1, \ldots, \mathsf{X}_\ell) \in \mathsf{CMSO}[\tau]$ are *equivalent* if for every τ-structure $\mathfrak{A}$, every $a_1, \ldots, a_k \in V(\mathfrak{A})$, and every $A_1, \ldots, A_\ell \subseteq V(\mathfrak{A})$, we have $\mathfrak{A} \models \varphi(a_1, \ldots, a_k, A_1, \ldots, A_\ell) \Leftrightarrow \mathfrak{A} \models \psi(a_1, \ldots, a_k, A_1, \ldots, A_\ell)$. We call the set $\{\mathfrak{A} \in \mathbb{STR}[\tau] \mid \mathfrak{A} \models \varphi\}$ the family of *models of* φ and we denote it by $\mathrm{Mod}(\varphi)$.

For simplicity, we use $\mathsf{CMSO}^{\mathsf{tw}}$ and FO as shortcuts for $\mathsf{CMSO}^{\mathsf{tw}}[\{\mathsf{E}, \mathsf{X}\}]$ and $\mathsf{FO}[\{\mathsf{E}\}]$, respectively.

Gaifman's Locality Theorem. We now aim to present one of the key tools of our proofs, *Gaifman's locality theorem*. For this, we first give some definitions. The *Gaifman graph* $G_{\mathfrak{A}}$ of a τ-structure $\mathfrak{A}$ is the graph with vertex set $V(\mathfrak{A})$ and an edge between two distinct vertices $a, b \in V(\mathfrak{A})$ if there is an $\mathsf{R} \in \tau$ of arity $r \in \mathbb{N}_{\geq 1}$ and a tuple $(a_1, \ldots, a_r) \in \mathsf{R}^{\mathfrak{A}}$ such that $a, b \in \{a_1, \ldots, a_r\}$. Notice that in the particular case of graphs (seen as structures), the original graph and its Gaifman graph are the same.

The *distance* $d_{\mathfrak{A}}(a, b)$ in $\mathfrak{A}$ between two elements $a, b \in V(\mathfrak{A})$ is the length of a shortest path in $G_{\mathfrak{A}}$ connecting a and b. Given an $r \geq 1$ and an $a \in V(\mathfrak{A})$, we define the *r-neighborhood* of a in $\mathfrak{A}$ to be the set $N^{(\leq r)}_{\mathfrak{A}}(a) = \{b \in V(\mathfrak{A}) \mid d_{\mathfrak{A}}(a, b) \leq r\}$. We use $d(a, b)$ instead of $d_{\mathfrak{A}}(a, b)$ and $N^{(\leq r)}(a)$ instead of $N^{(\leq r)}_{\mathfrak{A}}(a)$ when $\mathfrak{A}$ is clear from the context. A first-order formula $\psi(\mathsf{x})$ with one free variable x is called *r-local* if its validity at an element a in the universe of a structure $\mathfrak{A}$ only depends on the r-neighborhood of a in $\mathfrak{A}$, that is $\mathfrak{A} \models \psi(a) \Leftrightarrow \mathfrak{A}[N^{(\leq r)}_{\mathfrak{A}}(a)] \models \psi(a)$.

Observe that, for every $r \in \mathbb{N}$, there is a first-order formula $\delta_r(\mathsf{x}, \mathsf{y})$ such that for every τ-structure $\mathfrak{A}$ and $a, b \in V(\mathfrak{A})$ we have $\mathfrak{A} \models \delta_r(a, b)$ if and only if $d_{\mathfrak{A}}(a, b) \leq r$ (see [109, Lemma 2.4.2] for a proof). In what follows, we write $d(\mathsf{x}, \mathsf{y}) \leq r$ instead of $\delta_r(\mathsf{x}, \mathsf{y})$ and $d(\mathsf{x}, \mathsf{y}) > r$ instead of $\neg\delta_r(\mathsf{x}, \mathsf{y})$. Let $\ell, r \in \mathbb{N}_{\geq 1}$. A *basic local sentence with parameters ℓ and r* is a first-order sentence of the form

$$\exists \mathsf{x}_1 \ldots \exists \mathsf{x}_\ell \Big(\bigwedge_{1 \leq i < j \leq \ell} d(\mathsf{x}_i, \mathsf{x}_j) > 2r \wedge \bigwedge_{i=1}^{\ell} \psi(\mathsf{x}_i) \Big),$$

where ψ is r-local. A *Gaifman sentence* is a Boolean combination of basic local sentences.

Proposition 4 (Gaifman's locality theorem [52]). *Every first-order sentence σ is equivalent to a Gaifman sentence $\breve{\sigma}$. Moreover, $\breve{\sigma}$ can be computed effectively from σ.*

Here, "computed effectively" means that there is a computable function that maps σ to $\breve{\sigma}$. For every sentence $\sigma \in \mathsf{FO}[\tau]$, we use $\breve{\sigma}$ to denote a Gaifman sentence that is equivalent to σ.

2.4 Our Compound Logic

Translating Sentences. One can observe that stell is an MSO-transduction (see Appendix A). Using the *Backwards Translation Theorem* [30, Theorem 1.40] (see also [12, Lemma B.1]) one can obtain the following result (more formally, the next result follows by combining Lemma 19 and Proposition 17, which are presented in Appendix A).

Observation 5. *For every CMSO-sentence ψ on annotated graphs, there is a CMSO-formula $\psi|_{\mathsf{stell}_X}$ on annotated graphs, such that for every annotated graph (G, X), it holds that $(G, X) \models \psi|_{\mathsf{stell}_X} \Leftrightarrow (\mathsf{stell}_X(G, X), X) \models \psi$.*

We also set $\psi|_{\mathsf{rm}_\mathsf{X}}$ to be the sentence obtained from ψ after replacing, for each first-order variable, every occurrence of "$\exists/\forall\ \mathsf{x}$" with "$\exists/\forall\ \mathsf{x} \notin \mathsf{X}^{\mathfrak{A}}$" and, for each set variable Y, every occurrence of "$\exists/\forall\ \mathsf{Y}$" with "$\exists/\forall\ \mathsf{Y}(\mathsf{Y} \cap \mathsf{X}^{\mathfrak{A}} = \emptyset)$."

Compound Sentences. We define the operation $\triangleright$ as follows. Given $\beta \in \mathsf{CMSO}[\tau \cup \{\mathsf{X}\}]$ and $\sigma \in \mathsf{CMSO}[\tau]$, we define

$$\beta \triangleright \sigma := \exists \mathsf{X}\ (\beta|_{\mathsf{stell}_\mathsf{X}} \wedge \sigma|_{\mathsf{rm}_\mathsf{X}}). \tag{2}$$

As a byproduct of Observation 5, we get the following.

Corollary 6. *If $\beta \in \mathsf{CMSO}[\tau \cup \{\mathsf{X}\}]$ and $\sigma \in \mathsf{CMSO}[\tau]$, then $\beta \triangleright \sigma \in \mathsf{CMSO}[\tau]$.*

On a semantical level, given a formula $\beta(\mathsf{X}) \in \mathsf{CMSO}$ and a sentence $\sigma \in \mathsf{CMSO}$, G satisfies $\beta \triangleright \sigma$ if and only if there is a set $X \subseteq V(G)$ such that $(\mathsf{stell}(G, X), X) \models \beta$ and $G \setminus X \models \sigma$. For every sentence $\beta \triangleright \sigma$, we call β the *modulator* sentence of $\beta \triangleright \sigma$ and σ the *target sentence* of $\beta \triangleright \sigma$. We stress that, since $\beta \in \mathsf{CMSO}^{\mathbf{tw}}[\{\mathsf{E}, \mathsf{X}\}]$, the fact that $(\mathsf{stell}(G, X), X) \models \beta$ implies that there is a constant $c_\beta \in \mathbb{N}_{\geq 1}$, such that the graph $\mathsf{torso}(G, X)$ has treewidth at most c_β. We call c_β the *treewidth* of $\beta \triangleright \sigma$. We assume that c_β is upper bounded by a function of β.

3 Overview of the Proof

In this section we summarize some of the main ideas involved in the proof of Theorem 3, while keeping the description at an intuitive level. We would like to stress that some of the informal definitions given in this section are deliberately *imprecise*, since providing the precise ones would result in a huge overload of technicalities that would hinder the flow of the proof.

In Section 3.1 we present the general scheme of the algorithm (see Section 6, in particular Figure 4, for a more detailed presentation). In Section 3.2 we present the sketch of proof of correctness of the algorithm presented in Section 3.1 (see Section 7).

3.1 General Scheme of the Algorithm

We use the *irrelevant vertex technique* introduced by Robertson and Seymour [98]. Our overall strategy is the "typical" one when using this technique: if the treewidth of the input graph G is bounded by an appropriately chosen function, depending only on the sentence $\beta \triangleright \sigma$, then we use Courcelle's theorem [27–29] and solve the problem in linear time, using the fact that our compound formula $\beta \triangleright \sigma$ is a fragment of CMSO (see Section 2.4). Otherwise, we identify an irrelevant vertex in linear time, that is, a vertex whose removal produces an equivalent instance. Naturally, the latter case concentrates all our efforts and, in what follows, we sketch the main ingredients that we use in order to identify such an irrelevant vertex. In a nutshell, our approach is based on introducing a robust combinatorial framework for finding irrelevant vertices. In fact, what we find is *annotation-irrelevant flat territories*, building on our previous recent work [6, 45, 101–104], which is formulated with enough generality so as to allow for the application of powerful tools such as Gaifman's locality theorem (see Proposition 4) or a variant of Courcelle's theorem on boundaried graphs (see Proposition 11).

Flat Walls. An essential tool of our approach is the notion of *flat wall*, originating in the work of Robertson and Seymour [98]. Informally speaking, a flat wall W is a structure made up of (non-necessarily planar) pieces, called *flaps*, that are glued together in a bidimensional grid-like way defining the so-called *bricks* of the wall (see Figure B3). While such a structure may not be planar, it enjoys some topological properties that are similar (in spirit) to those of planar graphs. Namely, two paths that are not routed entirely inside a flap cannot "cross," except at a constant-sized vertex set A whose vertices are called *apices*. Hence, flat walls are only "locally non-planar," and after removing apices we can apply useful locality arguments, in the sense that two vertices that are in "distant" flaps should also be "distant" in the whole graph without the apices. One of the most celebrated results in the theory of Graph Minors by Robertson and Seymour [98, 99], known as the Flat Wall theorem (see Proposition 26 for a variant recently proved in [72, 104]), informally states that graphs of large treewidth contain either a large clique minor or a large flat wall. In this article we use the framework recently introduced in [104] that provides a more accurate view of some previously defined notions concerning flat walls, particularly in [72]. We provide these precise definitions in Subsection B.4, including the concepts of *flatness pair*, *regularity*, *tilt*, and *influence*, and we stress that they are *not* critical in order to understand the main technical contributions of the current article (however, they are critical for their formal correctness). In what follows, when considering a flat wall W with an apex set A in a graph G, for simplicity we refer to W by using indistinguishably the terms "wall" and "compass of a wall," which can be roughly described as the component containing W in the graph obtained from G by removing A and the "boundary" of W (see Section B.4 for the formal definition).

Working with an Annotated Version of the Problem. We start by defining a convenient equivalent version of the problem (see Section 4), by replacing our sentence $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ with an equivalent enhanced sentence $\theta_{\mathrm{R},\mathbf{c}}$. This is done in two steps, presented in Sections 4.1 and 4.2.

Assuming the existence of a flat wall and an apex set in our input graph G, we first transform (see Section 4.1) the question θ on G to a question on a structure obtained from G by "neutralizing" the apex set (Observation 7). The goal of this step is to ask the target FO-sentence σ of our sentence θ in a "flattened" structure, where apices can no longer "bring close" any distant parts of the wall. This transformation of the problem, which we call *apex-projection*, will allow for the application of the locality-based strategy discussed in the definition of the in-signature of a wall in Section 3.2. To do this, we introduce some additional constant symbols $\mathbf{c}$ to our vocabulary that will be interpreted as the apex vertices.

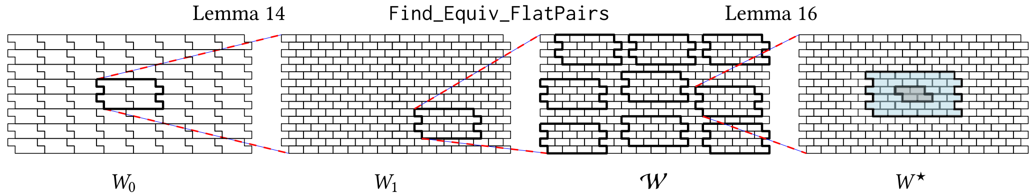

Fig. 3. Sequence of walls considered in the general scheme of our algorithm, along with the results used to obtain them, where the first wall is obtained by applying Proposition 26 to the input graph G.

The second step (Section 4.2) consists in defining an equivalent *annotated* version of the problem in order to deal with the target FO-sentence σ, inspired by the approach of [45]. To do so, we introduce a vertex set $R \subseteq V(G)$, and require that the vertices interpreting the variables of (the equivalent Gaifman sentence of) σ belong to the annotated set R. We prove that the initial sentence θ and the obtained sentence, denoted by $\theta_{\mathrm{R},\mathbf{c}}$ and called an *enhanced sentence*, are equivalent for any choice of the apex set interpreting $\mathbf{c}$ and when R is interpreted as the whole vertex set of the graph (see Lemma 9 and Lemma 10). This independence of the choice of the apex set is strongly used in the proofs since, as discussed below, we will consider a number of different flat walls, each of which is associated with a different apex set.

Our algorithms will work with the enhanced sentence $\theta_{\mathrm{R},\mathbf{c}}$. Starting with the input graph G with $V(G)$ as the annotated set R, we will create successive equivalent annotated instances, in which vertices from G are removed and such that the annotated set R is only reduced.

Zooming Inside a Flat Wall. Our next step is to find, in G, a large flat wall W_0 to work with. Since we work on graphs of bounded Hadwiger number, i.e., graphs that exclude K_c as a minor for some fixed constant $c \in \mathbb{N}$, we can apply Proposition 26 to the input graph G and, assuming that the treewidth of G is large enough as a function of c and θ, we can find in linear time a flat wall W_0 and an apex set A in G such that the height of W_0 is a sufficiently large function of c and θ. Moreover, another crucial property guaranteed by Proposition 26 is that the treewidth of W_0 is bounded from above by a function of c and θ. This will be exploited in Section 3.2 in order to compute the so-called *θ-characteristic* of a wall. We will now apply a series of "zooming" arguments to the wall W_0, which are illustrated in Figure 3 (see the proof of Lemma 15 for the precise constants).

As our next step, we apply Lemma 14 to W_0, and obtain in linear time a still large subwall W_1 such that its associated apex set A_1 is "tightly tied" to W_1, in the sense that the neighbors in W_1 of every vertex in A_1 are spread in a "bidimensional" way. This combinatorial technical condition is critically used in the proof of Lemma 16.

Finding an Irrelevant Subwall. So far, we have found a large wall W_1 that satisfies the conditions listed in the statement of Lemma 16. Now, in order to identify an irrelevant vertex inside W_1, we proceed as follows (see the algorithm `Find_Equiv_FlatPairs` discussed informally in Section 6.3 and presented with all details in Section 7.4). The strategy of the proof is to find, inside the wall W_1, a collection $\mathcal{W}$ of pairwise disjoint subwalls, and to associate each of these subwalls with an appropriately defined *θ-characteristic* that captures its behavior with respect to the partial satisfaction of the sentence θ. Then the idea is that, if there are sufficiently many subwalls in $\mathcal{W}$ with the same θ-characteristic (called *θ-equivalent*), then some subwall in the interior of one of them can be declared *annotation-irrelevant* and this implies some progress in simplifying the current problem instance.

The above strategy is formalized in Lemma 15, which allows to identify a subwall $W^\star$ inside $\mathcal{W}$ such that its central part can be removed from the annotated set R, and such that a central vertex of $W^\star$ can be removed from G (the blue subwall and the red vertex in the rightmost wall of Figure 3, respectively). The proof of Lemma 15 boils down to proving Lemma 16, which is the main technical part of this article, and whose full proof is postponed to Section 7. The proof is based on the algorithm `Find_Equiv_FlatPairs` mentioned above, which is in turn based on an appropriate definition of the *θ-characteristic* of a wall. A brief explanation of the proof strategy of Lemma 16 is given in Section 6.3 and in what follows we sketch the main ingredients and key ideas.

3.2 Defining the Characteristic of a Wall

In order to provide some intuition of the proof of Lemma 16, let us fix some notation. A sentence $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ in a semantical level is defined as follows: Given a general graph G as input, we seek for a vertex set $X \subseteq V(G)$, called *modulator*, such that, using the notation defined in the introduction, $\mathsf{stell}(G, X)$ satisfies the so-called *modulator sentence* β, which is a sentence in $\mathsf{CMSO}^{\mathbf{tw}}$, and the graph $G \setminus X$, satisfies the so-called *target sentence* σ, which is an FO-sentence.

Given the decomposition of θ into two questions (modulator and target), our "irrelevancy" arguments also decompose into two parts. For this, we need to define the *characteristic* of a wall with respect to θ, denoted by θ-char (see Equation (3)). This characteristic is composed of two parts: the *out-signature* (see Section 7.2) corresponding to the satisfiability of the sentence β, and the *in-signature* (see Section 7.3) corresponding to the FO-sentence σ. Let us now explain how we define the out-signature and the in-signature, and sketch why we can eventually declare a subwall irrelevant.

Splitting the Modulator into Two Parts. When X is a modulator, the fact that $\mathsf{torso}(G, X)$ has bounded treewidth implies that every connected component of $G \setminus X$ has a "small interface" to X and thus the flat wall W_0 (and any large subwall of it) is not significantly "damaged" by the removal of X (see Lemma 13). Intuitively (see Section 5.2 for the definition), this means that X intersects a small number of so-called "bags" of the wall. Informally, the *bags* of a wall W in a graph G with apex set A define a partition of $G \setminus A$ into connected sets, such that each bag, except the external one, contains the part of the wall W between two neighboring degree-3 vertices of the wall, as illustrated in Figure B4 (see Section B.6 for the definition). This is a property of every modulator X (as long as $\mathsf{torso}(G, X)$ has bounded treewidth) and it will be used to argue that every modulator X leaves an "intact buffer" in each large enough wall, as explained in the following paragraph.

Defining the Out-Signature of a Wall. Dealing with the irrelevancy with respect to the modulator formula β turns out to be the most interesting part of the proof of Lemma 16, and we introduce several ideas which are, in our opinion, some of the main conceptual contributions of this article. The goal is, for each wall W in the collection $\mathcal{W}$, to encode all the necessary information that concerns the satisfiability of β in the modulator X. To do this, for each $W \in \mathcal{W}$ with apex set A, we define a family of *ℓ-boundaried graphs* (i.e., graphs in which ℓ "boundary" vertices are equipped with labels), constructed as we describe below, and where ℓ depends only on θ. The boundary corresponds to (the boundary of) the part of modulator that is inside the wall. Also, we need to "guess" how to complement this boundary by the part of the modulator that is not inside the wall. Note that, since β is a CMSO-sentence, by a variant of Courcelle's theorem for boundaried graphs [27–29] (see Proposition 11), there exists a *finite* collection $\mathsf{rep}^{(\ell)}(\beta|_{\mathsf{stell}})$ of sentences on ℓ-boundaried graphs that are "representatives" of the sentence $\beta|_{\mathsf{stell}}$ and that can be effectively constructed. We next describe how these ℓ-boundaried graphs are constructed.

We observe that, by Lemma 13 (which uses the bounded-treewidth property of the modulator sentence β), there exists a "buffer" I in W, consisting of a set of consecutive layers of the wall, which is disjoint from a hypothetical modulator X. We guess with an integer d where this "buffer" I is placed in the wall and we denote its inner part by $I^{(d)}$. This naturally induces a partition of

X into X_{in} and X_{out}, with X_{in} being the part of X that is inside $I^{(d)}$ (see Figure 10). We also guess which subset of the apex set A will belong to the modulator X and we denote it by $V_L(\mathbf{a})$, where L is the set containing the indices of the corresponding apex vertices. Since parts of the modulator may lie outside the considered wall, we need to guess the part of the modulator (namely, its boundary towards the component) that lies outside the wall. More precisely, we need to guess as well which subset F' of X_{out}, other than $V_L(\mathbf{a})$, will belong to the neighborhood of the component containing the intact buffer. This is achieved by guessing all ways an (abstract) graph F' with a bounded number of vertices can extend the boundary (see Figure 5). We let F be the graph obtained from the union of $V_L(\mathbf{a})$ and F'. Finally, we also need to consider a set Z that corresponds to X_{in} together with the part inside $I^{(d)}$ that has been "chopped off" by the modulator X, that is, the part of W inside $I^{(d)}$ that will not belong to the component containing the buffer after the removal of the modulator X. We denote by $\partial(Z)$ the set of vertices in Z that have a neighbor in $I^{(d)}$. Altogether, these guesses result in the ℓ-boundaried graph $K^{(d,Z,L,F)}$ obtained from the graph induced by $I^{(d)}$ and the set F after contracting $I^{(d)} \setminus Z$ to a vertex. The boundary is the set $\partial(Z) \cup F$; see Figure 6 for an illustration of $K^{(d,Z,L,F)}$.

With each such a guess (d, L, Z, X) we associate the out-signature defined as follows and denoted by out-sig (see Section 7.2). Its elements are pairs $(\mathbf{H}, \bar{\theta})$, where $\mathbf{H}$ encodes how the set $V_L(\mathbf{a})$ in the boundary has been extended by the "abstract" graph F', and $\bar{\theta} \in \mathrm{rep}^{(\ell)}(\beta|_{\mathsf{stell}})$ prescribes the equivalence class, within the set of Courcelle's representatives mentioned above, of the considered ℓ-boundaried graph. This concludes the description of the out-signature.

While this out-signature indeed encodes the behavior of the considered wall with respect to the modulator sentence β, a crucial issue has been overlooked so far: in order to be able to identify an irrelevant subwall inside the collection $\mathcal{W}$ within the claimed running time, we need to be able to *compute* the (in- and out-) signature of a wall in linear time. To do this using Courcelle's theorem, we need to consider a graph that has treewidth bounded by a function of $\mathbf{hw}(G)$ and θ and has small boundary. By the condition guaranteed by Proposition 26 discussed in the paragraph above Figure 3, we have that the treewidth of W is bounded by a function of $\mathbf{hw}(G)$ and θ, hence the treewidth of the ℓ-boundaried "subwall" $K^{(d,Z,L,F)}$, for which we want to compute the out-signature, is also bounded by a function of $\mathbf{hw}(G)$ and θ. However, the graph $K^{(d,Z,L,F)} \setminus V(F)$ "lives" inside a *whole* component C of the graph $G \setminus X$, and we cannot guarantee that the treewidth of C is bounded by a function of $\mathbf{hw}(G)$ and θ. We overcome this problem with the following trick, which is an important tool in the proof of Claim 1. We observe that the satisfaction of $\beta|_{\mathsf{stell}}$ is preserved if, instead of the whole privileged component C, we consider the graph $K^{(d,Z,L,F)}$, which is obtained by "shrinking" C to a single vertex u_C and which has bounded treewidth as we need (compare the left part of Figure 11 with Figure 12). Indeed, after this modification and by adding edges from the "guessed extended boundary" F' to u_C in order to preserve connectivity (see Figure 6), the resulting graph $\mathsf{stell}(G, X)$ remains unchanged with this transformation, and therefore the satisfaction of the modulator sentence β is also preserved.

Defining the In-Signature of a Wall. To deal with the irrelevancy with respect to the FO-sentence σ, we use arguments strongly inspired by those of [45]. The core tool here is Gaifman's locality theorem (see Proposition 4), which states that every FO-sentence σ is a Boolean combination of basic local sentences $\sigma_1, \ldots, \sigma_p$, in the sense that the satisfaction of each σ_i depends only on the satisfaction of a set of sentences $\psi_1, \ldots, \psi_{\ell_i}$ evaluated on single vertices that can be assumed to be pairwise far apart (see Section 4.2). As discussed before, taking care of the domain of these vertices is the main reason why we consider an annotated version of the problem, corresponding to the enhanced sentence $\theta_{\mathsf{R},\mathbf{c}}$. Extending the approach of [45] (which does not deal with apices), the in-signature of a wall, denoted by in-sig, encodes all (partial) sets of variables, one set for each basic local sentence of the so-called Gaifman sentence $\breve{\sigma}$, such that these variables lie inside an "inner part" of the wall, they are scattered in the "apex-projection" of this inner part, and they satisfy the local sentences ψ_i; see Section 7.3 for the formal definition.

Declaring a Subwall Irrelevant. We now sketch the remaining of the proof of Lemma 16 for sentences in $\mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$, presented in Section 7.5. As mentioned above, suppose that we have already found, inside the collection $\mathcal{W}$, a large (as a function of θ) subcollection $\mathcal{W}' \subseteq \mathcal{W}$ of walls all having the same θ-characteristic. We pick one of these walls, say $W^\star \in \mathcal{W}'$, and we declare its central part irrelevant (see Figure 3). We need to prove that the input graph G satisfies θ, if and only if the graph G' obtained from G by removing the central part of $W^\star$, also satisfies θ. That is, given a modulator X in the original instance G, we need to construct another set $X' \subseteq V(G)$ that is disjoint from $W^\star$ and that is a modulator in G'. For this, we proceed as follows.

The cardinality of $\mathcal{W}'$ and the fact that X intersects few bags of the wall W_3 (see Lemma 13) imply that there exists a large (again, as a function of θ) subcollection $\mathcal{W}'' \subseteq \mathcal{W}'$ of walls that are disjoint from X. We take such a wall $\hat{W} \in \mathcal{W}''$ and, using the fact that $W^\star$ and $\hat{W}$ have the same θ-characteristic, we show that we can "replace" the part of the modulator X that intersects $W^\star$ with another part in $\hat{W}$ (see Figure 15), together with an alternative assignment of variables that satisfies the corresponding sentences. This results in another set X' that is a modulator in G', hence yielding the annotation irrelevancy of (the central part of) $W^\star$.

Showing these facts is far from being easy and we need a number of technical details that are structured into two parts, corresponding to Claim 1 and Claim 2. Each of these claims deals, respectively, with the irrelevancy with respect to β and σ. In particular, an important idea in the proof of Claim 1 is that, changing from X to X', we obtain a new boundaried graph, which is in fact the same graph but with a new boundary (see Figure 15). In the proof of Claim 2, the replacement arguments for the in-signature work because of the aforementioned distance-preservation property of the apex-projection.

4 An Annotated Version of the Problem

In this section we aim to define an enhanced version of every $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$. This is done in two steps. In Section 4.1, we focus on "neutralizing" a tuple $\mathbf{a}$ of vertices of a graph G and transforming a question on G to a question on the structure obtained after "neutralizing" $\mathbf{a}$ (Observation 7). We will apply this tool under the existence of an apex set and a flat wall in the Gaifman graph of our structure, in order to "neutralize" the apex set by adding additional colors in the vertices of our graph and ask the final FO-question of our sentence in a "flattened" colored graph, where apices can no longer "bring close" any distant parts of the wall. This transformation of the problem will allow the application of the "locality-based" strategy that uses Gaifman's locality theorem. In Section 4.2 we define an enhanced version of the problem, by replacing, in a given $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$, the target sentence σ of θ with the sentence obtained from σ after (i) "projecting" it with respect to a set $\mathbf{c}$ of constant symbols (using the definitions in Section 4.1), (ii) taking a Gaifman equivalent sentence of the obtained sentence, and (iii) requiring that the "scattered" variables of the basic local sentences of the Gaifman sentence belong to an annotated set R. We prove that the initial sentence θ and the obtained sentence, denoted by $\theta_{\mathsf{R},\mathbf{c}}$, are "equivalent" for any choice of $\mathbf{a}$ interpreting $\mathbf{c}$ and when R is interpreted as the whole universe of the given structure (Lemma 9). Our algorithms will work with the sentence $\theta_{\mathsf{R},\mathbf{c}}$.

4.1 Dealing with Apices

In this subsection we introduce all necessary tools to handle the (possible) apices in the Gaifman graph of the input graph. As we mentioned in the overview (see Section 3), apices are an obstacle to the locality arguments needed for the part of the proof that concerns FO. To be able to work in a "flat" graph, without the presence of the apices that possibly connect "distant" parts of the graph, we introduce an *apex-projection* of our graph and the corresponding *apex-projection* of a given FO-sentence. This trick appears in [42] and gives an equivalent sentence (see Observation 7).

Apex-Tuples of Structures. Let τ be a vocabulary, let $\mathfrak{A}$ be a τ-structure, and let $l \in \mathbb{N}$. Tuples of the form $\mathbf{a} = (a_1, \dots, a_l)$ where each $a_i \in V(\mathfrak{A}) \cup \{\emptyset\}$, we also call them *apex-tuples* of $\mathfrak{A}$ of size l. We use $V(\mathbf{a})$ for the set containing the non-$\emptyset$ elements in $\mathbf{a}$. Also, if $S \subseteq V(\mathfrak{A})$, we define $\mathbf{a} \cap S = (a'_1, \dots, a'_l)$ so that if $a_i \in S$, then $a'_i = a_i$, and otherwise $a'_i = \emptyset$. We also define $\mathbf{a} \setminus S = \mathbf{a} \cap (V(\mathfrak{A}) \setminus S)$.

Constant-Projections of Vocabularies. Let τ be a vocabulary of annotated graphs, let $l \in \mathbb{N}$, and let $\mathbf{c}$ be a collection of l constant symbols. We define the *constant-projection* $\tau^{\langle \mathbf{c} \rangle}$ *of* $(\tau \cup \mathbf{c})$ to be the vocabulary obtained from $(\tau \cup \mathbf{c})$ by adding l new unary relation symbols $\mathsf{C}_1, \dots, \mathsf{C}_l$.

Projecting a Structure with Respect to an Apex-Tuple. Let $h, l \in \mathbb{N}$. Let $\tau = \{\mathsf{E}, \mathsf{R}_1, \dots, \mathsf{R}_h\}$ be a vocabulary of annotated graphs and let $\mathbf{c} = \{\mathsf{c}_1, \dots, \mathsf{c}_l\}$ be a collection of l constant symbols. Let also $\tau^{\langle \mathbf{c} \rangle}$ be the constant-projection of $(\tau \cup \mathbf{c})$. Given a $(\tau \cup \mathbf{c})$-structure $(\mathfrak{A}, \mathbf{a})$, where $\mathbf{a} = (a_1, \dots, a_l)$ is an apex-tuple of $\mathfrak{A}$ of size l and, for every $i \in [l]$, $\mathsf{c}_i^{\mathfrak{A}} = a_i$, we define the structure $\mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})$ to be the $\tau^{\langle \mathbf{c} \rangle}$-structure obtained as follows:

- $V(\mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})) = V(\mathfrak{A})$,
- $\mathsf{E}^{\mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})} = \mathsf{E}^{\mathfrak{A}} \cap (V(\mathfrak{A}) \setminus V(\mathbf{a}))^2$,
- for every $i \in [h]$, $\mathsf{R}_i^{\mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})} = \mathsf{R}_i^{\mathfrak{A}}$,
- for every $i \in [l]$, $\mathsf{c}_i^{\mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})} = \mathsf{c}_i^{\mathfrak{A}} = a_i$, and
- for every $i \in [l]$, $\mathsf{C}_i^{\mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})} = \{\mathbf{x} \in V(\mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})) \mid \{a_i, x\} \in \mathsf{E}^{\mathfrak{A}}\}$.

Notice that if $a_i = \emptyset$, C_i is interpreted in $\mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})$ as the empty set. It is crucial to see that the Gaifman graph of $\mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})$ is a subgraph of $G_{\mathfrak{A}}$. In fact, $G_{\mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})}$ is obtained from $G_{\mathfrak{A}}$ after removing every edge that is incident to a vertex in $V(\mathbf{a})$. This removal permits us to deal with "flat structures" that are amenable to the application of Gaifman's Theorem.

Apex-Projected Sentences. Let τ be a vocabulary of annotated graphs, let $l \in \mathbb{N}$, and let $\mathbf{c} = \{\mathsf{c}_1, \dots, \mathsf{c}_l\}$ be a collection of l constant symbols. For every sentence $\sigma \in \mathsf{FO}[\tau]$, we define its *$l$-apex-projected sentence* σ^l to be the sentence obtained from σ by replacing every term $\mathsf{E}(\mathsf{x}, \mathsf{y})$ by

$$\mathsf{E}(\mathsf{x}, \mathsf{y}) \vee \bigvee_{i \in [l]} \Big(\big(\mathsf{x} = \mathsf{c}_i \wedge \mathsf{y} \in \mathsf{C}_i \big) \vee \big(\mathsf{y} = \mathsf{c}_i \wedge \mathsf{x} \in \mathsf{C}_i \big) \Big).$$

The definition of the l-apex-projected sentence σ^l implies the following (see [42, Lemma 26]).

Observation 7. *Let τ be a vocabulary of annotated graphs, let $l \in \mathbb{N}$, and let $\mathbf{c}$ be a collection of l constant symbols. For every $\sigma \in \mathsf{FO}[\tau]$, every τ-structure $\mathfrak{A}$, and every apex-tuple $\mathbf{a}$ of $\mathfrak{A}$ of size l, it holds that $\mathfrak{A} \models \sigma \Leftrightarrow \mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a}) \models \sigma^l$ (where $\mathbf{c}$ is interpreted as $\mathbf{a}$).*

Backwards Translating an Apex-Projected Sentence. The above transformation can be expressed in terms of FO-transductions (see Observation 20 in Appendix A). Therefore, given a vocabulary τ, an $l \in \mathbb{N}$, a collection $\mathbf{c}$ of l constant symbols, and a sentence $\sigma \in \mathsf{FO}[\tau]$, we can find a sentence $\sigma' \in \mathsf{FO}[\tau \cup \mathbf{c}]$ such that for every τ-structure $\mathfrak{A}$ and every apex-tuple $\mathbf{a}$ of $\mathfrak{A}$ of size l, $(\mathfrak{A}, \mathbf{a}) \models \sigma' \Leftrightarrow \mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a}) \models \sigma^l$. Again, following the Backwards Translation Theorem (Proposition 17), we get the following:

Table 1. List of Notations Used on Formulas, with Their Respective Meaning and the Results Indicating Their Relation to an Initial Formula σ

Formulas	Relation with σ	Supporting results
$\breve{\sigma}$	Equivalent Gaifman sentence of a sentence $\sigma \in \mathsf{FO}[\tau]$	Proposition 4
σ^l	Formula obtained after "projecting" w.r.t. a tuple $\mathbf{c}$ of size l	Observation 7
$\sigma\vert_{\mathsf{ap}_\mathbf{c}}$	$(\mathfrak{A},\mathbf{a}) \models \sigma\vert_{\mathsf{ap}_\mathbf{c}} \Leftrightarrow \mathsf{ap}_\mathbf{c}(\mathfrak{A},\mathbf{a}) \models \sigma$	Corollary 8

Corollary 8. *Let τ be a vocabulary of annotated graphs, let $l \in \mathbb{N}$, let $\mathbf{c}$ be a collection of l constant symbols, and let $\tau^{\langle \mathbf{c}\rangle}$ be the constant-projection of $\tau \cup \mathbf{c}$. For every sentence $\varphi \in \mathsf{FO}[\tau^{\langle \mathbf{c}\rangle}]$, there exists a sentence $\varphi|_{\mathsf{ap}_\mathbf{c}} \in \mathsf{FO}[\tau \cup \mathbf{c}]$ such that for every $\tau^{\langle \mathbf{c}\rangle}$-structure $\mathfrak{B}$, if $\mathfrak{B} = \mathsf{ap}_\mathbf{c}(\mathfrak{A},\mathbf{a})$ for some $(\tau \cup \mathbf{c})$-structure $(\mathfrak{A},\mathbf{a})$, it holds that $(\mathfrak{A},\mathbf{a}) \models \varphi|_{\mathsf{ap}_\mathbf{c}} \Leftrightarrow \mathsf{ap}_\mathbf{c}(\mathfrak{A},\mathbf{a}) \models \varphi$.*

Concluding this subsection, we present Table 1 that summarizes the notations introduced above for the different kinds of formulas that we consider.

4.2 Introducing an Annotation

In this subsection we present a way to "slightly modify" our sentences in order to construct an enhanced version of every sentence in $\mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$. Let $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ and let σ be its target FO-sentence. Based on the results of Section 4.1, we first consider the l-apex-projected sentence σ^l of σ. We then take an equivalent Gaifman sentence of σ^l. Finally, we add an additional unary relation symbol R to our vocabulary and we ask that the interpretations of the "scattered" variables of each Gaifman sentence belong to the interpretation of R in our structure. This idea is borrowed from [45] but here, on the top of it, we also incorporate the "apex-projection" in order to be able to apply locality arguments inside a "flat" graph.

Restricting the Domain of Variables. Let $l \in \mathbb{N}$, let $\mathbf{c}$ be a collection of l constant symbols, and let R be a unary relation symbol. We now describe how to define an *enhanced version* $\theta_{\mathsf{R},\mathbf{c}}$ of a sentence $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$. The sentence $\theta_{\mathsf{R},\mathbf{c}}$ will be evaluated on $(\{\mathsf{E},\mathsf{R}\} \cup \mathbf{c})$-structures.

Let $\sigma \in \mathsf{FO}[\{\mathsf{E}\}]$ be the target sentence of θ. We consider the l-apex-projected sentence $\sigma^l \in \mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c}\rangle}]$ and we denote it by ζ. By Proposition 4, there is a Gaifman sentence $\breve{\zeta} \in \mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c}\rangle}]$ that is equivalent to ζ. Since $\breve{\zeta}$ is a Gaifman sentence, there exist $p \in \mathbb{N}_{\geq 1}, r_1, \dots, r_p, \ell_1, \dots, \ell_p \in \mathbb{N}_{\geq 1}$, and a collection of sentences $\zeta_1, \dots, \zeta_p \in \mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c}\rangle}]$ such that $\breve{\zeta}$ is a Boolean combination of $\zeta_1, \dots, \zeta_p$ and, for every $h \in [p]$, every ζ_h is a basic local sentence with parameters ℓ_h and r_h, i.e.,

$$\zeta_h = \exists \mathsf{x}_1 \dots \exists \mathsf{x}_{\ell_h} \left(\bigwedge_{1 \leq i < j \leq \ell_h} d(\mathsf{x}_i, \mathsf{x}_j) > 2r_h \wedge \bigwedge_{i \in [\ell_h]} \psi_h(\mathsf{x}_i) \right),$$

where ψ_h is an r_h-local formula in $\mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c}\rangle}]$ with one free variable. Given a Gaifman sentence $\breve{\zeta} \in \mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c}\rangle}]$ as above that is a Boolean combination of sentences $\zeta_1, \dots, \zeta_p \in \mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c}\rangle}]$, we define the sentence $\breve{\zeta}_\mathsf{R}$ to be the sentence in $\mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c}\rangle} \cup \{\mathsf{R}\}]$ that is the same Boolean combination of sentences $\tilde{\zeta}_1, \dots, \tilde{\zeta}_p \in \mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c}\rangle} \cup \{\mathsf{R}\}]$ such that, for every $h \in [p]$,

$$\tilde{\zeta}_h = \exists \mathsf{x}_1 \dots \exists \mathsf{x}_{\ell_h} \left(\bigwedge_{i \in [\ell_h]} \mathsf{x}_i \in \mathsf{R} \wedge \bigwedge_{1 \leq i < j \leq \ell_h} d(\mathsf{x}_i, \mathsf{x}_j) > 2r_h \wedge \bigwedge_{i \in [\ell_h]} \psi_h(\mathsf{x}_i) \right).$$

We define an *enhanced version* $\theta_{\mathsf{R},\mathbf{c}}$ of θ to be a sentence obtained from θ after replacing σ by $\breve{\zeta}_\mathsf{R}|_{\mathsf{ap}_\mathbf{c}}$, where $\zeta = \sigma^l$, i.e., $\theta_{\mathsf{R},\mathbf{c}} = \beta \triangleright \breve{\zeta}_\mathsf{R}|_{\mathsf{ap}_\mathbf{c}}$. Note that since $\breve{\zeta}_\mathsf{R} \in \mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c}\rangle} \cup \{\mathsf{R}\}]$ and $\{\mathsf{E}\}^{\langle \mathbf{c}\rangle} \cup \{\mathsf{R}\} = \{\mathsf{E},\mathsf{R}\}^{\langle \mathbf{c}\rangle}$,

Table 2. List of Formulas to Define an Enhanced Version of a Sentence $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$

Formulas	Meaning
β	the modulator $\mathsf{CMSO}^{\mathbf{tw}}$-sentence of θ
σ	the target FO-sentence of θ
ζ	the l-apex-projected sentence σ^l of σ
$\breve{\zeta}$	a Gaifman sentence equivalent to ζ
ψ_h	r-local formulas of the basic local sentences of $\breve{\zeta}$
$\breve{\zeta}_{\mathsf{R}}$	the Gaifman sentence $\breve{\zeta}$ after adding R (whose model is of the form $\mathsf{ap}_{\mathbf{c}}(G, R, \mathbf{a})$)
$\breve{\zeta}_{\mathsf{R}}\vert_{\mathsf{ap}_{\mathbf{c}}}$	the "backwards translation" of $\breve{\zeta}_{\mathsf{R}}$ to structures without "projecting" $\mathbf{c}$
$\theta_{\mathsf{R},\mathbf{c}}$	$\beta \triangleright \breve{\zeta}_{\mathsf{R}}\vert_{\mathsf{ap}_{\mathbf{c}}}$

it holds that $\breve{\zeta}_{\mathsf{R}}|_{\mathsf{ap}_{\mathbf{c}}} \in \mathsf{FO}[\{\mathsf{E},\mathsf{R}\} \cup \mathbf{c}]$, which in turn implies that $\theta_{\mathsf{R},\mathbf{c}} \in \mathsf{CMSO}[\{\mathsf{E},\mathsf{R}\} \cup \mathbf{c}]$. We also stress that, because of Gaifman's theorem (Proposition 4), for every sentence ζ, there may exist many different Gaifman sentences that are equivalent to ζ. Due to this fact, a sentence $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ can have many enhanced versions. However, all the enhanced versions of θ are equivalent. On the other hand, the proof of Gaifman's theorem implies that there is one effectively computable Gaifman sentence that is equivalent to the given sentence ζ.

We now prove the equivalence between θ and an enhanced version $\theta_{\mathsf{R},\mathbf{c}}$ of θ.

Lemma 9. *Let* R *be a unary relation symbol, and* $\mathbf{c}$ *be a collection of* l *constant symbols, where* $l \in \mathbb{N}_{\geq 1}$. *Also, let* $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ *and let* $\theta_{\mathsf{R},\mathbf{c}}$ *be an enhanced version of* θ. *For every graph* G *and every apex-tuple* $\mathbf{a}$ *of* G *of size* l, *it holds that* $G \models \theta \Leftrightarrow (G, V(G), \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}}$, *where* R *is interpreted as* $V(G)$ *and* $\mathbf{c}$ *is interpreted as* $\mathbf{a}$.

Proof. Let σ be the target sentence of θ and let $\zeta = \sigma^l$. By Observation 7, for every graph G and every apex-tuple $\mathbf{a}$ of G of size l, it holds that $G \models \sigma \Leftrightarrow \mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}) \models \zeta$, where $\mathbf{c}$ is interpreted as $\mathbf{a}$. Also, observe that since R is a unary relation symbol and by the definition of the function $\mathsf{ap}_{\mathbf{c}}$, the structures $(\mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}), V(G))$ and $\mathsf{ap}_{\mathbf{c}}(G, V(G), \mathbf{a})$ are the same. This implies that $\mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}) \models \zeta \Leftrightarrow \mathsf{ap}_{\mathbf{c}}(G, V(G), \mathbf{a}) \models \breve{\zeta}_{\mathsf{R}}$, where R is interpreted as $V(G)$. Thus, by Corollary 8, $G \models \sigma \Leftrightarrow (G, V(G), \mathbf{a}) \models \breve{\zeta}_{\mathsf{R}}|_{\mathsf{ap}_{\mathbf{c}}}$. □

Observe that, for every FO-sentence σ, by Observation 7, for every graph G and every two apex-tuples $\mathbf{a}_1, \mathbf{a}_2$ of G of size l, it holds that $\mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}_1) \models \sigma^l \Leftrightarrow \mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}_2) \models \sigma^l$. Therefore, using Corollary 8, it is easy to prove the following:

Lemma 10. *Let* R *be a unary relation symbol and let* $\mathbf{c}$ *be a collection of* l *constant symbols, where* $l \in \mathbb{N}_{\geq 1}$. *Also, let* $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ *and let* $\theta_{\mathsf{R},\mathbf{c}}$ *be an enhanced version of* θ. *For every graph* G, *for every* $R \subseteq V(G)$, *and every two apex-tuples* $\mathbf{a}_1, \mathbf{a}_2$ *of* G *of size* l, *it holds that* $(G, R, \mathbf{a}_1) \models \theta_{R,\mathbf{c}} \Leftrightarrow (G, R, \mathbf{a}_2) \models \theta_{\mathsf{R},\mathbf{c}}$.

In Table 2, we present all formulas needed to define $\theta_{\mathsf{R},\mathbf{c}}$.

5 Preliminary Tools

In this section we present a series of preliminary results required for our algorithm and its proof of correctness.

Our first tool, presented in Section 5.1, deals with boundaried structures (a generalization of boundaried graphs). Given a sentence $\varphi \in \mathsf{CMSO}[\tau]$, we define an equivalence relation on boundaried structures with respect to the (partial) satisfaction of φ. A variant of Courcelle's theorem (Proposition 11) indicates that there is a finite set of sentences that are evaluated on boundaried structures and are "representatives" of the equivalence classes defined by the above equivalence relation. These "representatives" will help us to "finitize" the way a sentence is partially satisfied (or not) in a boundaried part of our structure.

In Section 5.2 we define a concept that measures the "dispersion" of X inside the "bidimensional territories" of a flatness pair. By flatness pair, here, we mean a flat wall W together with a tuple $\mathfrak{R}$ that certifies its flatness, as defined in [104]; see Section B.4 for a formal definition. We present two results on these notions, namely Lemma 13 and Lemma 14. These two results will be crucial for our algorithm and its correctness.

5.1 A Variant of Courcelle's Theorem

In this subsection we aim to present a variant of Courcelle's theorem (Proposition 11). We start with some definitions on *boundaried structures*.

Boundaried Structures. Given a vocabulary τ of annotated graphs and a non-negative integer ℓ, an *ℓ-boundaried τ-structure* is a tuple $(\mathfrak{A}, x_1, \dots, x_\ell)$, also denoted by $(\mathfrak{A}, \mathbf{x})$, where $\mathfrak{A}$ is a τ-structure and $x_i \in V(\mathfrak{A})$, $i \in [\ell]$. A *boundaried τ-structure* is an ℓ-boundaried τ-structure, for some $\ell \in \mathbb{N}$. We denote by $\mathcal{B}_\tau$ the class of all boundaried τ-structures and, given an $\ell \in \mathbb{N}$, we denote by $\mathcal{B}_\tau^{(\ell)}$ the class of all ℓ-boundaried τ-structures. We treat CMSO-sentences evaluated on ℓ-boundaried τ-structures, as sentences in $\mathsf{CMSO}[\tau \cup \{\mathrm{b}_1, \dots, \mathrm{b}_\ell\}]$, where $\mathrm{b}_1, \dots, \mathrm{b}_\ell$ are constant symbols not contained in τ.

Let $\ell \in \mathbb{N}$. We say that two ℓ-boundaried τ-structures $(\mathfrak{A}, \mathbf{x}), (\mathfrak{B}, \mathbf{y}) \in \mathcal{B}_\tau^{(\ell)}$ are *compatible* if the function that maps x_i to y_i, for every $i \in [\ell]$ is an isomorphism from $\mathfrak{A}[V(\mathbf{x})]$ to $\mathfrak{B}[V(\mathbf{y})]$. Given two compatible ℓ-boundaried τ-structures $(\mathfrak{A}, \mathbf{x})$ and $(\mathfrak{B}, \mathbf{y})$, we define $(\mathfrak{A}, \mathbf{x}) \oplus (\mathfrak{B}, \mathbf{y})$ as the τ-structure obtained if we take the disjoint union of $\mathfrak{A}$ and $\mathfrak{B}$ and then, for every $i \in [\ell]$, identify the elements x_i and y_i, i.e., remove $y_1, \dots, y_\ell$ from the universe of the structure and replace for every $i \in [\ell]$ each occurence of y_i with x_i in the tuples in the interpretation of each relational symbol.

Let τ be a vocabulary and let $\varphi \in \mathsf{CMSO}[\tau]$. We say that two ℓ-boundaried τ-structures $(\mathfrak{A}, \mathbf{x}), (\mathfrak{B}, \mathbf{y}) \in \mathcal{B}_\tau^{(\ell)}$ are *(φ, ℓ)-equivalent*, and we denote it by $(\mathfrak{A}, \mathbf{x}) \equiv_{\varphi,\ell} (\mathfrak{B}, \mathbf{y})$, if they are compatible and for every $(\mathfrak{C}, \mathbf{z}) \in \mathcal{B}_\tau^{(\ell)}$ that is also compatible with $(\mathfrak{A}, \mathbf{x})$ (and $(\mathfrak{B}, \mathbf{y})$) it holds that

$$(\mathfrak{C}, \mathbf{z}) \oplus (\mathfrak{A}, \mathbf{x}) \models \varphi \Leftrightarrow (\mathfrak{C}, \mathbf{z}) \oplus (\mathfrak{B}, \mathbf{y}) \models \varphi.$$

Note that $\equiv_{\varphi,\ell}$ is an equivalence relation on $\mathcal{B}_\tau^{(\ell)}$.

The following result is a variant of Courcelle's theorem [27–29]. It essentially says that the dynamic programming tables constructed by the proof of Courcelle's theorem are also definable in CMSO. This fact is implicit in the proof of Courcelle's theorem. For instance, it can easily be derived from the proof of [10, Lemma 3.2].

Proposition 11 (Courcelle). *There is a function $f : \mathbb{N}^3 \to \mathbb{N}$ such that for every vocabulary τ, every $\varphi \in \mathsf{CMSO}[\tau]$, and every $\ell \in \mathbb{N}$, it holds that $|\mathcal{B}_\tau^{(\ell)}/_{\equiv_{\varphi,\ell}}| \leq f(|\varphi|, \ell, |\tau|)$.*

An alternative way to see Proposition 11 is to say that, for every vocabulary τ, every $\varphi \in \mathsf{CMSO}[\tau]$, and every $\ell \in \mathbb{N}$, there is a collection $\mathsf{rep}_\tau^{(\ell)}(\varphi) = \{\varphi_1, \dots, \varphi_m\}$ of sentences on ℓ-boundaried τ-structures (i.e., sentences in $\mathsf{CMSO}[\tau \cup \{\mathrm{b}_1, \dots, \mathrm{b}_\ell\}]$) where $m \leq f(|\varphi|, \ell, |\tau|)$ and such that

- for every $(\mathfrak{A}, \mathbf{x}) \in \mathcal{B}_\tau^{(\ell)}$ there *exists exactly one* $i \in [m]$ such that $(\mathfrak{A}, \mathbf{x}) \models \varphi_i$ and

– for every compatible $(\mathfrak{A}, \mathbf{x}), (\mathfrak{B}, \mathbf{y}) \in \mathcal{B}_{\tau}^{(\ell)}$ and every $i \in [m]$, if $(\mathfrak{A}, \mathbf{x}) \models \varphi_i$ and $(\mathfrak{B}, \mathbf{y}) \models \varphi_i$, then $(\mathfrak{A}, \mathbf{x}) \equiv_{\varphi,\ell} (\mathfrak{B}, \mathbf{y})$.

The elements of $\mathsf{rep}_{\tau}^{(\ell)}(\varphi)$ are called *types* and can be seen as an CMSO-definable encoding of the tables of the dynamic programming generated by Courcelle's theorem. This representation of φ, in what concerns boundary structures, provides an abstract representation that does not depend on the "internal part" of a boundary graph and will be used as a key ingredient of the encodings in Section 7.

5.2 Dispersion of Sets in Flatness Pairs

Brambles. Let G be a graph. Two sets $V_1, V_2 \subseteq V(G)$ are said to *touch* if they have a vertex in common or there is an edge $\{v_1, v_2\} \in E(G)$ with $v_1 \in V_1$ and $v_2 \in V_2$. A set $\mathcal{B}$ of pairwise touching vertex sets of $V(G)$ that induce connected subgraphs of G is called a *bramble* of G. The *order* of a bramble $\mathcal{B}$ is the minimum size of a hitting set of $\mathcal{B}$, i.e., a vertex set that intersects every element of $\mathcal{B}$.

The following relation between treewidth and a maximum order bramble is proved in [108] (see also [7, Theorem 5]).

Proposition 12. *Let k be a non-negative integer, let G be a graph. The treewidth of G is at most k if and only if every bramble of G has order at most $k + 1$.*

We will prove the following key result. We refer to Section B.4 for the definition of flatness pairs and to Section B.6 for the definition of canonical partitions.

Lemma 13. *Let $t \in \mathbb{N}$, let G be a graph and let $(W, \mathfrak{R})$ be a flatness pair of G. For every set $X \subseteq V(G)$, if $\mathsf{torso}(G, X)$ has treewidth at most t, then X intersects at most $(t + 1)^2$ internal bags of any $(W, \mathfrak{R})$-canonical partition of G.*

Proof. Let $X \subseteq V(G)$ such that $\mathsf{torso}(G, X)$ has treewidth at most t. Also, let $\tilde{\mathcal{Q}}$ be a $(W, \mathfrak{R})$-canonical partition of G. We will show that X intersects at most $(t + 1)^2$ internal bags of $\tilde{\mathcal{Q}}$. Let h be the height of $(W, \mathfrak{R})$. For every $i \in [h]$, let P_i be the union of the vertex sets of all internal bags of $\tilde{\mathcal{Q}}$ that intersect the i-th horizontal path of W, i.e., $P_i := \bigcup_{j \in [2,h-1]} V(Q^{(i,j)})$. Also, let L_i be the union of the vertex sets of all internal bags of $\tilde{\mathcal{Q}}$ that intersect the ith vertical path of W, i.e., $L_i := \bigcup_{j \in [2,h-1]} V(Q^{(j,i)})$. We also define $T^{(i,j)} := P_i \cup L_j, i, j \in [h]$. We let $T_X^{(i,j)} := T^{(i,j)} \cap X$. We now consider the collection

$$\mathcal{D} = \{\mathsf{torso}(G, X)[T_X^{(i,j)}] \mid i, j \in [h] \text{ and } T_X^{(i,j)} \neq \emptyset\}.$$

We will prove that $\mathcal{D}$ is a bramble of $\mathsf{torso}(G, X)$. For this, we have to prove that $\mathcal{D}$ consists of pairwise touching connected subgraphs of $\mathsf{torso}(G, X)$. We use the fact that the graph $\mathsf{compass}_{\mathfrak{R}}(W)$ is connected (see Section B.4 for the formal definition of $\mathsf{compass}_{\mathfrak{R}}(W)$). Notice that if $v, u \in X$ and there is a path P in $\mathsf{compass}_{\mathfrak{R}}(W)$ connecting v and u such that no internal vertex of P is in X, then $\{v, u\} \in E(\mathsf{torso}(G, X))$. This implies that every $D \in \mathcal{D}$ is connected and every two $D_1, D_2 \in \mathcal{D}$ are touching, thus $\mathcal{D}$ is a bramble.

By Proposition 12, we have that $\mathbf{tw}(\mathsf{torso}(G, X)) \leq t$ implies that $\mathcal{D}$ has order at most $t + 1$. This, in turn, implies that X intersects at most $(t + 1)^2$ internal bags of $\tilde{\mathcal{Q}}$. $\square$

The next result intuitively states that given a flat wall and some apices, we can find another flat wall inside the first one such that the set of apices that are adjacent to the compass of the new flat wall are adjacent to "many enough" internal bags of every canonical partition of the graph defined by the latter flat wall. We refer the reader to Section B.4 for the definition of the tilt of a wall inside a flatness pair.

Lemma 14. *There is a function $f_1 : \mathbb{N}^3 \to \mathbb{N}$ and an algorithm that receives as an input two integers $l, d \in \mathbb{N}$, an odd integer $r \geq 3$, a graph G, a set $A \subseteq V(G)$ of size at most l, and a flatness pair $(W, \mathfrak{R})$ of $G \setminus A$ of height $f_1(r, l, d)$, and outputs, in time $O_{r,l,d}(n)$, a set $A' \subseteq A$ and a flatness pair $(\tilde{W}, \tilde{\mathfrak{R}})$ of $G \setminus A'$ of height at least r that is a W'-tilt of some subwall W' of W and for every $a \in A'$ is adjacent to at least d internal bags of every $(\breve{W}, \breve{\mathfrak{R}})$-canonical partition of $G \setminus A'$.*

Proof. Let $l, d \in \mathbb{N}$ and let an odd integer $r \geq 3$. We define the function $f_1 : \mathbb{N}^3 \to \mathbb{N}$ so that, for every $x, z \in \mathbb{N}$, $f_1(x, 0, z) := x$, while, for $y \geq 1$, we set $f_1(x, y, z) := \mathrm{odd}(\lceil \sqrt{z+1} \rceil) \cdot (f_1(x, y-1, z)+2)$.

Let G be a graph, let $A \subseteq V(G)$ of size at most l, and let $(W, \mathfrak{R})$ be a flatness pair of $G \setminus A$ of height $f_1(r, l, d)$. We will prove the lemma by induction on l. In the case that $l = 0$, W has height $f_1(r, 0, d) = r$ and $A = \emptyset$, so the lemma holds trivially for $(W, \mathfrak{R})$. Suppose now that $l \geq 1$ and that the lemma holds for smaller values of l. We set $q := f_1(r, l-1, d)$. Let $\tilde{Q}$ be a $(W, \mathfrak{R})$-canonical partition of $G \setminus A$. If every vertex in A is adjacent in G to at least d internal bags of $\tilde{Q}$, then the algorithm outputs A and $(W, \mathfrak{R})$. Otherwise, there is a vertex $a \in A$ that is adjacent, in G, to less than d internal bags of $\tilde{Q}$. In this case, we consider a collection $\mathcal{W} = \{W_1, \ldots, W_{d+1}\}$ of $d+1$ subwalls of W of height q such that, for every $i, j \in [d+1]$, $i \neq j$, if $(\tilde{W}_i, \tilde{\mathfrak{R}}_i)$ and $(\tilde{W}_j, \tilde{\mathfrak{R}}_j)$ are some W_i-tilt and W_j-tilt of $(W, \mathfrak{R})$ respectively, then $V(\bigcup \mathsf{influence}_{\tilde{\mathfrak{R}}_i}(\tilde{W}_i))$ and $V(\bigcup \mathsf{influence}_{\tilde{\mathfrak{R}}_j}(\tilde{W}_j))$ are disjoint. The existence of this collection is guaranteed by the fact that $f_1(r, l, d) \geq \lceil \sqrt{d+1} \rceil \cdot (q+2)$ and it can be found in time $O_{r,l,d}(n)$. Now notice that since a is adjacent, in G, to less than d internal bags of $\tilde{Q}$, then there is a wall W_i, $i \in [d+1]$ in $\mathcal{W}$ such that a is adjacent, in G, to no internal bag of any $(\tilde{W}_i, \tilde{\mathfrak{R}}_i)$-canonical partition of $G \setminus A$. From the induction hypothesis, we have that we can compute, in time $O_{r,l,d}(n)$, a W''-tilt $(\breve{W}, \breve{\mathfrak{R}})$ of $(W, \mathfrak{R})$, for some W'' that is a subwall of W_i (and therefore of W), that has height at least r and set $A' \subseteq A \setminus \{a\}$ of G of size $l' < l$ such that every vertex in A' is adjacent, in G, to at least d internal bags of every $(\breve{W}, \breve{\mathfrak{R}})$-canonical partition of $G \setminus A'$. □

6 The Algorithm

In this section we aim to present the general scheme of our algorithm for Theorem 3. In Section 6.1, we present the main subroutine of our algorithm that reduces the annotated set R and removes a vertex from the graph under the presence of a flatness pair of "big enough" height in our structure, which is a certificate that the treewidth of the structure is "big enough" (Lemma 15). The proof of Lemma 15 is an almost direct corollary of Lemma 16, whose proof is the main technical part of this article and is postponed to Section 7. A brief explanation of the proof idea is given in Section 6.3. Assuming the claimed algorithm of Lemma 15, in Section 6.2 we show how to use this subroutine to design an algorithm for Theorem 3 and we provide the proof of the latter.

6.1 Reducing the Instance

As we mention in the overview of the proof presented in Section 3, we use the irrelevant vertex technique to reduce the problem to instances of bounded treewidth. This idea is materialized in the next lemma that provides an algorithm that, given an instance $(G, R, \mathbf{a})$, where $\mathbf{a}$ is an apex-tuple of G, and a regular flatness pair $(W, \mathfrak{R})$ of $G \setminus V(\mathbf{a})$ of "big enough" height, such that $\mathsf{compass}_{\mathfrak{R}}(W)$ has bounded treewidth, outputs an instance $(G', R', \mathbf{a})$ such that $V(G') \subsetneq V(G)$, $R' \subsetneq R$, and $(G, R, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}} \Leftrightarrow (G', R', \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}}$.

Lemma 15. *Let R be a unary relation symbol and let $\mathbf{c}$ be a collection of l constant symbols, where $l \in \mathbb{N}_{\geq 1}$. There is a function $f_2 : \mathbb{N}^2 \to \mathbb{N}$ and an algorithm that receives as an input*

- *an enhanced version $\theta_{\mathsf{R},\mathbf{c}}$ of a sentence $\theta \in \mathsf{CMSO}^{\mathsf{tw}} \triangleright \mathsf{FO}$,*
- *a $t \in \mathbb{N}$,*

—*a graph* G*, a set* $R \subseteq V(G)$*, and an apex-tuple* $\mathbf{a}$ *of* G *of size* l*, and*
—*a regular flatness pair* $(W, \mathfrak{R})$ *of* $G \setminus V(\mathbf{a})$ *of height* $f_2(|\theta|, l)$ *such that* $\text{compass}_{\mathfrak{R}}(W)$ *has treewidth at most* t*,*

and outputs, in time $O_{|\theta_{\mathsf{R},\mathbf{c}}|,l,t}(n)$*, a vertex set* $Y \subseteq V(\mathfrak{A}) \setminus V(\mathbf{a})$ *and a vertex* $v \in Y$ *such that* $(\mathfrak{A}, R, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}} \Leftrightarrow (\mathfrak{A} \setminus v, R \setminus Y, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}}$.

To prove Lemma 15, we aim to reduce the annotated set R and to characterize some non-annotated vertex as "irrelevant" to the existence of a solution to the problem, which allows us to reduce our problem to "simpler" equivalent instances. Since our problem has two basic elements[3]

1) the satisfaction of the modulator formula in the modulator sets and
2) the satisfaction of the target FO-sentence in the remaining "terminal part" of the structure,
our "irrelevancy" arguments also decompose into two parts.

Lemma 16. *Let* R *be a unary relation symbol and* $\mathbf{c}$ *be a collection of* l *constant symbols, where* $l \in \mathbb{N}_{\geq 1}$*. There are two functions* $f_3 : \mathbb{N}^3 \to \mathbb{N}$ *and* $f_4 : \mathbb{N} \to \mathbb{N}$ *and an algorithm that receives as an input*

(1) *a sentence* $\theta \in \mathsf{CMSO}^{tw} \triangleright \mathsf{FO}$ *and an enhanced version* $\theta_{\mathsf{R},\mathbf{c}}$ *of* θ*,*
(2) *a* $z \in \mathbb{N}$*,*
(3) *a graph* G*, a set* $R \subseteq V(G)$*, and an apex-tuple* $\mathbf{a}$ *of* G *of size* l*, and*
(4) *a regular flatness pair* $(W, \mathfrak{R})$ *of* $G \setminus V(\mathbf{a})$ *of height at least* $f_3(\mathbf{tw}(\theta), c, l)$*, where* c *is the size of the target sentence of* θ*, such that*
—$\text{compass}_{\mathfrak{R}}(W)$ *has treewidth at most* z *and*
—*every* $a \in V(\mathbf{a})$ *is adjacent to at least* $f_4(\mathbf{tw}(\theta))$ *internal bags of every* $(W, \mathfrak{R})$*-canonical partition of* $G \setminus V(\mathbf{a})$*,*

and outputs, in time[4] $O_{|\theta|,l,z}(n)$*, a set* $Y \subseteq V(G) \setminus V(\mathbf{a})$ *and a vertex* $v \in Y$ *such that* $(G, R, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}} \Leftrightarrow (G \setminus v, R \setminus Y, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}}$.

The proof of Lemma 16 is based on the algorithm `Find_Equiv_FlatPairs`, presented in Section 7.4 (also informally sketched in Section 6.3).

We now provide the proof of Lemma 15, assuming the correctness of Lemma 16. See the down-right green rectangle of Figure 4 for a summary of the main ideas and supporting results of the proof of Lemma 15.

Proof of Lemma 15. Let c be the size of the target sentence of θ. We set

$$d := f_4(\mathbf{tw}(\theta)), \qquad r := f_3(\mathbf{tw}(\theta), c, l), \text{ and } \qquad f_2(|\theta|, l) := f_1(r, l, d).$$

We apply the algorithm of Lemma 14 for $r, l, d, G, V(\mathbf{a})$, and $(W, \mathfrak{R})$. In time $O_{|\theta_{\mathsf{R},\mathbf{c}}|}(n)$, we obtain an apex-tuple $\mathbf{a}'$ of G of size $l' \leq l$ and a flatness pair $(\tilde{W}, \tilde{\mathfrak{R}})$ of $G \setminus V(\mathbf{a}')$ of height r with the following properties: 1) $(\tilde{W}, \tilde{\mathfrak{R}})$ is a $W^\star$-tilt of a subwall $W^\star$ of W and 2) every $a \in V(\mathbf{a}')$, is adjacent to at least d internal bags of any $(\tilde{W}, \tilde{\mathfrak{R}})$-canonical partition of $G \setminus V(\mathbf{a}')$. Note that by Observation 23, $(\tilde{W}, \tilde{\mathfrak{R}})$ is also regular. The wall $\tilde{W}$ corresponds to the selected wall inside the second wall of Figure 3. By applying Lemma 16, we can find, in time $O_{|\theta_{\mathsf{R},\mathbf{c}}|,z}(n)$, a set $Y \subseteq V(G) \setminus V(\mathbf{a}')$ and a vertex $v \in Y$ such that $(G, R, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}} \Leftrightarrow (G \setminus v, R \setminus Y, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}}$. □

[3]Throughout the reminder of the article, we use consistently this color coding using blue/green to easily identify these two parts of our problem.
[4]Given two functions $\chi, \psi \colon \mathbb{N} \to \mathbb{N}$, we write $\chi(n) = O_x(\psi(n))$ to denote that there exists a computable function $f \colon \mathbb{N} \to \mathbb{N}$ such that $\chi(n) = O(f(x) \cdot \psi(n))$.

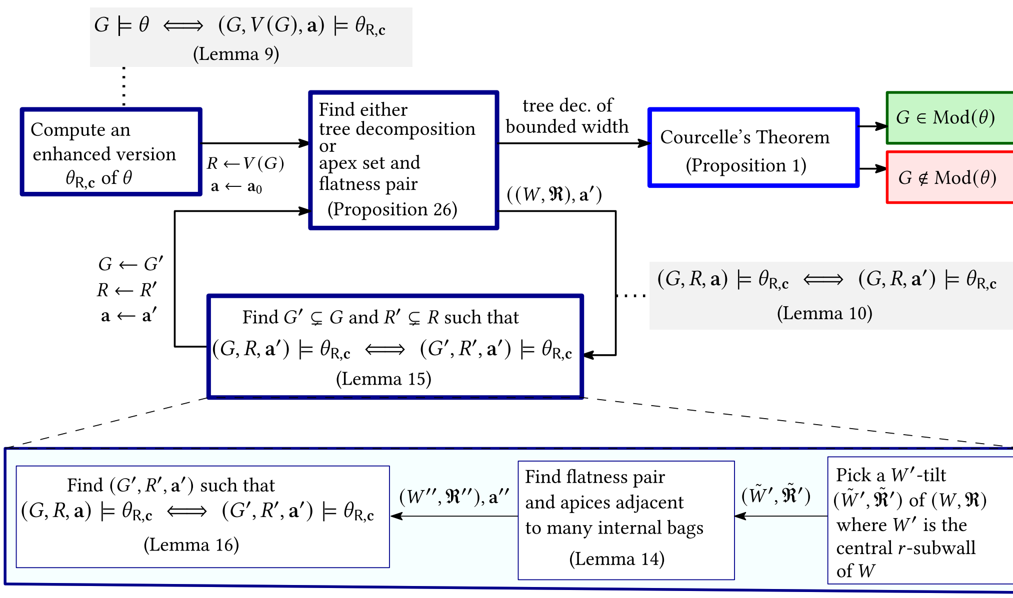

Fig. 4. The flow of the algorithm in the proof of Theorem3 along with the supporting results.

6.2 The Algorithm of Theorem 3

We are now ready to present the proof of Theorem 3 (assuming the correctness of Lemma 16 and therefore of Lemma 15 as well).

Proof of Theorem 3. Given a sentence $\theta \in \mathsf{CMSO}^{\mathsf{tw}} \triangleright \mathsf{FO}$, we set $c = \mathbf{hw}(G)$,

$$l := f_6(c) \text{ where } f_6 \text{ is the function of Proposition 26, and}$$
$$r := f_2(|\theta|, l, 3).$$

Our algorithm consists of four steps, which are summarized in Figure 4, along with the supporting results:

Step 1. Consider an enhanced version $\theta_{\mathsf{R},\mathbf{c}}$ of θ. Consider an arbitrary apex-tuple $\mathbf{a}_0$ of G of size l. By Lemma 9, we have that $G \models \theta \Leftrightarrow (G, V(G), \mathbf{a}_0) \models \theta_{\mathsf{R},\mathbf{c}}$, where R is interpreted as $V(G)$ and $\mathbf{c}$ is interpreted as $\mathbf{a}_0$. We set $R_0 := V(G)$ and we proceed to Step 2.

Step 2. Run the algorithm of Proposition 26 for G, r, and c. Since $K_c \not\leq_{\mathsf{m}} G$, this algorithm outputs, in linear time, either a tree decomposition of G of width at most $f_7(c) \cdot r$, or a set $A \subseteq V(G)$, where $|A| \leq l$, a regular flatness pair $(W, \mathfrak{R})$ of $G \setminus A$ of height r, and a tree decomposition of $\mathsf{compass}_{\mathfrak{R}}(W)$ of width at most $f_7(c) \cdot r$. In the first possible output, i.e., a tree decomposition of G of width at most $f_7(c) \cdot r$, proceed to Step 4. In the second possible output, proceed to Step 3.

Step 3. We first consider an ordering $a_1, \ldots, a_l$ of the vertices in A, and set $\mathbf{a} = (a_1, \ldots, a_l)$. By Lemma 10, we have that $(G, R_0, \mathbf{a}_0) \models \theta_{\mathsf{R},\mathbf{c}} \Leftrightarrow (G, R_0, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}}$. We run the algorithm of Lemma 15 for $\theta_{\mathsf{R},\mathbf{c}}$, G, R_0, $\mathbf{a}$, and $(W, \mathfrak{R})$, and we obtain, in linear time, a set $Y \subseteq V(G) \setminus V(\mathbf{a})$ and a vertex $v \in X$ such that $(G, R_0, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}} \Leftrightarrow (G \setminus v, R_0 \setminus Y, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}}$. Then, we set $G := G \setminus v$, $\mathbf{a}_0 := \mathbf{a}$, $R_0 := R_0 \setminus Y$, and we run again Step 2.

Step 4. Given a tree decomposition of G of width at most $f_7(c) \cdot r$, and since $\theta_{\mathsf{R},\mathbf{c}} \in \mathsf{CMSO}[\{\mathsf{E}, \mathsf{R}\} \cup \mathbf{c}]$, we decide whether $(G, R_0, \mathbf{a}_0) \models \theta_{\mathsf{R},\mathbf{c}}$ in linear time by using Courcelle's theorem.

Observe that the second and the third step of the algorithm are executed in linear time and they can be repeated no more than a linear number of times. Therefore, the overall algorithm runs in quadratic time, as claimed.

□

6.3 Sketch of Proof of Lemma 16

In the next section, we aim to provide a proof for Lemma 16. In this subsection we give a brief description of the main ideas of this proof. Consider a sentence $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$, a formula $\beta \in \mathsf{CMSO}^{\mathbf{tw}}$, and a sentence $\sigma \in \mathsf{FO}$ such that θ can be written as $\beta \triangleright \sigma$. Also, consider an enhanced version $\theta_{\mathsf{R},\mathbf{c}}$ of θ. This can be written as $\beta \triangleright \breve{\zeta}_{\mathsf{R}}|_{\mathsf{ap}_\mathbf{c}}$. To deal with the sentences β and $\breve{\zeta}_{\mathsf{R}}|_{\mathsf{ap}_\mathbf{c}}$, we define the out-signature (Section 7.2) and the in-signature (Section 7.3) of a flatness pair, respectively, and the combination of these two constitutes the characteristic of a flatness pair. This characteristic is an "encoding" of the partial satisfaction of β and $\breve{\zeta}_{\mathsf{R}}|_{\mathsf{ap}_\mathbf{c}}$ inside the flatness pair, and it is worth noting that it is CMSO-definable. After defining this characteristic, we use the following algorithm, that is formally presented in Section 7.4.

The Algorithm `Find_Equiv_FlatPairs`. The algorithm has the following four steps.

- Compute a collection of z subwalls $W_1, \dots, W_z$ of W, where z is some "big enough" integer depending on the sentence θ, such that the compasses of all W_i-tilts of $(W, \mathfrak{R})$ are pairwise disjoint (this collection of walls virtually corresponds to the walls inside the third wall of Figure 3).
- Compute a W_i-tilt of $(W, \mathfrak{R})$ for each $i \in [z]$. These define a collection $\tilde{\mathcal{W}}$ of z flatness pairs.
- For each of the flatness pairs in $\tilde{\mathcal{W}}$, compute its characteristic.
- Output a collection $\tilde{\mathcal{W}}'$ of at least m flatness pairs, a vertex subset Y, and a vertex v with the following properties:
  (1) all flatness pairs in $\tilde{\mathcal{W}}'$ have the same characteristic,
  (2) the set Y is the vertex set of $\mathsf{compass}_{\breve{\mathfrak{R}}'}(\breve{W}')$, where $(\breve{W}', \breve{\mathfrak{R}}')$ is a $\breve{W}$-tilt of $(W, \mathfrak{R})$ and $\breve{W}$ is the central j'-subwall of W_0, for some $(W_0, \mathfrak{R}_0) \in \tilde{\mathcal{W}}'$ (W_0 virtually corresponds to the fourth wall in Figure 3), and
  (3) v is a central vertex v of W_0.

  In the fourth wall of Figure 3, Y corresponds to the light blue area and v belongs to the innermost part of the wall.

After detecting Y and v, it remains to prove that $(G, R, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}} \Leftrightarrow (G \setminus v, R \setminus Y, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}}$. The proof of the above is presented in Subsection 7.5 and is split into two parts, corresponding to Claim 1 and Claim 2.

7 Proof of Lemma 16

This section is structured as follows. In Section 7.1, we define the *extended compass* of a flatness pair, that is a tuple that contains all necessary information around a flatness pair. In Sections 7.2 and 7.3, we define the *out-signature* and the *in-signature* of the extended compass of a flatness pair that encodes how a partial solution (partial assignment of vertices to the variables) satisfies the two parts of θ respectively. Finally, in Section 7.4, we present the algorithm `Find_Equiv_FlatPairs` and, in Section 7.5, we prove that this algorithm correctly returns the claimed output of Lemma 16.

7.1 Extended Compasses of Flatness Pairs

Extended Compasses of Flatness Pairs. Let $l, r \in \mathbb{N}$ and $j \in \mathbb{N}_{\geq 3}$. Let G be a graph, let $\mathbf{a}$ be an apex-tuple of G of size l, and let $(W, \mathfrak{R})$ be a flatness pair of $G \setminus V(\mathbf{a})$ of height $2r + j$. For every subwall W' of W, we denote by $\mathsf{influence}_{\mathfrak{R}}(W')$ the set of the flaps of the flat wall W that either contain an edge of the perimeter of W' or are "embedded" inside the disk "cropped" by the perimeter of W'. Intuitively, $\mathsf{influence}_{\mathfrak{R}}(W')$ contains all flaps "captured" by the wall W'. The graph $\mathsf{compass}_{\mathfrak{R}}(W)$ is always assumed to be connected. See Section B.4 for a formal definition of the above notions. We set $K := \mathsf{compass}_{\mathfrak{R}}(W)$ and $K^{\mathbf{a}} := G[V(\mathbf{a}) \cup V(K)]$. Also, for every $i \in [r]$, let $I^{(i)} = V(\bigcup \mathsf{influence}_{\mathfrak{R}}(W^{(2i+j)}))$ and let $\mathbf{I} = (I^{(1)}, \dots, I^{(r)})$.

Let G be a graph, let $\mathbf{a} = (a_1, \dots, a_l)$ be an apex-tuple of G, and let $(W, \mathfrak{R})$ be a flatness pair of $G \setminus V(\mathbf{a})$ of height $2r + j$. We call the tuple $\mathfrak{K} = (G[V(K^{\mathbf{a}})], \mathbf{a}, \mathbf{I})$ *the extended compass of the flatness pair* $(W, \mathfrak{R})$ *of* $G \setminus V(\mathbf{a})$. Given a $Z \subseteq V(K)$, we define $\partial_{\mathfrak{K}}(Z)$ to be the set $\partial_K(Z)$ (we remind here that $\partial_K(Z)$ is the set of vertices in Z that are adjacent to vertices of $K \setminus Z$). Also, if $L \subseteq [l]$, then $V_L(\mathbf{a})$ contains all non-$\emptyset$ elements in $\mathbf{a}$ indexed by L.

Intuitively, $\mathfrak{K} = (G[V(K^{\mathbf{a}})], \mathbf{a}, \mathbf{I})$ contains all the "useful information" around the flatness pair $(W, \mathfrak{R})$. The structure $G[V(K^{\mathbf{a}})]$ induced by the union of the $\mathfrak{R}$-compass K of W, the apices $V(\mathbf{a})$, the homocentric zones of influence of the layers of W (away from its j-central part).

7.2 Out-Signature

In this subsection we aim to "encode" all necessary information that concerns the satisfiability of the sentence β in the "modulator" part of the input structure. To do this, given an extended compass of a flatness pair, we define a certain boundaried structure and we will use the finite set of representatives of $\beta|_{\mathsf{stell}_X}$ given by Proposition 11, to "associate" this boundaried structure with a representative. Since parts of the modulator X may lie outside the extended compass of the considered flatness pair, we may have to "extend" the boundary of our structure in order to encode which disjoint parts of the modulator have neighbors in the same connected components of $G \setminus X$. This is achieved by "guessing" all ways an (abstract) graph with a bounded number of vertices can "extend" the boundary.

Before presenting some additional definitions, we first set up the sentences and the constants in which we will build the out-signature. Let R be a unary relation symbol and let $\mathbf{c}$ be a collection of l constant symbols, where $l \in \mathbb{N}_{\geq 1}$. Let $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$, let $q := (\mathbf{tw}(\theta) + 1)^2 + 1$, and let $\theta_{\mathsf{R},\mathbf{c}}$ be an enhanced version of θ. Recall that $\theta_{\mathsf{R},\mathbf{c}} = \beta \triangleright \check{\zeta}_{\mathsf{R}}|_{\mathsf{ap}_{\mathbf{c}}}$ and $\check{\zeta}_{\mathsf{R}}|_{\mathsf{ap}_{\mathbf{c}}} \in \mathsf{FO}[\{\mathsf{E}, \mathsf{R}\} \cup \mathbf{c}]$.

For the rest of this subsection, keep in mind that $q = (\mathbf{tw}(\theta) + 1)^2 + 1$. Let $r \in \mathbb{N}$ and an odd $j \in \mathbb{N}_{\geq 3}$. Also, we set $w = (r + 2) \cdot q$.

To give an intuition for the above, let us explain what j and w represent: First, j is the size of the wall whose compass we want to declare annotation-irrelevant in the proof of Lemma 16, i.e., the set Y will be the $\check{\mathfrak{R}}'$-compass of a wall $\check{W}'$ of height j. We will consider a wall W of height $2w + j$ that, apart from its j-central part (corresponding to Y), contains q "annulus buffers" of thickness $(r + 2)$. We stress that while, for now, r is a given constant, in Section 7.3, it will be a particular constant depending on the parameters of the basic local sentences in the definition of $\check{\zeta}_{\mathsf{R}}$. As a modulator X cannot affect more than $(\mathbf{tw}(\theta) + 1)^2 = q-1$ of these q "annulus buffers" (Lemma 13), one of them will not be affected by the solution. This "buffer" will allow us to focus on a single connected component of the graph $G \setminus X$ (the one that contains this buffer) and take the union of all parts of X that are "cropped" by this buffer and all other components of $G \setminus X$ containing some piece of the compass to obtain a single set Z. This trick will allow us to argue that the graph induced by Z (with a small number some extra vertices) has a small boundary (since the neighborhood of X to a component of $G \setminus X$ is bounded by $\mathbf{tw}(\theta)$, due to the fact that the torso of X should have bounded treewidth).

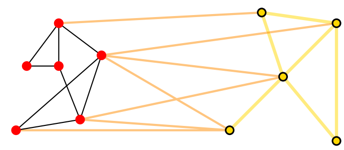

Fig. 5. A set S of 6 vertices of a graph, a graph F' on 5 vertices and a set E of edges between vertices of S and of $V(F')$. The graph $(F' \cup (S \cup V(F'), E))$ belongs to $\mathcal{F}_{11}^{S}$.

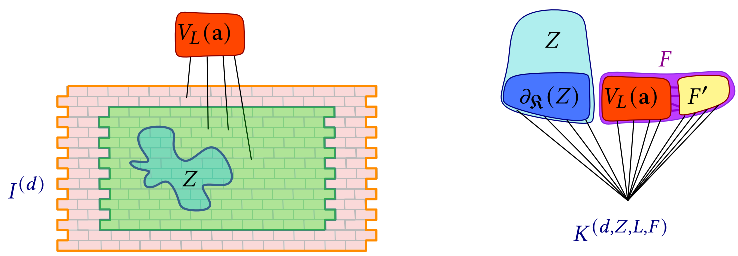

Fig. 6. Left: The position of Z in $I^{(d)}$, where Z is depicted in blue and the set $I^{(d-r)}$ is depicted in green. Right: The graph $K^{(d,Z,L,F)}$.

Guessing an Extension of a Vertex Set. Let G be a graph. Given a set of vertices $S \subseteq V(G)$ and an $\ell \in \mathbb{N}$, we define the collection of graphs $\mathcal{F}_{\ell}^{S}$, such that $F \in \mathcal{F}_{\ell}^{S}$ if and only if there exists a graph F' on $\max\{\ell - |S|, 0\}$ vertices and a set E of edges each with one endpoint in S and the other endpoint in $V(F')$, such that $F = F' \cup (S \cup V(F'), E)$ (see Figure 5 for an example). We stress that, for every $F \in \mathcal{F}_{\ell}^{S}$, $F[S]$ is an edgeless graph.

Notice that if $\ell \leq |S|$ then $\mathcal{F}_{\ell}^{S}$ contains only the graph with vertex set S and no edges.

Towards Constructing a Boundaried Structure. Let q, j, r, w, and l as above. Let G be a graph, let $\mathbf{a}$ be an apex-tuple of G of size l, and let $(W, \mathfrak{R})$ be a flatness pair of $G \setminus V(\mathbf{a})$ of height $2w + j$. Also, let $\mathfrak{K} = (G[V(K^{\mathbf{a}})], \mathbf{a}, \mathbf{I})$ be the extended compass of the flatness pair $(W, \mathfrak{R})$ of $G \setminus V(\mathbf{a})$ and let $\ell \in [0, \mathbf{tw}(\theta) - 1]$. Given a $d \in [r, w]$, an $L \subseteq [l]$, a vertex set $Z \subseteq I^{(d-r+1)}$, and a graph $F \in \mathcal{F}_{\ell}^{V_L(\mathbf{a})}$, we define the graph $K^{(d,Z,L,F)}$ as the one obtained from $K^{\mathbf{a}}[Z \cup V_L(\mathbf{a})] \cup F$ by adding an extra vertex v and making it adjacent to all vertices in $\partial_{\mathfrak{R}}(Z) \cup V(F)$. See Figure 6 for a visualization of $K^{(d,Z,L,F)}$. Intuitively, it contains the set Z that is the part of $I^{(d-r)}$ that is cropped out after the removal of the modulator X (except of the part that contains the buffer $I^{(d)} \setminus I^{(d-r)}$), the apices that we guess that will belong to X, and the part F' of F that corresponds to X_{out}, i.e., the portion of the modulator X that will not be part of $I^{(d)}$. The graph F in Figure 6 is the graph containing all vertices $V_L(\mathbf{a})$ and the "extra" guessed part F' together with the extra edges from $V(F')$ to $V_L(\mathbf{a})$. Let us explain the motivation behind adding these extra edges: The reason we consider the graph $K^{(d,Z,L,F)}$ is to "focus" inside $I^{(d)}$ and temporarily "forget" what happens outside $I^{(d)}$. However, we need to keep record of the fact that $I^{(d)} \setminus Z$ is in the same connected component as $V(F')$. This is why we add the extra vertex and we make it adjacent to $\partial_{\mathfrak{R}}(Z) \cup V(F)$.

Strongly Isomorphic Graphs. Let G be a graph. A *nice 3-partition* of G is an ordered partition $\mathcal{V} = (V_1, V_2, V_3)$ of $V(G)$ such that $(V_1 \cup V_2, V_2 \cup V_3)$ is a separation of G (see Figure 7 for an example).

For every $\ell \in \mathbb{N}$, let

$$\mathcal{H}^{(\ell)} = \{(H, \mathcal{V}) \mid H \text{ is a graph on } \ell \text{ vertices and } \mathcal{V} \text{ is a nice 3-partition of } H\}.$$

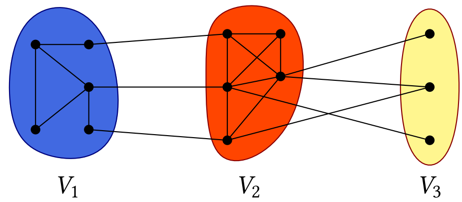

Fig. 7. An example of a nice 3-partition (V_1, V_2, V_3) of a graph.

Let G and H be two graphs and $\mathcal{V} = (V_1, V_2, V_3)$ and $\mathcal{U} = (U_1, U_2, U_3)$ be nice 3-partitions of G and H, respectively. We say that G is *strongly isomorphic to H with respect to $(\mathcal{V}, \mathcal{U})$*, if G is isomorphic to H, $G[V_1 \cup V_2]$ is isomorphic to $H[U_1 \cup U_2]$, $G[V_2 \cup V_3]$ is isomorphic to $H[U_2 \cup U_3]$, and these two last isomorphisms are identical when restricted to V_2.

Let us exlain why we introduce strongly isomorphic graphs. Having defined the graph $K^{(d,Z,L,F)}$, we aim to define a boundaried structure that we will associate with a representative of $\beta|_{\mathsf{stell}}$. The boundary of our boundaried structure will be the set $\partial_{\mathfrak{R}}(Z) \cup V(F)$. By definition, the graph induced by this set has an obvious nice 3-partition (since there is no edge between F' and $\partial_{\mathfrak{R}}(Z)$). The information we want to store is not just a boundary but the "inner-structure" of this boundary, which is mirrored by the nice 3-partition. We demand this "stronger" notion of isomorphism to be able to find another boundaried structure that corresponds to the same representative of β and still its boundary is "nicely 3-partitioned" in the same way as the boundary of the initial boundaried structure, since (as we will see later in the course of the proof) the set $V(F)$ remains "invariant" no matter of which flatness pair the extended compass we consider—thus, our isomorphism needs to keep $V(F)$ "intact."

The Out-Signature of an Extended Compass. We now define the out-signature of an extended compass. First, we encode all possible sets $F \in \mathcal{F}^{V_L(\mathbf{a})}_{i-|\partial_{\mathfrak{R}}(Z)|}$, where $i \in [0, \mathbf{tw}(\theta) - 1]$. We also consider some set $S \subseteq Z$ such that $\partial_{\mathfrak{R}}(Z) \subseteq S$ (this set corresponds to the part of X that is contained in Z). Then, we consider all representatives $\bar{\varphi}$ of $\beta|_{\mathsf{stell}}$ (recall that $\beta|_{\mathsf{stell}}$ is a formula on annotated graphs) such that when extending $\partial_{\mathfrak{R}}(Z) \cup V_L(\mathbf{a})$ to $\partial_{\mathfrak{R}}(Z) \cup V(F)$, the boundaried structure obtained from $(G^{(d,Z,L,F)}, S \cup V(F))$ after considering $\partial_{\mathfrak{R}}(Z) \cup V(F)$ as its boundary, satisfies $\bar{\varphi}$.

We set $\tau' := \{\mathsf{E}, \mathsf{X}\}$ and, for every $\ell \in [0, \mathbf{tw}(\theta) - 1]$, following Proposition 11, we consider the collection $\mathsf{rep}^{(\ell)}_{\tau'}(\beta|_{\mathsf{stell}_\mathsf{X}})$ of sentences on ℓ-boundaried τ'-structures that are "representatives" of the sentence $\beta|_{\mathsf{stell}_\mathsf{X}}$ (that is a sentence in $\mathsf{CMSO}[\tau']$). We set

$$\mathsf{SIG}_{\mathsf{out}} := \{(\mathbf{H}, \bar{\varphi}) \mid \exists \ell \in [0, \mathbf{tw}(\theta) - 1] \text{such that } \mathbf{H} \in \mathcal{H}^{(\ell)} \text{ and } \bar{\varphi} \in \mathsf{rep}^{(\ell)}_{\tau'}(\beta|_{\mathsf{stell}_\mathsf{X}})\}.$$

Let $\mathfrak{R} = (G[V(K^{\mathbf{a}})], \mathbf{a}, \mathbf{I})$ be the extended compass of a flatness pair $(W, \mathfrak{R})$ of $G \setminus V(\mathbf{a})$ of height $2w + j$, $R \subseteq V(K^{\mathbf{a}})$, $d \in [r, w]$, $L \subseteq [l]$, $Z \subseteq I^{(d-r+1)}$, and $S \subseteq Z$ such that $\partial_{\mathfrak{R}}(Z) \subseteq S$.

We define

$$\begin{aligned}\mathsf{out\text{-}sig}(\mathfrak{R}, R, d, L, Z, S) = \{(\mathbf{H}, \bar{\varphi}) \in \mathsf{SIG}_{\mathsf{out}} \quad | \quad & \exists\, F \in \mathcal{F}^{V_L(\mathbf{a})}_{|V(H)|-|\partial_{\mathfrak{R}}(Z)|}, \text{ such that if } \mathbf{H} = (H, \mathcal{U}) \\ & \quad \text{and } \mathcal{V} = (\partial_{\mathfrak{R}}(Z), V_L(\mathbf{a}), V(F) \setminus V_L(\mathbf{a})), \text{ then} \\ & \quad \mathcal{V} \text{ is a nice 3-partition of } K^{\mathbf{a}}[\partial_{\mathfrak{R}}(Z) \cup V_L(\mathbf{a})] \cup F \\ & \quad \text{and } K^{\mathbf{a}}[\partial_{\mathfrak{R}}(Z) \cup V_L(\mathbf{a})] \cup F \text{ is strongly isomorphic} \\ & \quad \text{to } H \text{ with respect to } (\mathcal{V}, \mathcal{U}), \text{ and} \\ & \exists \text{ an ordering } \mathbf{b} \text{ of } \partial_{\mathfrak{R}}(Z) \cup V(F) \text{ such that} \\ & \quad (G^{(d,Z,L,F)}, S \cup V(F), \mathbf{b}) \models \bar{\varphi}\}.\end{aligned}$$

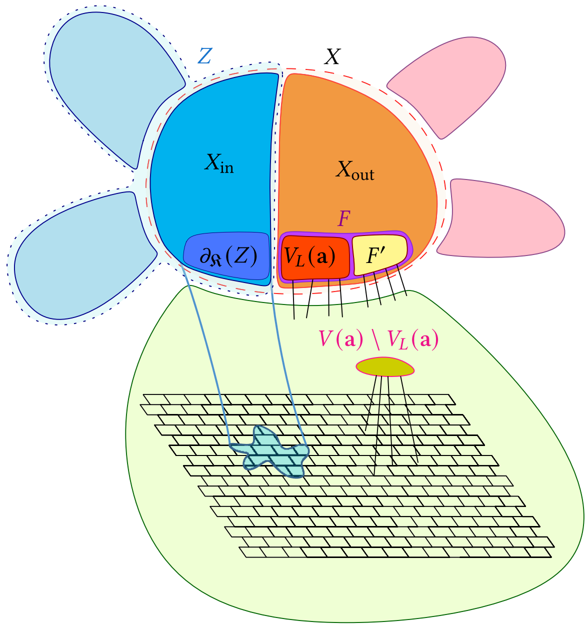

Fig. 8. The set X and the connected components of $G \setminus X$.

Intuitively, the set S is the portion of the solution that will be part of $I^{(d)}$. Also, for each $\mathbf{H} \in \mathcal{H}^{(\ell)}$, where $\ell \in [0, \mathbf{tw}(\theta) - 1]$ and $\mathbf{H}$ is a graph H together with a nice 3-partition, and each $\bar{\varphi} \in \mathsf{rep}^{(\ell)}_{\tau'}(\beta|_{\mathsf{stell}})$, we are asked to guess two objects: a graph $F \in \mathcal{F}^{V_L(\mathbf{a})}_{|V(H)|-|\partial_{\mathfrak{R}}(Z)|}$ and an ordering $\mathbf{b}$ of $\partial_{\mathfrak{R}}(Z) \cup V(F)$. The guessed additional part F' of F represents the boundary of X_{out} that is the portion of the modulator that will be away from $I^{(d)}$. The set $\partial_{\mathfrak{R}}(Z) \cup V(F)$ is the boundary of the boundaried structure $\big(G^{(d,Z,L,F)}, S \cup V(F), \mathbf{b}\big)$. This $|\partial_{\mathfrak{R}}(Z) \cup V(F)|$-boundaried structure should be a model of $\bar{\varphi}$ and its boundary (that is the union of F and $K^{\mathbf{a}}[\partial_{\mathfrak{R}}(Z) \cup V_L(\mathbf{a})]$) should be isomorphic to H. The ordering $\mathbf{b}$ of the boundary is guessed. See Figures 8 and 9 for the situation of these sets inside the d-layer. Also, keep in mind that, since $\beta|_{\mathsf{stell}} \in \mathsf{CMSO}[\{\mathsf{E}, \mathsf{X}\}]$ and $\bar{\varphi} \in \mathsf{rep}^{(\ell)}_{\tau'}(\beta|_{\mathsf{stell}})$, we have that $\bar{\varphi} \in \mathsf{CMSO}[\{\mathsf{E}, \mathsf{X}\} \cup \{\mathsf{b}_1, \ldots, \mathsf{b}_\ell\}]$, where $\mathsf{b}_1, \ldots, \mathsf{b}_\ell$ are constant symbols. When asking whether $\big(G^{(d,Z,L,F)}, S \cup V(F), \mathbf{b}\big) \models \bar{\varphi}$, we interpret $\mathsf{b}_1, \ldots, \mathsf{b}_\ell$ by $\mathbf{b}$.

In the proof of Lemma 16, we will find two extended compasses $(\mathfrak{R}, R)$, $(\mathfrak{R}', R')$ with the same out-sig for a particular choice of d and L and some choices Z and Z', respectively. In the proof, Z and Z' will be exchanged. Here it is important to notice that the graph F is always the same (for both $(\mathfrak{R}, R)$ and $(\mathfrak{R}', R')$) and constitutes the fictitious "invariant" part of the graph, that is not affected during this exchange. See Figure 8 for the great picture—what is F will not be exchanged, while $\partial_{\mathfrak{R}}(Z)$ will be substituted by the isomorphic $\partial_{\mathfrak{R}'}(Z')$ (see also Figure 15).

7.3 In-Signature

Recall that $\breve{\zeta}_{\mathsf{R}}$ is a Gaifman sentence in $\mathsf{FO}[\{\mathsf{E}\}^{\langle \mathsf{c} \rangle} \cup \{\mathsf{R}\}]$. Thus, there are $p, \ell_1, \ldots, \ell_p, r_1, \ldots, r_p \in \mathbb{N}_{\geq 1}$, and sentences $\tilde{\zeta}_1, \ldots, \tilde{\zeta}_p \in \mathsf{FO}[\{\mathsf{E}\}^{\langle \mathsf{c} \rangle} \cup \{\mathsf{R}\}]$ such that $\breve{\zeta}_{\mathsf{R}}$ is a Boolean combination of $\tilde{\zeta}_1, \ldots, \tilde{\zeta}_p$ and for every $h \in [p]$, $\tilde{\zeta}_h$ is a basic local sentence with parameters ℓ_h and r_h, i.e.,

$$\tilde{\zeta}_h = \exists \mathsf{x}_1 \ldots \exists \mathsf{x}_{\ell_h} \left(\bigwedge_{i \in [\ell_h]} \mathsf{x}_i \in \mathsf{R} \wedge \bigwedge_{1 \leq i < j \leq \ell_h} d(\mathsf{x}_i, \mathsf{x}_j) > 2r_h \wedge \bigwedge_{i \in [\ell_h]} \psi_h(\mathsf{x}_i) \right),$$

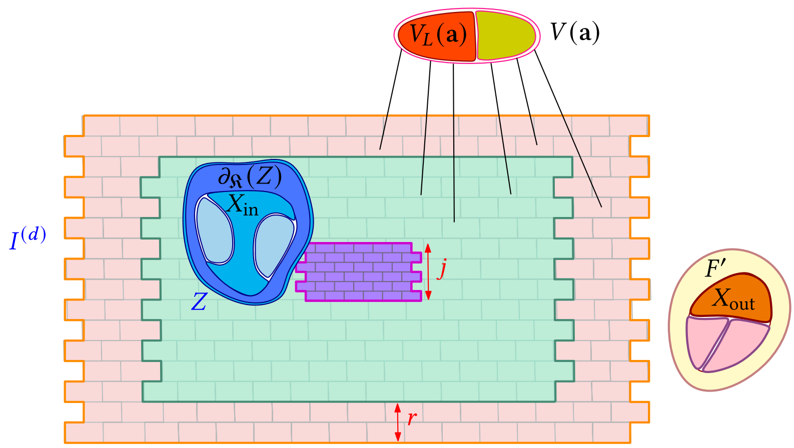

Fig. 9. An example of a set $I^{(d)}$ inside the extended compass of a flatness pair of a given graph and the position of $Z, X_{\text{in}}, X_{\text{out}}, V_L(\mathbf{a})$, and F'.

where ψ_h is an r_h-local formula in $\mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c} \rangle}]$ with one free variable. Keep in mind that, since $\check{\zeta}_{\mathsf{R}} \in \mathsf{FO}[\{\mathsf{E}\}^{\langle \mathbf{c} \rangle} \cup \{\mathsf{R}\}]$, distances are measured in the Gaifman graph of $\{\mathsf{E}\}^{\langle \mathbf{c} \rangle}$-structures.

We set $\hat{r} := \max_{h \in [p]}\{r_h\}$, $\hat{\ell} := \max_{h \in [p]}\{\ell_h\}$, and $r := 2 \cdot (\hat{\ell}+3) \cdot \hat{r}$. As in the previous subsection, we set $q = (\mathbf{tw}(\theta) + 1)^2 + 1$ and $w = (r+2) \cdot q$. The reason that r is set to be equal to $2 \cdot (\hat{\ell}+3) \cdot \hat{r}$ will be clear in the proof of Lemma 16 and is based on an idea already present in [45].

Scattered Sets in Structures. Let $\mathfrak{A}$ be a τ-structure and let $X \subseteq V(\mathfrak{A})$. We say that X is (ℓ, r)-*scattered in* $\mathfrak{A}$, if $|X| = \ell$ and for every two distinct vertices in X, their distance in the Gaifman graph $G_{\mathfrak{A}}$ is more than $2r$, i.e., for every $a, b \in X$, $a \neq b$, it holds that $d_{\mathfrak{A}}(a, b) > 2r$.

The In-Signature of an Extended Compass. We now define the in-signature of an extended compass. In this, using the approach of [45], we encode all (partial) sets of variables, one set for each basic local sentence of the Gaifman sentence $\check{\zeta}_{\mathsf{R}}$, such that these variables are lying inside an "inner part" of the compass, they are scattered in this inner part, and they satisfy the local formulas ψ_i. These arguments are always applied in some $\{\mathsf{E}\}^{\langle \mathbf{c} \rangle}$-structure of the form $\mathbf{ap}_{\mathbf{c}}(G, \mathbf{a})$. We define

$$\mathsf{SIG}_{\mathsf{in}} = 2^{[\ell_1]} \times \cdots \times 2^{[\ell_p]} \times [w].$$

Let $\mathfrak{R} = (G[V(K^{\mathbf{a}})], K^{\mathbf{a}}, \mathbf{a}, \mathbf{I})$ be an extended compass of the flatness pair $(W, \mathfrak{R})$ of $G \setminus V(\mathbf{a})$ of height $2w + j$, $R \subseteq V(K^{\mathbf{a}})$, $d \in [r, w]$, $L \subseteq [l]$, $Z \subseteq I^{(d-r+1)}$, and $S \subseteq Z$ such that $\partial_{\mathfrak{R}}(Z) \subseteq S$. We set

$$\begin{aligned}\mathsf{in\text{-}sig}(\mathfrak{R}, R, d, L, Z, S) := \{(Y_1, \ldots, Y_p, t) \in \mathsf{SIG}_{\mathsf{in}} \mid t \leq d \text{ and } \exists(\tilde{X}_1, \ldots, \tilde{X}_p) \text{ such that } \forall h \in [p] \\ &\tilde{X}_h = \{x_i^h \mid i \in Y_h\}, \\ &\tilde{X}_h \subseteq (I^{(t-\hat{r}+1)} \setminus Z) \cap R, \text{ and} \\ &\text{if } \mathbf{a}' = \mathbf{a} \setminus V_L(\mathbf{a}), \text{ then} \\ &\tilde{X}_h \text{ is } (|Y_h|, r_h)\text{-scattered in } \mathbf{ap}_{\mathbf{c}}(G, \mathbf{a}')[I^{(t)} \setminus S] \\ &\text{and } \mathbf{ap}_{\mathbf{c}}(G, \mathbf{a}')[I^{(d)} \setminus S] \models \bigwedge_{x \in \tilde{X}_h} \psi_h(x)\}.\end{aligned}$$

Notice that $I^{(d)} \setminus S$ is a more restricted part of the component of $G \setminus S$ that contains $I^{(d)} \setminus I^{(d-r)}$, but it is also "flat." Then, we guess how the scattered sets of each of the basic local sentences of

Table 3. List of Formulas Used in the Proof of Lemma 16 for Sentences in $\mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ with Their Respective Meanings

Formulas	Meaning
β	modulator sentence expressing CMSO-property on bounded treewidth structures
σ	target FO-sentence
θ	$\exists \mathsf{X}\ \beta\vert_{\mathsf{stell}_\mathsf{X}} \wedge \sigma\vert_{\mathsf{rm}_\mathsf{X}}$ or, alternatively, $\beta \triangleright \sigma$
ζ	the l-apex-projected sentence σ^l of σ
$\breve{\zeta}$	a Gaifman sentence equivalent to ζ
ψ_h	r-local formulas of the basic local sentences of $\breve{\zeta}$
$\breve{\zeta}_\mathsf{R}$	the Gaifman sentence $\breve{\zeta}$ after adding R (whose model is of the form $\mathsf{ap}_\mathbf{c}(G, R)$)
$\breve{\zeta}_\mathsf{R}\vert_{\mathsf{ap}_\mathbf{c}}$	the "backwards translation" of $\breve{\zeta}_\mathsf{R}$ to structures without "projecting" $\mathbf{c}$
$\theta_{\mathsf{R},\mathbf{c}}$	the sentence $\beta \triangleright \breve{\zeta}_\mathsf{R}\vert_{\mathsf{ap}_\mathbf{c}}$
$\bar{\varphi}$	a representative of $\beta\vert_{\mathsf{stell}_\mathsf{X}}$ given by Courcelle's theorem
ψ_h	r-local formulas of the basic local sentences of the Gaifman sentence $\breve{\zeta}$

the Gaifman sentence can intersect this graph (a buffer that "crops" the area that contains the vertices that intersect an inner-area of C corresponds to t, and the numbers of the selected vertices correspond to the sets $Y_1, \dots, Y_p$) and how these variables satisfy the r_i-local formulas ψ_i. The "scatteredness" and the satisfaction of ψ_i are evaluated on the structure after "projecting" with respect to the apices.

We finally define

$$\mathsf{CHAR} = [r, w] \times 2^{[l]} \times 2^{\mathsf{SIG}_{\mathsf{out}}} \times 2^{\mathsf{SIG}_{\mathsf{in}}}$$

and

$$\begin{aligned}\theta\text{-}\mathsf{char}(\mathfrak{R}, R) = \{(d, L, \mathsf{sig}_{\mathsf{out}}, \mathsf{sig}_{\mathsf{in}}) \in \mathsf{CHAR} \mid \exists\, Z \subseteq I^{(d-r+1)} &\text{ and } \exists\, S \subseteq Z \text{ such that} \\ &\partial_{\mathfrak{R}}(Z) \subseteq S, \\ &\mathsf{out\text{-}sig}(\mathfrak{R}, R, d, L, Z, S) = \mathsf{sig}_{\mathsf{out}}, \text{ and} \\ &\mathsf{in\text{-}sig}(\mathfrak{R}, R, d, L, Z, S) = \mathsf{sig}_{\mathsf{in}}\}.\end{aligned} \tag{3}$$

Observe that $|\mathsf{CHAR}| = \mathcal{O}_{|\theta|,l,q,j'}(1)$.

7.4 An Algorithm for Finding Equivalent Flatness Pairs

In this subsection we present an algorithm `Find_Equiv_FlatPairs` that will serve as the algorithm for Lemma 16. Given the inputs in Lemma 16, the algorithm `Find_Equiv_FlatPairs` will return, in linear time, a set $Y \subseteq V(G) \setminus V(\mathbf{a})$ and a vertex $v \in Y$, where G is the input graph and $\mathbf{a}$ is an apex-tuple of G, with the property that $(G, R, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}} \Leftrightarrow (G \setminus v, R \setminus Y, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}}$. The proof of correctness of the algorithm `Find_Equiv_FlatPairs` will prove Lemma 16 and it is in Section 7.5.

Before presenting the the algorithm `Find_Equiv_FlatPairs`, we present Table 3 that summarizes all different formulas that we consider up to this point, with their corresponding meanings.

The Algorithm `Find_Equiv_FlatPairs`. The algorithm has four steps. First, recall that there exist $p \in \mathbb{N}_{\geq 1}$, $\ell_1, \dots, \ell_p, r_1, \dots, r_p \in \mathbb{N}_{\geq 1}$, and sentences $\tilde{\zeta}_1, \dots, \tilde{\zeta}_p \in \mathsf{FO}[\tau^{\langle \mathbf{c} \rangle} \cup \{\mathsf{R}\}]$ such that $\breve{\zeta}_\mathsf{R}$ is a Boolean combination of $\tilde{\zeta}_1, \dots, \tilde{\zeta}_p$ and for every $h \in [p]$, $\tilde{\zeta}_h$ is a basic local sentence with parameters ℓ_h and r_h, i.e.,

$$\tilde{\zeta}_h = \exists \mathsf{x}_1 \ldots \exists \mathsf{x}_{\ell_h} \left(\bigwedge_{i \in [\ell_h]} \mathsf{x}_i \in \mathsf{R} \wedge \bigwedge_{1 \leq i < j \leq \ell_h} d(\mathsf{x}_i, \mathsf{x}_j) > 2r_h \wedge \bigwedge_{i \in [\ell_h]} \psi_h(\mathsf{x}_i) \right),$$

where ψ_h is an r_h-local formula in $\mathsf{FO}[\tau^{\langle \mathrm{c} \rangle}]$ with one free variable.

Let $\hat{r} := \max_{h \in [p]}\{r_h\}$ and $\hat{\ell} := \max_{h \in [p]}\{\ell_h\}$. We set c to be the size of the FO-target sentence σ of θ,

$$\begin{aligned} q &:= (\mathbf{tw}(\theta) + 1)^2 + 1, \\ f_4(\mathbf{tw}(\theta)) &:= \max\{q, 2\}, \\ z &:= 2\hat{r} + 3, \\ r &:= 2 \cdot (\hat{\ell} + 3) \cdot \hat{r}, \\ w &:= (r+2) \cdot q, \\ m &:= 2^{|\mathsf{CHAR}|} \cdot q \cdot (\hat{\ell} + 3), \text{ and} \\ f_3(\mathbf{hw}(\theta), \mathbf{tw}(\theta), c, l, z) &:= \lceil (2w + z) \cdot \sqrt{m} \rceil. \end{aligned}$$

Step 1. We first find a "packing" of subwalls of W, i.e., a collection $\mathcal{W}$ of m $(2w + j)$-subwalls of W such that their influences are pairwise disjoint. This collection exists because W has height at least $f_3(\mathbf{hw}(\theta), \mathbf{tw}(\theta), |\sigma|, l, z) = \lceil (2w + j) \cdot \sqrt{m} \rceil$ and because, due to Observation 21, for every distinct $W_i, W_j \in \mathcal{W}$, there are no cells of $\mathfrak{R}$ that are both W_i-perimetric and W_j-perimetric. Observe that the collection $\mathcal{W}$ can be computed in linear time.

Step 2. For every wall $W_i \in \mathcal{W}$, we compute a W_i-tilt of $(W, \mathfrak{R})$, which we denote by $(\tilde{W}_i, \tilde{\mathfrak{R}}_i)$, and we consider the collection $\tilde{\mathcal{W}} := \{(\tilde{W}_i, \tilde{\mathfrak{R}}_i) \mid W_i \in \mathcal{W}\}$ of m flatness pairs of $G \setminus V(\mathbf{a})$ of height $2w + z$. Note that $\tilde{\mathcal{W}}$ can be computed in time $\mathcal{O}(n)$, due to Proposition 24.

Step 3. For every $i \in [m]$, let $K_i := \mathsf{compass}_{\tilde{\mathfrak{R}}_i}(\tilde{W}_i)$ and $K_i^{\mathbf{a}} := G[V(\mathbf{a}) \cup V(K_i)]$. Also, for every $d \in [w]$, let $I_i^{(d)} := V(\bigcup \mathsf{influence}_{\tilde{\mathfrak{R}}_i}(W_i^{(2d+z)}))$ and let $\mathbf{I}_i := (I_i^{(1)}, \ldots, I_i^{(w)})$. Let $\mathfrak{K}_i := (G[V(K_i^{\mathbf{a}})], \mathbf{a}, \mathbf{I}_i)$ be the extended compass of $(\tilde{W}_i, \tilde{\mathfrak{R}}_i)$ in $G \setminus V(\mathbf{a})$, $R_i := R \cap V(K_i^{\mathbf{a}})$, and observe that for every $i, j \in [m]$, $R_i \cap R_j = R \cap V(\mathbf{a})$. After defining the above collection $\{(\mathfrak{K}_1, R_1), \ldots, (\mathfrak{K}_m, R_m)\}$ of extended compasses of flatness pairs of $G \setminus V(\mathbf{a})$, we compute their characteristics: Since, by the hypothesis of the lemma, $K_i, i \in [m]$ has treewidth at most t, by Courcelle's theorem (Proposition 1), θ-$\mathsf{char}(\mathfrak{K}_i, R_i)$ can be computed in time $\mathcal{O}_{|\theta|}(n)$. We say that two flatness pairs $(\tilde{W}_i, \tilde{\mathfrak{R}}_i), (\tilde{W}_j, \tilde{\mathfrak{R}}_j) \in \tilde{\mathcal{W}}$ are *θ-equivalent* if θ-$\mathsf{char}(\mathfrak{K}_i, R_i) = \theta$-$\mathsf{char}\ (\mathfrak{K}_j, R_j)$.

Step 4. Since $m = 2^{|\mathsf{CHAR}|} \cdot q \cdot (\hat{\ell} + 3)$ and for every $i \in [m]$, θ-$\mathsf{char}(\mathfrak{K}_i, R_i) \subseteq \mathsf{CHAR}$, we can find a collection $\tilde{\mathcal{W}}' \subseteq \tilde{\mathcal{W}}$ of pairwise θ-equivalent flatness pairs such that $|\tilde{\mathcal{W}}'| = q \cdot (\hat{\ell} + 3)$. Without loss of generality, we assume that $(\tilde{W}_1, \tilde{\mathfrak{R}}_1) \in \tilde{\mathcal{W}}'$. We set $\breve{W}$ to be the central z-subwall of $\tilde{W}_1$ and keep in mind that $z = 2\hat{r} + 2$. Note that $\breve{W}$ is also the central z-subwall of W_1 and, therefore, it is a subwall of W of height z. Again, using Proposition 24, we compute, in time $\mathcal{O}(n)$, a $\breve{W}$-tilt $(\breve{W}', \breve{\mathfrak{R}}')$ of $(W, \mathfrak{R})$ and a central vertex v of W_1. We set $Y := V(\mathsf{compass}(\breve{W}', \breve{\mathfrak{R}}'))$. We output the set Y and the vertex v.

Observe that the overall algorithm runs in linear time.

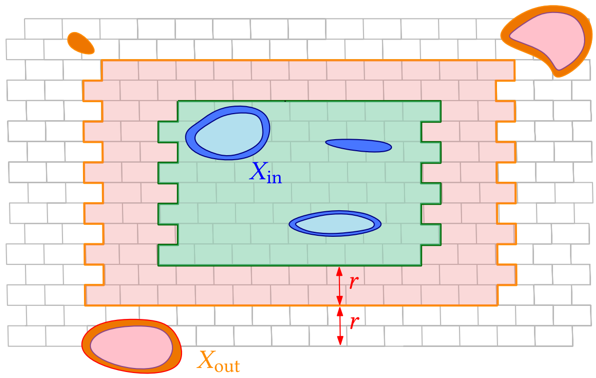

Fig. 10. The wall $\tilde{W}_1$ together with some "zones" of r consecutive layers. The area bounded by the orange layer corresponds to the set $I_1^{(i\cdot r)}$, while the area bounded by the green layer corresponds to the set $I_1^{(i\cdot r-r+1)}$. The sets X_{in} and X_{out} are depicted with blue and orange, repectively. With light blue (resp. pink) we depict the "non-privileged" connected components of $G \setminus X$ that are adjacent to vertices of X_{in} (resp. X_{out}).

7.5 Proof of Correctness of the Algorithm

In order to complete the proof of Lemma 16 for a sentence $\theta \in \mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$, we have to prove that $(G, R, \mathbf{a}) \models \theta_{\mathrm{R},\mathbf{c}} \Leftrightarrow (G \setminus v, R \setminus Y, \mathbf{a}) \models \theta_{\mathrm{R},\mathbf{c}}$. Suppose that $(G, R, \mathbf{a}) \models \theta_{\mathrm{R},\mathbf{c}}$. This means that there exists a set $X \subseteq V(G)$ such that $(\mathsf{stell}(G, X), X) \models \beta$ and $(G, R, \mathbf{a})[V(G) \setminus X] \models \breve{\zeta}_{\mathrm{R}}|_{\mathsf{ap}_{\mathbf{c}}}$.

Observations on the Collection $\tilde{\mathcal{W}}'$. Recall that for every $(\tilde{W}_i, \tilde{\mathfrak{R}}_i) \in \tilde{\mathcal{W}}'$, K_i denotes the graph $\mathsf{compass}_{\tilde{\mathfrak{R}}_i}(\tilde{W}_i)$. Also, recall that for every $i, j \in [m], i \neq j$, the walls W_i and W_j in $\mathcal{W}$ have disjoint influences. This implies that $V(K_i) \cap V(K_j) = \emptyset$. Moreover, observe that if $\tilde{Q}$ is a $(W, \mathfrak{R})$-canonical partition of $G \setminus V(\mathbf{a})$, then no internal bag of $\tilde{Q}$ intersects both $V(\bigcup \mathsf{influence}_{\mathfrak{R}}(W_i))$ and $V(\bigcup \mathsf{influence}_{\mathfrak{R}}(W_j))$, for every $i, j \in [m], i \neq j$.

Finding a θ-Equivalent Extended Compass That Is Disjoint from X. Recall that $\tilde{\mathcal{W}}'$ is a collection of $q \cdot (\hat{\ell} + 3)$ flatness pairs of $G \setminus V(\mathbf{a})$ of height $2w + z$ that are θ-equivalent to $(\tilde{W}_1, \tilde{\mathfrak{R}}_1)$. The fact that $(\mathsf{stell}(G, X), X) \models \beta$ and $\beta \in \mathsf{CMSO}^{\mathbf{tw}}$ implies that $\mathsf{torso}(G, X)$ has treewidth at most $\mathbf{tw}(\theta)$. Therefore, by Lemma 13, X intersects at most $(\mathbf{tw}(\theta) + 1)^2 = q - 1$ internal bags of every $(W, \mathfrak{R})$-canonical partition of $G \setminus V(\mathbf{a})$. This, together with the fact that $|\tilde{\mathcal{W}}'| = q \cdot (\hat{\ell} + 3)$ and that, if $\tilde{Q}$ is a $(W, \mathfrak{R})$-canonical partition of $G \setminus V(\mathbf{a})$, then no internal bag of $\tilde{Q}$ intersects both the vertex set of the influence of W_i and of the influence of W_j, for every $i, j \in [m], i \neq j$, implies that there is a collection $\tilde{\mathcal{W}}'' \subseteq \tilde{\mathcal{W}}'$ of size $\hat{\ell} + 2$ such that $(\tilde{W}_1, \tilde{\mathfrak{R}}_1) \notin \tilde{\mathcal{W}}''$, every flatness pair in $\tilde{\mathcal{W}}''$ is θ-equivalent to $(\tilde{W}_1, \tilde{\mathfrak{R}}_1)$, and the vertex set of its influence is disjoint from X. Assume, without loss of generality, that $(\tilde{W}_2, \tilde{\mathfrak{R}}_2) \in \tilde{\mathcal{W}}''$, which implies that $\theta\text{-char}(\mathfrak{K}_1, R_1) = \theta\text{-char}\ (\mathfrak{K}_2, R_2)$ and $I_2^{(w)} \cap X = \emptyset$.

Every Modulator Leaves an Intact Buffer. We fix some $(\tilde{W}_1, \tilde{\mathfrak{R}}_1)$-canonical partition $\tilde{Q}$ of $G \setminus V(\mathbf{a})$. By Lemma 13, X intersects at most $q - 1$ bags of $\tilde{Q}$. This implies that, given that $\tilde{W}_1$ has height $2w + z$ and $w = (r + 2) \cdot q$, there is an $i \in [q]$ such that $X \cap (I_1^{(i\cdot r-1)} \setminus I_1^{(i\cdot r-r)}) = \emptyset$. Let $X_{\text{in}} = X \cap I_1^{(i\cdot r-r)}$ and $X_{\text{out}} = X \setminus I_1^{(i\cdot r-1)}$ (see Figure 10 for a visualization of an example). We set $d := i \cdot r - 1$.

We remind the reader that, to prove our lemma, our objective is to show that we can replace X_{in} by another set X' that is "away" from a central part Y of W_1 and therefore "away" from any central vertex v of W_1, so that $X_{\text{out}} \cup X'$ is also a "solution," i.e., $(\mathsf{stell}(G, X_{\text{out}} \cup X'), X_{\text{in}} \cup X') \models \beta$ and

$\mathsf{ap}_{\mathbf{c}}(G,R,\mathbf{a})[V(G)\setminus(X_{\text{out}}\cup X')]\models\check{\zeta}_{\mathsf{R}}$. This will allow us to argue that $(G,R,\mathbf{a})$ and $(G\setminus v, R\setminus Y,\mathbf{a})$ are equivalent with respect to the satisfaction of $\theta_{\mathsf{R},\mathbf{c}}$. To do this, we will use the equality of characteristics of the extended compasses of the above collection of flatness pairs. Before this, we need to prove that $X\cap V(\mathbf{a})\subseteq\partial_G(X)$. This will allow us to keep the part of $\mathbf{a}$ that intersects X "intact," meaning that the same vertices will be present also in $X_{\text{out}}\cup X'$ and therefore the apex set will not change after replacing X by $X_{\text{out}}\cup X'$.

All Apices in X Are Adjacent to $V(G)\setminus X$. We set $V_L(\mathbf{a})=X\cap V(\mathbf{a})$ and L be the set of indices of the vertices of $\mathbf{a}$ in X. We claim that $V_L(\mathbf{a})\subseteq\partial_G(X)$. More generally, we show that for every set S that intersects at most $q-1$ bags of $\tilde{\mathcal{Q}}$, $V_L(\mathbf{a})\subseteq\partial_G(S)$.

Let $S\subseteq V(G)$ be a set that intersects at most $q-1$ bags of $\tilde{\mathcal{Q}}$. By assumption, every vertex in $V(\mathbf{a})$ is adjacent, in G, to at least q internal bags of $\tilde{\mathcal{Q}}$. Therefore, for every $a\in V(\mathbf{a})$, there is an internal bag Q of $\tilde{\mathcal{Q}}$ that is disjoint from S and to which a is adjacent, i.e., $V(Q)\subseteq V(G\setminus S)$ and a is adjacent, in G, to a vertex in $V(Q)$. This implies that every $a\in V(\mathbf{a})$ is either in $\partial_G(S)$ or belongs to $V(G)\setminus S$. Therefore, $V_L(\mathbf{a})\subseteq\partial_G(S)$.

Defining the Set Z. Let $\check{C}$ be the connected component of $G\setminus X$ that contains $I_1^{(d)}\setminus I_1^{(d-r+1)}$. Let $Z=I_1^{(d-r+1)}\setminus\check{C}$ and observe that $\partial_{\mathfrak{R}_1}(Z)\subseteq X_{\text{in}}$ (in Figure 10, Z corresponds to the union of the set X_{in} and all connected components of $G\setminus X$ that are depicted in yellow). Note that Z is the union of X_{in} and of every $C\in\mathsf{cc}(G,X)\setminus\{\check{C}\}$ that contains a vertex that is adjacent to a vertex of X_{in}. Keep in mind that $Z\setminus X\subseteq V(G)\setminus X$.

The fact that θ-$\mathsf{char}(\mathfrak{R}_1,R_1)=\theta$-$\mathsf{char}\,(\mathfrak{R}_2,R_2)$ implies that there is a $Z'\subseteq I_2^{(d-r+1)}$ and a set $X'\subseteq Z'$ such that $\partial_{\mathfrak{R}_2}(Z')\subseteq X'$,

— $\mathsf{out\text{-}sig}(\mathfrak{R}_1,R_1,d,L,Z,X_{\text{in}})=\mathsf{out\text{-}sig}(\mathfrak{R}_2,R_2,d,L,Z',X')$, and
— $\mathsf{in\text{-}sig}(\mathfrak{R}_1,R_1,d,L,Z,X_{\text{in}})=\mathsf{in\text{-}sig}(\mathfrak{R}_2,R_2,d,L,Z',X')$.

Note that $Z'\cap I_1^{(w)}=\emptyset$.

We first prove the following:

Claim 1. $(\mathsf{stell}(G,X),X)\models\beta\Leftrightarrow(\mathsf{stell}(G\setminus v,X_{\text{out}}\cup X'),X_{\text{out}}\cup X')\models\beta$.

Proof of Claim 1: Let $h:=|N_G(\check{C})|$. Recall that $\mathsf{torso}(G,X)$ has treewidth at most $\mathbf{tw}(\theta)$. Since the set $N_G(\check{C})$ induces a complete graph on h vertices in the graph $\mathsf{torso}(G,X)$, we have that $h\in[0,\mathbf{tw}(\theta)-1]$.

The idea here is to build an h-boundaried annotated graph (with respect to W_1) to fit the $\mathsf{out\text{-}sig}$. By Proposition 11, this annotated graph is associated with a sentence $\bar{\varphi}\in\mathsf{rep}^{(h)}_{\{\mathsf{E},\mathsf{X}\}}(\beta|_{\mathsf{stell}_{\mathsf{X}}})$. Next, we will consider another h-boundaried annotated graph (with respect to W_2) that satisfies the same sentence, using the fact that W_1 and W_2 have the same $\mathsf{out\text{-}sig}$. The fact that both these boundaried annotated graphs, when "completed" from the other side by the same annotated graph, give the same annotated graph, will imply that they are $(\beta|_{\mathsf{stell}_{\mathsf{X}}},h)$-equivalent.

Defining the Boundary of Our Boundaried Structure. We set F' to be the graph $G[(X_{\text{out}}\setminus V_L(\mathbf{a}))\cap N_G(\check{C})]$. In other words, F' is the subgraph of G induced by the vertices of X_{out} that are not apices and are adjacent to vertices in $\check{C}$ (see Figure 8 for an example). Also, we set $F^\star$ to be the graph obtained from $G[V_L(\mathbf{a})\cup V(F')]$ after removing every edge that has both endpoints in $V_L(\mathbf{a})$. Intuitively, we extend F' to $F^\star$ by adding the vertices in $V_L(\mathbf{a})$ and the edges connecting vertices of $V_L(\mathbf{a})$ and $V(F')$, but not the edges that have both endpoints in $V_L(\mathbf{a})$. This graph $F^\star$ will be later associated with a graph $F\in\mathcal{F}^{V_L(\mathbf{a})}_{h-|\partial_{\mathfrak{R}_1}(Z)|}$.

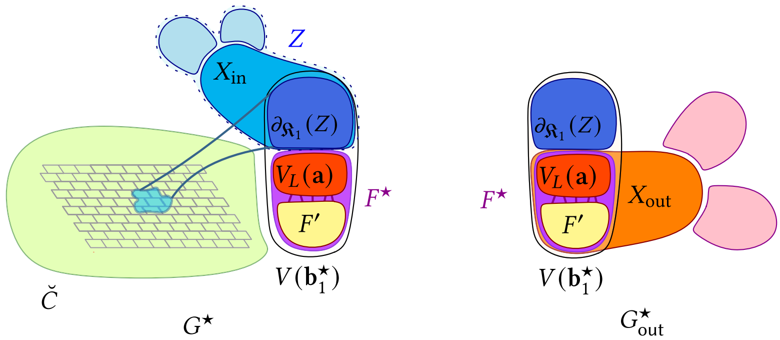

Fig. 11. The graph G from Figures 8 and 9, "separated" into two parts. Left: The graph $G^\star$ with the set $V(\mathbf{b}_1^\star)$ as boundary. Right: The graph $G^\star_{\text{out}}$ with the set $V(\mathbf{b}_1^\star)$ as boundary.

Separating G into Two Boundaried Structures. We now aim to "break" G into two boundaried structures, to be able to encode, using representatives, the "partial satisfaction" of β inside $\check{C}$ and Z. Let

$$G^\star_{\text{out}} = G \setminus (\check{C} \cup (Z \setminus \partial_{\mathfrak{R}_1}(Z))) \text{ and } G^\star = G[\check{C} \cup Z \cup V(F^\star)].$$

Keep in mind that $Z = I_1^{(d-r+1)} \setminus \check{C}$. See Figure 11 to get some intuition of the graphs $G^\star_{\text{out}}$ and $G^\star$. Verbally, the graph $G^\star_{\text{out}}$ is obtained from G by removing from its vertex set the sets $\check{C}$ and $Z \setminus \partial_{\mathfrak{R}_1}(Z)$. In other words, apart from the vertices in $\partial_{\mathfrak{R}_1}(Z) \cup V(F^\star)$ (that are in the vertex set of both $G^\star_{\text{out}}$ and $G^\star$), the graph $G^\star_{\text{out}}$ corresponds to the part of G that is "away" from $\check{C}$ and Z, while the graph $G^\star$ corresponds to the part of G induced by the union of $\check{C}$, Z and $V(F^\star)$. Keep in mind that $V(G^\star_{\text{out}}) \cap V(G^\star) = \partial_{\mathfrak{R}_1}(Z) \cup V(F^\star)$. Next, we will define two boundaried structures corresponding to $G^\star_{\text{out}}$ and $G^\star$, whose boundary will be the set $\partial_{\mathfrak{R}_1}(Z) \cup V(F^\star)$.

An Ordering on the (Common) Boundary of the Two Structures. We next claim that $\partial_{\mathfrak{R}_1}(Z) \cup V(F^\star) = N_G(\check{C})$, which directly implies that $|\partial_{\mathfrak{R}_1}(Z) \cup V(F^\star)| = h$. To see why $\partial_{\mathfrak{R}_1}(Z) \cup V(F^\star) = N_G(\check{C})$, first observe that, since $\check{C} \in \text{cc}(G, X)$, it holds that $N_G(\check{C}) \subseteq X$ and also notice that $X_{\text{in}} \cap N_G(\check{C}) = \partial_{\mathfrak{R}_1}(Z)$. Since $V_L(\mathbf{a}) = V(\mathbf{a}) \cap N_G(\check{C})$ and $V(F') = (X_{\text{out}} \setminus V_L(\mathbf{a})) \cap N_G(\check{C})$, we have that $N_G(\check{C}) = \partial_{\mathfrak{R}_1}(Z) \cup V_L(\mathbf{a}) \cup V(F') = \partial_{\mathfrak{R}_1}(Z) \cup V(F^\star)$. Therefore, $|\partial_{\mathfrak{R}_1}(Z) \cup V(F^\star)| = h$. Consider an ordering $\mathbf{b}_1^\star = (v_1, \ldots, v_h)$ of the vertices in $\partial_{\mathfrak{R}_1}(Z) \cup V(F^\star)$ and recall that $V(G^\star_{\text{out}}) \cap V(G^\star) = \partial_{\mathfrak{R}_1}(Z) \cup V(F^\star)$. Now, consider the h-boundaried graph $(G^\star_{\text{out}}, \mathbf{b}_1^\star)$ and $(G^\star, \mathbf{b}_1^\star)$. Notice that $(G^\star_{\text{out}}, \mathbf{b}_1^\star)$ and $(G^\star, \mathbf{b}_1^\star)$ are compatible and that $(G^\star_{\text{out}}, \mathbf{b}_1^\star) \oplus (G^\star, \mathbf{b}_1^\star) = G$.

Adding $V(F^\star)$ to X_{in}. Let $\tilde{X}_1^\star = X_{\text{in}} \cup V(F^\star)$. Now, the fact that $\partial_{\mathfrak{R}_1}(Z) \subseteq X_{\text{in}}$ implies that $V(\mathbf{b}_1^\star) \subseteq \tilde{X}_1^\star$. Since $V(F^\star) \subseteq X_{\text{out}}$, it holds that $X_{\text{out}} \cup \tilde{X}_1^\star = X$.

Separating (G, X) into Two Boundaried Structures. Observe that

$$(G, X) = (G^\star_{\text{out}}, X_{\text{out}}, \mathbf{b}_1^\star) \oplus (G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star). \tag{4}$$

We set $\tilde{X}_2^\star := X' \cup V(F^\star)$. Our aim is to prove that there is an ordering $\mathbf{b}_2^\star$ of $\partial_G(Z') \cup V(F^\star)$ such that

$$(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star) \text{ and } (G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star) \text{ are } (\beta|_{\mathsf{stell_X}}, h)\text{-equivalent.}$$

Let $\bar{\varphi} \in \text{rep}^{(h)}_{\{\mathsf{E},\mathsf{X}\}}(\beta|_{\mathsf{stell_X}})$ such that $\left(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star\right) \models \bar{\varphi}$.

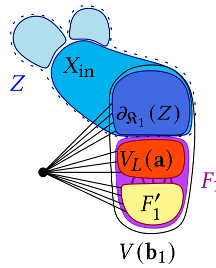

Fig. 12. The graph G_1.

Shifting from $G^\star$ to $G^{(d,Z,L,F_1)}$. Now, consider a graph $F_1 \in \mathcal{F}^{V_L(\mathbf{a})}_{h-|\partial_{\mathfrak{R}_1}(Z)|}$ that is isomorphic[5] to $F^\star$, via a bijection $\xi : V(F_1) \leftrightarrow V(F^\star)$ that maps every $a \in V_L(\mathbf{a})$ to itself. Let $F_1' := F_1 \setminus V_L(\mathbf{a})$. We set $\mathcal{V}_1 := (\partial_{\mathfrak{R}_1}(Z), V_L(\mathbf{a}), V(F_1'))$ and observe that $\mathcal{V}_1$ is a nice 3-partition of $K_1^{\mathbf{a}}[\partial_{\mathfrak{R}_1}(Z) \cup V_L(\mathbf{a})] \cup F_1$. Also, observe that the graph $V(K_1^{\mathbf{a}}[\partial_{\mathfrak{R}_1}(Z) \cup V_L(\mathbf{a})] \cup F_1)$ has h vertices and therefore $(K_1^{\mathbf{a}}[\partial_{\mathfrak{R}_1}(Z) \cup V_L(\mathbf{a})] \cup F_1, \mathcal{V}_1) \in \mathcal{H}^{(h)}$. Let $\mathbf{H} := (K_1^{\mathbf{a}}[\partial_{\mathfrak{R}_1}(Z) \cup V_L(\mathbf{a})] \cup F_1, \mathcal{V}_1)$.

A Boundaried Structure of Bounded Treewidth That Satisfies $\bar{\varphi}$. Let $\mathbf{b}_1$ be the tuple obtained from $\mathbf{b}_1^\star$ after replacing, in $\mathbf{b}_1^\star$, each vertex $v \in V(F^\star)$ with the vertex $\xi^{-1}(v) \in V(F_1)$. Also, let $G_1 = K_1^{(d,Z,L,F_1)}$ (see Figure 12 for a visualization of G_1). We set $\tilde{X}_1 := X_{\text{in}} \cup V(F_1)$. Keep in mind that $\tilde{X}_1$ is obtained from $\tilde{X}_1^\star$ after replacing $V(F^\star)$ with $V(F_1)$, i.e., $\tilde{X}_1 = (\tilde{X}_1^\star \setminus V(F^\star)) \cup V(F_1)$. Also, observe that $\tilde{X}_1 \subseteq V(G_1)$.

We aim to show that $(\mathbf{H}, \bar{\varphi}) \in$ out-sig$(\mathfrak{R}_1, R_1, d, L, Z)$. To show this, by the definition of out-sig it remains to prove that $(G_1, \tilde{X}_1, \mathbf{b}_1) \models \bar{\varphi}$. To prove the latter, first notice that, since F_1 and $F^\star$ are isomorphic, we have that $G_1[V(\mathbf{b}_1)]$, $G^\star[V(\mathbf{b}_1^\star)]$, and $G^\star_{\text{out}}[V(\mathbf{b}_1^\star)]$ are (pairwise) isomorphic. This implies that $(G_1, \mathbf{b}_1)$, $(G^\star, \mathbf{b}_1^\star)$, and $(G^\star_{\text{out}}, \mathbf{b}_1^\star)$ are (pairwise) compatible. We consider the h-boundaried annotated graphs $(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star)$ and $(G_1, \tilde{X}_1, \mathbf{b}_1)$. These h-boundaried annotated graphs are also compatible. We now show that, moreover, they are $(\beta|_{\mathsf{stell}_\mathsf{X}}, h)$-equivalent, which will imply that $(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star) \models \bar{\varphi} \Leftrightarrow (G_1, \tilde{X}_1, \mathbf{b}_1) \models \bar{\varphi}$.

Subclaim. $(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star)$ and $(G_1, \tilde{X}_1, \mathbf{b}_1)$ are $(\beta|_{\mathsf{stell}_\mathsf{X}}, h)$-equivalent.

Proof of Subclaim. Let G° be a graph, $X^\circ \subseteq V(G^\circ)$, and $\mathbf{b}^\circ$ be an apex-tuple of G° of size h, such that $(G^\circ, X^\circ, \mathbf{b}^\circ)$ is an h-boundaried annotated graph that is compatible with the h-boundaried annotated graphs $(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star)$ and $(G_1, \tilde{X}_1, \mathbf{b}_1)$. We set $D^\star := (G^\circ, \mathbf{b}^\circ) \oplus (G^\star, \mathbf{b}_1^\star)$ and $D := (G^\circ, \mathbf{b}^\circ) \oplus (G_1, \mathbf{b}_1)$. Our goal is to show that

$$(D^\star, X^\circ \cup \tilde{X}_1^\star) \models \beta|_{\mathsf{stell}_\mathsf{X}} \Leftrightarrow (D, X^\circ \cup \tilde{X}_1) \models \beta|_{\mathsf{stell}_\mathsf{X}}.$$

To prove this, we will argue that the graph $\mathsf{stell}_\mathsf{X}(D, X^\circ \cup \tilde{X}_1)$ is isomorphic to the structure $\mathsf{stell}_\mathsf{X}(D^\star, X^\circ \cup \tilde{X}_1^\star)$. To see this, let $v^\star$ be the vertex of $V(\mathsf{stell}_\mathsf{X}(D^\star, X^\circ \cup \tilde{X}_1^\star)) \setminus (X^\circ \cup \tilde{X}_1^\star)$ that corresponds to the component of $D^\star \setminus (X^\circ \cup \tilde{X}_1^\star)$ that contains $V(G^\star) \setminus (Z \cup V(F^\star))$. Also, let v' be the vertex of $V(\mathsf{stell}_\mathsf{X}(D, X^\circ \cup \tilde{X}_1)) \setminus (X^\circ \cup \tilde{X}_1)$ that corresponds to the component of $D \setminus (X^\circ \cup \tilde{X}_1)$ that contains the vertex in $V(G_1) \setminus (Z \cup V(F_1))$. Observe that $\mathsf{stell}_\mathsf{X}(D^\star, X^\circ \cup \tilde{X}_1^\star) \setminus v^\star$

[5]In the rest of the proof of the claim, we will usually consider a subgraph of G, or a structure with universe $V(G)$, and isomorphic graphs/structures of them, and the latter will be "abstract" graphs/structures. For example, here we consider an "abstract" graph F_1 that is isomorphic to the graph $F^\star$ that is a subgraph of G. We will always use superscript "$\star$" in order to denote the subgraphs/structures that are being given by the graph, while the lack of superscript reflects to the corresponding isomorphic "abstract" graphs/structures.

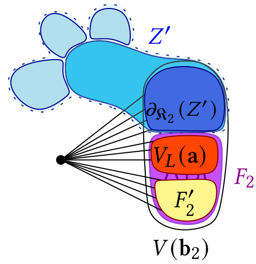

Fig. 13. The graph G_2.

and $\mathsf{stell}_{\mathsf{X}}(D, X^\circ \cup \tilde{X}_1) \setminus v'$ are isomorphic graphs, obtained by replacing $F^\star$ by F_1 (and vice-versa), which are isomorphic. Moreover, note that this isomorphism can be extended by mapping $v^\star$ to v'. Thus, $(D^\star, X^\circ \cup \tilde{X}_1^\star) \models \beta|_{\mathsf{stell}_{\mathsf{X}}} \Leftrightarrow (D, X^\circ \cup \tilde{X}_1) \models \beta|_{\mathsf{stell}_{\mathsf{X}}}$ and the subclaim follows. $\diamond$

Since by the above subclaim, $(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star)$ and $(G_1, \tilde{X}_1, \mathbf{b}_1)$ are $(\beta|_{\mathsf{stell}_{\mathsf{X}}}, h)$-equivalent,

$$(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star) \models \bar{\varphi} \Leftrightarrow (G_1, \tilde{X}_1, \mathbf{b}_1) \models \bar{\varphi}. \tag{5}$$

Therefore, F_1 and $\mathbf{b}_1$ certify that $(\mathbf{H}, \bar{\varphi}) \in \text{out-sig}(\mathfrak{R}_1, R_1, d, L, Z, X_{\text{in}})$. Equality of out-signatures implies that $(\mathbf{H}, \bar{\varphi}) \in \text{out-sig}(\mathfrak{R}_2, R_2, d, L, Z', X')$. Thus,

(a) there is an $F_2 \in \mathcal{F}^{V_L(\mathbf{a})}_{h-|\partial_{\mathfrak{R}_2}(Z')|}$ such that if $\mathbf{H} = (H, \mathcal{U})$ and $\mathcal{V}_2 = (\partial_{\mathfrak{R}_2}(Z'), V_L(\mathbf{a}), V(F_2) \setminus V_L(\mathbf{a}))$, then $\mathcal{V}_2$ is a nice 3-partition of $K_2^{\mathbf{a}}[\partial_{\mathfrak{R}_2}(Z') \cup V_L(\mathbf{a})] \cup F_2$ and $K_2^{\mathbf{a}}[\partial_{\mathfrak{R}_2}(Z') \cup V_L(\mathbf{a})] \cup F_2$ is strongly isomorphic to H with respect to $(\mathcal{V}_2, \mathcal{U})$ and
(b) there is an ordering $\mathbf{b}_2$ of $\partial_{\mathfrak{R}_2}(Z') \cup V(F_2)$ such that $(K_2^{(d,Z',L,F_2)}, X' \cup V(F_2), \mathbf{b}_2) \models \bar{\varphi}$.

Observe that, since $K_1^{\mathbf{a}}[\partial_{\mathfrak{R}_1}(Z) \cup V_L(\mathbf{a})] \cup F_1$ and $K_2^{\mathbf{a}}[\partial_{\mathfrak{R}_2}(Z') \cup V_L(\mathbf{a})] \cup F_2$ are strongly isomorphic to H with respect to $(\mathcal{V}_1, \mathcal{U})$ and $(\mathcal{V}_2, \mathcal{U})$, respectively, we also have that $K_1^{\mathbf{a}}[\partial_{\mathfrak{R}_1}(Z) \cup V_L(\mathbf{a})] \cup F_1$ is strongly isomorphic to $K_2^{\mathbf{a}}[\partial_{\mathfrak{R}_2}(Z') \cup V_L(\mathbf{a})] \cup F_2$ with respect to $(\mathcal{V}_1, \mathcal{V}_2)$. We now set $F_2' = F_2 \setminus V_L(\mathbf{a})$ and $G_2 = K_2^{(d,Z',L,F_2)}$ (see Figure 13 for a visualization of G_2).

Notice that the fact that $K_1^{\mathbf{a}}[\partial_{\mathfrak{R}_1}(Z) \cup V_L(\mathbf{a})] \cup F_1$ is strongly isomorphic to $K_2^{\mathbf{a}}[\partial_{\mathfrak{R}_2}(Z') \cup V_L(\mathbf{a})] \cup F_2$ with respect to $(\mathcal{V}_1, \mathcal{V}_2)$ implies that the h-boundaried annotated graphs $(G_2, \tilde{X}_2, \mathbf{b}_2)$ and $(G_1, \tilde{X}_1, \mathbf{b}_1)$ are compatible. Thus, given that $(G_2, \tilde{X}_2, \mathbf{b}_2) \models \bar{\varphi}$, we have that $(G_2, \tilde{X}_2, \mathbf{b}_2)$ and $(G_1, \tilde{X}_1, \mathbf{b}_1)$ are $(\beta|_{\mathsf{stell}_{\mathsf{X}}}, h)$-equivalent. Therefore,

$$(G_1, \tilde{X}_1, \mathbf{b}_1) \models \bar{\varphi} \Leftrightarrow (G_2, \tilde{X}_2, \mathbf{b}_2) \models \bar{\varphi}. \tag{6}$$

At this point, to give some intuition, we underline that even if $(G_2, \tilde{X}_2, \mathbf{b}_2)$ and $(G_1, \tilde{X}_1, \mathbf{b}_1)$ are $(\beta|_{\mathsf{stell}_{\mathsf{X}}}, h)$-equivalent, we did not yet provide a boundaried structure that is a substructure of (G, X) and that is $(\beta|_{\mathsf{stell}_{\mathsf{X}}}, h)$-equivalent to $(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star)$. To find such a substructure $(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star)$ of (G, X), we have to "shift" from $(G_2, \tilde{X}_2, \mathbf{b}_2)$ to $(G_2^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star)$, by replacing $V(F_2)$ with $V(F^\star)$. This substructure $(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star)$ will replace $(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star)$ in Equation (4), thus proving the claim.

Defining a Substructure of the Initial Structure with a Different Boundary. Let us now define this substructure $(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star)$ from $(G_2, \tilde{X}_2, \mathbf{b}_2)$. We set $\mathbf{b}_2^\star$ to be the tuple obtained from $\mathbf{b}_2$ after replacing each $v \in V(F_2)$ with the corresponding $u \in V(F^\star)$. See Figure 14 for an example. We stress that the h-boundaried graph $(G^\star, \mathbf{b}_2^\star)$ of Figure 14 can also be defined as the one obtained from $(G^\star, \mathbf{b}_1^\star)$ of Figure 11 after replacing $\partial_{\mathfrak{R}_1}(Z)$ with $\partial_{\mathfrak{R}_2}(Z')$ in the boundary. Also, let $\tilde{X}_2^\star = (\tilde{X}_2 \setminus V(F_2)) \cup V(F^\star)$.

For the tuple $(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star)$ to be an h-boundaried annotated graph, we need to show that $\tilde{X}_2^\star \subseteq V(G^\star)$ and $V(\mathbf{b}_2^\star) \subseteq V(G^\star)$. To show that $V(\mathbf{b}_2^\star) \subseteq V(G^\star)$, we first notice that $I_2^{(d)} \subseteq \check{C}$. Therefore,

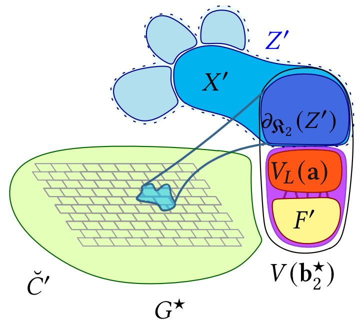

Fig. 14. The graph $G^\star$, when adding $\partial_{\mathfrak{R}_2}(Z')$ to the boundary. The set X' is the set $\tilde{X}_2 \setminus V(F^\star)$ and $\check{C}'$ is the set $V(G^\star) \setminus (Z' \cup V(F^\star))$.

since $Z' \subseteq I_2^{(d-r+1)}$, we have that $Z' \subseteq \check{C}$. The latter implies that $\partial_{\mathfrak{R}_2}(Z')$ is a subset of $V(G^\star)$. By the definition of $\mathbf{b}_2^\star$ and since $V(\mathbf{b}_2) = \partial_{\mathfrak{R}_2}(Z') \cup V(F_2)$, we have that $V(\mathbf{b}_2^\star) = \partial_{\mathfrak{R}_2}(Z') \cup V(F^\star)$. Hence, given that $\partial_{\mathfrak{R}_2}(Z') \subseteq V(G^\star)$ and $V(F^\star) \subseteq V(G^\star)$, it holds that $V(\mathbf{b}_2^\star) \subseteq V(G^\star)$. Also, observe that $\tilde{X}_2^\star \subseteq V(G^\star)$, since $\tilde{X}_2^\star = (\tilde{X}_2 \setminus V(F_2)) \cup V(F^\star)$, $\tilde{X}_2 \setminus V(F_2) \subseteq Z'$, and $Z' \subseteq V(G^\star)$.

All Considered Boundaried Structures Are $(\beta|_{\mathsf{stell}_\mathsf{X}}, h)$-Equivalent. As a next step, we argue that the h-boundaried annotated graphs $(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star)$, $(G_2, \tilde{X}_2, \mathbf{b}_2)$, and $(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star)$ are (pairwise) compatible. To see why this holds, notice that, since $K_1^{\mathbf{a}}[\partial_{\mathfrak{R}_1}(Z) \cup V_L(\mathbf{a})] \cup F_1$ is strongly isomorphic to $K_2^{\mathbf{a}}[\partial_{\mathfrak{R}_2}(Z') \cup V_L(\mathbf{a})] \cup F_2$ with respect to $(\mathcal{V}_1, \mathcal{V}_2)$, it holds that F_1 and F_2 are isomorphic. This, together with the fact that F_1 is isomorphic to $F^\star$, implies that F_2, F_1, and $F^\star$ are pairwise isomorphic graphs. Therefore, the structures $G^\star[V(\mathbf{b}_2^\star)]$, $G_2[V(\mathbf{b}_2)]$, and $G^\star[V(\mathbf{b}_1^\star)]$ are (pairwise) isomorphic. By following the exactly symmetric arguments as in the proof of the subclaim above, it is easy to show that $(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star)$ and $(G_2, \tilde{X}_2, \mathbf{b}_2)$ are $(\beta|_{\mathsf{stell}_\mathsf{X}}, h)$-equivalent. This implies that

$$\left(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star\right) \models \bar{\varphi} \Leftrightarrow \left(G_2, \tilde{X}_2, \mathbf{b}_2\right) \models \bar{\varphi}. \tag{7}$$

Another Way to Add a Boundary in the Initial Structure. Combining Equations (5), (6), and (7), we conclude that the h-boundaried annotated graphs $\left(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star\right)$ and $\left(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star\right)$ are $(\beta|_{\mathsf{stell}_\mathsf{X}}, h)$-equivalent. Recall that, by Equation (4),

$$(G^\star_{\text{out}}, X_{\text{out}}, \mathbf{b}_1^\star) \oplus (G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star) = (G, X),$$

and $\mathsf{stell}((G,X),X) \models \beta$. Since the h-boundaried annotated graphs $\left(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star\right)$ and $\left(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star\right)$ are $(\beta|_{\mathsf{stell}_\mathsf{X}}, h)$-equivalent,

$$(G^\star_{\text{out}}, X_{\text{out}}, \mathbf{b}_1^\star) \oplus \left(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star\right) \models \beta|_{\mathsf{stell}_\mathsf{X}}. \tag{8}$$

Observe that the h-boundaried annotated graphs $(G^\star_{\text{out}}, X_{\text{out}}, \mathbf{b}_1^\star)$ and $\left(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star\right)$ are compatible. Therefore, we can consider the annotated graph $(G^\star_{\text{out}}, X_{\text{out}}, \mathbf{b}_1^\star) \oplus \left(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star\right)$, which is the same as $(G, X_{\text{out}} \cup X')$. See Figure 15 for an example of how $\left(G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star\right)$ is "transformed" to $\left(G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star\right)$.

Finally, we have that $(G, X) \models \beta|_{\mathsf{stell}_\mathsf{X}} \Leftrightarrow (G, X_{\text{out}} \cup X') \models \beta|_{\mathsf{stell}_\mathsf{X}}$. To conclude the proof of Claim 1, it remains to prove that if v is a central vertex of W_1, then

$$(G, X_{\text{out}} \cup X') \models \beta|_{\mathsf{stell}_\mathsf{X}} \Leftrightarrow (G \setminus v, X_{\text{out}} \cup X') \models \beta|_{\mathsf{stell}_\mathsf{X}}.$$

Observe that since v is a central vertex of W_1 and $(X_{\text{out}} \cup X') \cap I_1^{(d)} = \emptyset$, we have that $N_G(v) \cap (X_{\text{out}} \cup X') \subseteq V(\mathbf{a})$. Also, recall that, by the hypothesis of the lemma, $N_G(V(\mathbf{a}))$ intersects at least $f_4(\mathbf{tw}(\theta)) \geq 2$ many internal bags of any $(W, \mathfrak{R})$-canonical partition of $G \setminus A$. Therefore, every

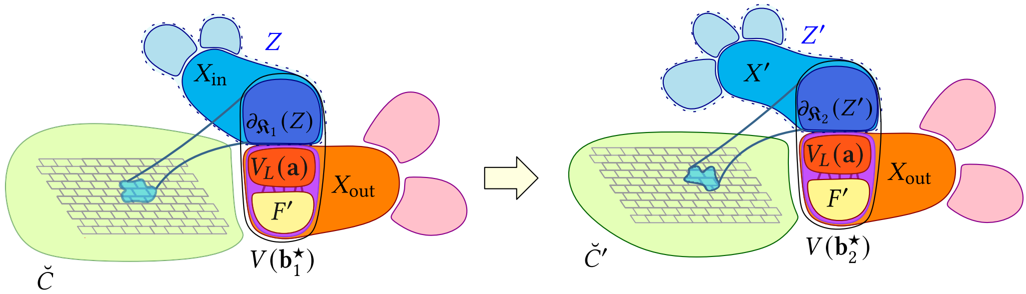

Fig. 15. On the left, the annotated graph $(G^\star_{\text{out}}, X_{\text{out}}, \mathbf{b}_1^\star) \oplus (G^\star, \tilde{X}_1^\star, \mathbf{b}_1^\star)$, while on the right, the annotated graph $(G^\star_{\text{out}}, X_{\text{out}}, \mathbf{b}_1^\star) \oplus (G^\star, \tilde{X}_2^\star, \mathbf{b}_2^\star)$.

$u \in X_{\text{out}} \cup X'$ that is adjacent to v in G, is also adjacent to a vertex in $V(\text{compass}_{\mathfrak{R}}(W))$ that is different from v. This implies that $(\text{stell}_{\mathsf{X}}(G, X_{\text{out}} \cup X'), X_{\text{out}} \cup X') \models \beta \Leftrightarrow (\text{stell}_{\mathsf{X}}(G \setminus v, X_{\text{out}} \cup X'), X_{\text{out}} \cup X') \models \beta$. Claim 1 follows. $\diamond$

Let $\check{C}' = V(G^\star) \setminus (Z' \cup V(F^\star))$ and observe that, since $I_1^{(d\cdot r)} \subseteq V(G^\star)$ and $I_1^{(d\cdot r)} \cap Z' = \emptyset$, we have that $I_1^{(d\cdot r)} \subseteq \check{C}'$. Also, since every vertex in $V(\mathbf{a})$ is adjacent, in G, to at least q internal bags of $\tilde{Q}$, we have that every $a \in V(\mathbf{a})$ is either in $V_L(\mathbf{a})$ or belongs to both $\check{C}$ and $\check{C}'$. Therefore, $\mathbf{a} \cap \check{C} = \mathbf{a} \cap \check{C}' = \mathbf{a} \cap (V(G) \setminus X) = \mathbf{a} \cap (V(G) \setminus (X_{\text{out}} \cup X'))$. We set $\mathbf{a}' := \mathbf{a} \cap \check{C}$. We aim to prove the following:

$\boxed{\text{Claim 2.}}$ $\mathit{ap}_{\mathbf{c}}((G, R, \mathbf{a})[V(G) \setminus X]) \models \check{\zeta}_R \Leftrightarrow \mathit{ap}_{\mathbf{c}}((G \setminus v, R \setminus Y, \mathbf{a})[V(G) \setminus (X_{\text{out}} \cup X')]) \models \check{\zeta}_R$.

Proof of Claim 2: We first show that $\text{ap}_{\mathbf{c}}((G, R, \mathbf{a})[V(G) \setminus X]) \models \check{\zeta}_{\mathsf{R}} \Leftrightarrow \text{ap}_{\mathbf{c}}((G, R \setminus Y, \mathbf{a})[V(G) \setminus (X_{\text{out}} \cup X')]) \models \check{\zeta}_{\mathsf{R}}$, i.e., without removing v from $V(G)$.

We set $(\mathfrak{B}, R_{\mathfrak{B}}) := \text{ap}_{\mathbf{c}}((G, R, \mathbf{a})[V(G) \setminus X])$, where $R_{\mathfrak{B}} := R \cap (V(G) \setminus X)$. Keep in mind that $\mathfrak{B}$ is a $\{\mathsf{E}\}^{\langle \mathsf{c} \rangle}$-structure. Since the Gaifman graphs of $\mathfrak{B}$ and of $(\mathfrak{B}, R_{\mathfrak{B}})$ are the same, in the rest of the proof we will use $G_{\mathfrak{B}}$ to denote both of them. Also, to get some intuition, notice that $G_{\mathfrak{B}}$ is obtained from $G \setminus X$ after removing some edges (namely, the edges of $G \setminus X$ that connect the vertices in $V(\mathbf{a}')$ with $(V(G) \setminus X) \setminus V(\mathbf{a}')$). We also set $(\mathfrak{B}', R'_{\mathfrak{B}}) := \text{ap}_{\mathbf{c}}((G, R, \mathbf{a})[V(G) \setminus (X_{\text{out}} \cup X')])$, where $R'_{\mathfrak{B}} := R \cap (V(G) \setminus (X_{\text{out}} \cup X'))$. Thus, one can rewrite the statement of the claim as $(\mathfrak{B}, R_{\mathfrak{B}}) \models \check{\zeta}_{\mathsf{R}} \implies (\mathfrak{B}', R'_{\mathfrak{B}} \setminus Y) \models \check{\zeta}_{\mathsf{R}}$.

Since $\check{\zeta}_{\mathsf{R}}$ is a Boolean combination of the basic local sentences $\check{\zeta}_1, \ldots, \check{\zeta}_p$, there is a set $J \subseteq [p]$ such that for every model $\mathfrak{C}$ of $\check{\zeta}_{\mathsf{R}}$, it holds that $\mathfrak{C} \models \check{\zeta}_j$, for every $j \in J$, while $\mathfrak{C} \models \neg\check{\zeta}_j$, for every $j \notin J$. We will show that for every $j \in J$ it holds that $(\mathfrak{B}, R_{\mathfrak{B}}) \models \check{\zeta}_j \Leftrightarrow (\mathfrak{B}', R'_{\mathfrak{B}} \setminus Y) \models \check{\zeta}_j$ and that for every $j \notin J$ it holds that $(\mathfrak{B}, R_{\mathfrak{B}}) \models \neg\check{\zeta}_j \Leftrightarrow (\mathfrak{B}', R'_{\mathfrak{B}} \setminus Y) \models \neg\check{\zeta}_j$. We proceed by distinguishing these two cases.

Case 1. $j \in J$.

We aim to prove that $(\mathfrak{B}, R_{\mathfrak{B}}) \models \check{\zeta}_j \Leftrightarrow (\mathfrak{B}', R'_{\mathfrak{B}} \setminus Y) \models \check{\zeta}_j$. First, suppose that $(\mathfrak{B}, R_{\mathfrak{B}}) \models \check{\zeta}_j$. Since $\check{\zeta}_j$ is a basic local sentence with parameters r_j and ℓ_j, we have that

$$(\mathfrak{B}, R_{\mathfrak{B}}) \models \check{\zeta}_j \Leftrightarrow \exists X_j \subseteq R_{\mathfrak{B}} \text{that is} (\ell_j, r_j)\text{-scattered in} \mathfrak{B} \text{and} \mathfrak{B} \models \bigwedge_{x \in X_j} \psi_j(x).$$

We prove the following, which intuitively states that, given the set X_j, we can find another set X'_j that "behaves" in the same way as X_j but also "avoids" some inner part of $K_2^{\mathbf{a}}$.

Subclaim. There exists a $t \in [d - \frac{r}{2} + 2\hat{r} + 1, d - \hat{r}]$ and a set X'_j that is (ℓ_j, r_j)-scattered in $\mathfrak{B}$ such that $X_j \subseteq R_{\mathfrak{B}}$, $\mathfrak{B} \models \bigwedge_{x \in X_j} \psi_j(x) \Leftrightarrow \mathfrak{B} \models \bigwedge_{x \in X'_j} \psi_j(x)$, and $X'_j \cap I_2^{(t)} = \emptyset$.

Proof of Subclaim. Our goal is to find a flatness pair, say $(\tilde{W}_3, \tilde{\mathfrak{R}}_3)$, that is θ-equivalent to $(\tilde{W}_2, \tilde{\mathfrak{R}}_2)$, and a proper "buffer" t so as to replace the part of X_j that is in $I_2^{(t)}$ to an "equivalent" one that is inside $I_3^{(t)}$. For this replacement to be "safe," we first have to demand that the the influence of $(\tilde{W}_3, \tilde{\mathfrak{R}}_3)$, i.e., the set $I_3^{(w)}$, is disjoint from both the modulator $X_{\text{out}} \cup X'$ and the set X_j. Recall that $\tilde{\mathcal{W}}''$ is a collection of $\hat{\ell} + 2$ flatness pairs of $G \setminus V(\mathbf{a})$ that are θ-equivalent to $(\tilde{W}_1, \tilde{\mathfrak{R}}_1)$ and the vertex sets of their influences are disjoint from $X_{\text{out}} \cup X'$. Therefore, since X_j has size at most $\hat{\ell}$, there exists a flatness pair in $\tilde{\mathcal{W}}'' \setminus \{(\tilde{W}_2, \tilde{\mathfrak{R}}_2)\}$, say $(\tilde{W}_3, \tilde{\mathfrak{R}}_3)$, such that $I_3^{(w)} \cap (X_{\text{out}} \cup X' \cup X_j) = \emptyset$.

We now focus on the set $I_2^{(d)} \setminus I_2^{(d-r+1)}$. Recall that for the set $X_{\text{out}} \cup X'$ it holds that $X' \subseteq Z' \subseteq I_2^{(d-r)}$ and $X_{\text{out}} \cap I_2^{(w)} = \emptyset$. Therefore, $I_2^{(d)} \setminus I_2^{(d-r+1)}$ does not intersect the set $X_{\text{out}} \cup X'$. Since $r = 2 \cdot (\hat{\ell} + 3) \cdot \hat{r}$ and $|X_j| \leq \hat{\ell}$, there exists a $t \in [d - \frac{r}{2} + 2\hat{r} + 1, d - \hat{r}]$ such that X_j does not intersect $I_2^{(t)} \setminus I_2^{(t-\hat{r}+1)}$. Intuitively, we partition the r layers of $\tilde{W}_2$ that are in $I_2^{(d)} \setminus I_2^{(d-r+1)}$ into two parts, the first $r/2$ layers and the second $r/2$ layers, and then we find some layer among the "$\hat{r}$-central" $(\hat{\ell}+1)\hat{r}$ layers of the second part. This layer together with its preceding $\hat{r} - 1$ layers define a "buffer" of size $\hat{r}$ that X_j "avoids" - that is $I_2^{(t)} \setminus I_2^{(t-\hat{r}+1)}$. Notice that $I_2^{(t)} \setminus I_2^{(t-\hat{r}+1)}$ is a subset of $I_2^{(d)} \setminus I_2^{(d-r+1)}$ and therefore $I_2^{(t)} \setminus I_2^{(t-\hat{r}+1)}$ intersects neither X_j nor $X_{\text{out}} \cup X'$.

We set $X_j^\star := X_j \cap I_2^{(t-\hat{r}+1)}$ and $Y_j \subseteq [\ell_j]$ to be the set of indices of the vertices in $X_j^\star$. Notice that $X_j^\star \subseteq R_2$, given that $X_j^\star = X_j \cap I_2^{(t-\hat{r}+1)} \subseteq R_{\mathfrak{B}} \cap I_2^{(t-\hat{r}+1)}$ and $R_{\mathfrak{B}} \cap I_2^{(t-\hat{r}+1)} \subseteq R_2$. Therefore, since $X_j^\star = X_j \cap I_2^{(t-\hat{r}+1)}$, $\psi_j(x)$ is an r_j-local formula (where "r_j-local" refers to distances in $G_{\mathfrak{B}}$), and $\hat{r} \geq r_j$, we have that $\mathfrak{B} \models \bigwedge_{x \in X_j^\star} \psi_j(x) \Leftrightarrow \mathfrak{B}[I_2^{(t)}] \models \bigwedge_{x \in X_j^\star} \psi_j(x)$. To sum up, we observe that the set $X_j^\star$ is a subset of $I_2^{(t-\hat{r}+1)} \cap R_2$ that is $(|Y_j|, r_j)$-scattered in $\mathfrak{B}[I_2^{(t)}]$ (since X_j is $(|Y_j|, r_j)$-scattered in $\mathfrak{B}$) and

$$\mathfrak{B} \models \bigwedge_{x \in X_j^\star} \psi_j(x) \Leftrightarrow \mathfrak{B}[I_2^{(t)}] \models \bigwedge_{x \in X_j^\star} \psi_j(x). \tag{9}$$

Also, notice that $\mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}')[I_2^{(t)}] = \mathfrak{B}[I_2^{(t)}]$. Using the fact that $(\tilde{W}_2, \tilde{\mathfrak{R}}_2)$ is θ-equivalent to $(\tilde{W}_3, \tilde{\mathfrak{R}}_3)$, we now aim to find a set $\tilde{X}_j$ that "equivalent" (in $I_3^{(t)}$) to $X_j^\star$. Since $(\tilde{W}_2, \tilde{\mathfrak{R}}_2)$ is θ-equivalent to $(\tilde{W}_3, \tilde{\mathfrak{R}}_3)$, we have that $\text{in-sig}(\mathfrak{R}_2, R_2, t', L, \emptyset, \emptyset) = \text{in-sig}(\mathfrak{R}_3, R_3, t', L, \emptyset, \emptyset)$, for every $t' \in [w]$. Therefore, we have that $\text{in-sig}(\mathfrak{R}_2, R_2, t, L, \emptyset, \emptyset) = \text{in-sig}(\mathfrak{R}_3, R_3, t, L, \emptyset, \emptyset)$ for the particular value t given above. This implies that there exists a set $\tilde{X}_j \subseteq I_3^{(t-\hat{r}+1)} \cap R_3$ such that $\tilde{X}_j$ is $(|Y_j|, r_j)$-scattered in $\mathfrak{B}[I_3^{(t)}]$ and $\mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}')[I_2^{(t)}] \models \bigwedge_{x \in X_j^\star} \psi_j(x) \Leftrightarrow \mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}')[I_3^{(t)}] \models \bigwedge_{x \in \tilde{X}_j} \psi_j(x)$. Observe that $\mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}')[I_3^{(t)}] = \mathfrak{B}[I_3^{(t)}]$. Thus,

$$\mathfrak{B}[I_2^{(t)}] \models \bigwedge_{x \in X_j^\star} \psi_j(x) \Leftrightarrow \mathfrak{B}[I_3^{(t)}] \models \bigwedge_{x \in \tilde{X}_j} \psi_j(x). \tag{10}$$

Given that $X \cap I_3^{(w)} = \emptyset$, we have that $I_3^{(w)} \subseteq V(G) \setminus X$. Also, since $\tilde{X}_j \subseteq I_3^{(t-\hat{r}+1)}$, for every $x \in \tilde{X}_j$ it holds that $N_{G_{\mathfrak{B}}}^{(\leq \hat{r})}(x) \subseteq I_3^{(t)}$. Thus, since $\psi_j(\mathrm{x})$ is r_j-local, it follows that

$$\mathfrak{B}[I_3^{(t)}] \models \bigwedge_{x \in \tilde{X}_j} \psi_j(x) \Leftrightarrow \mathfrak{B} \models \bigwedge_{x \in \tilde{X}_j} \psi_j(x). \tag{11}$$

We now consider the set

$$X'_j := \left(X_j \setminus X^\star_j\right) \cup \tilde{X}_j.$$

Since $I_3^{(w)} \cap X_j = \emptyset$ and $\hat{r} \geq r_j$, for every $x \in X_j$, and thus, for every $x \in X_j \setminus X^\star_j$, it holds that $N^{(\leq r_j)}_{G_{\mathfrak{B}}}(x) \cap I_3^{(w-\hat{r}+1)} = \emptyset$. Also, since $t \leq w - \hat{r}$ and $\tilde{X}_j \subseteq I_3^{(t-\hat{r}+1)}$, for every $x \in \tilde{X}_j$ it holds that $N^{(\leq r_j)}_{G_{\mathfrak{B}}}(x) \subseteq I_3^{(w-\hat{r}+1)}$. Thus, for every $x \in X_j \setminus X^\star_j$ and $x' \in \tilde{X}_j$ we have that $N^{(\leq r_j)}_{G_{\mathfrak{B}}}(x) \cap N^{(\leq r_j)}_{G_{\mathfrak{B}}}(x') = \emptyset$. The latter, together with the fact that the set $X_j \setminus X^\star_j$ is $(\ell_j - |Y_j|, r_j)$-scattered in $\mathfrak{B}$ and $\tilde{X}_j$ is $(|Y_j|, r_j)$-scattered in $\mathfrak{B}[I_3^{(t)}]$, implies that X'_j is an (ℓ_j, r_j)-scattered set in $\mathfrak{B}$. Moreover, by definition, we have that $X'_j \subseteq R_{\mathfrak{B}} \cup R_3 = R_{\mathfrak{B}}$ (the latter equality holds since $I_3^{(w)} \subseteq V(\mathfrak{B})$) and X'_j does not intersect $I_2^{(t)}$, while, by Equations (9), (10), and (11), we have that $\mathfrak{B} \models \bigwedge_{x\in X_j} \psi_j(x) \Leftrightarrow \mathfrak{B} \models \bigwedge_{x\in X'_j} \psi_j(x)$. The subclaim follows. $\diamond$

Following the above subclaim, let $t \in [d - \frac{r}{2} + 2\hat{r} + 1, d - \hat{r}]$ and let X'_j be a set that is (ℓ_j, r_j)-scattered in $\mathfrak{B}$ such that $X'_j \subseteq R_{\mathfrak{B}}$, $\mathfrak{B} \models \bigwedge_{x\in X_j} \psi_j(x) \Leftrightarrow \mathfrak{B} \models \bigwedge_{x\in X'_j} \psi_j(x)$, and $X'_j \cap I_2^{(t)} = \emptyset$.

Since $r = 2 \cdot (\hat{\ell} + 3) \cdot \hat{r}$ and $|X'_j| \leq \hat{\ell}$, there exists a $t' \in [d - r + 2\hat{r} + 1, d - \frac{r}{2} - \hat{r}]$ such that X'_j does not intersect $I_1^{(t')} \setminus I_1^{(t'-\hat{r}+1)}$. Intuitively, here, we partition the r layers of $\tilde{W}_1$ that are in $I_1^{(d)} \setminus I_1^{(d-r+1)}$ into two parts, the first $r/2$ layers and the second $r/2$ layers, and then we find some layer among the "$\hat{r}$-central" $(\hat{\ell} + 1)\hat{r}$ layers of the first part. This layer together with its preceding $\hat{r} - 1$ layers define a "buffer" of size $\hat{r}$ that X'_j "avoids"—that is $I_1^{(t')} \setminus I_1^{(t'-\hat{r}+1)}$.

Now, consider the set $U_1 := X'_j \cap (I_1^{(t'-\hat{r}+1)} \setminus X_{\text{in}})$. Observe that $U_1 \subseteq R_1$ and therefore $U_1 \subseteq (I_1^{(t'-\hat{r}+1)} \setminus X_{\text{in}}) \cap R_1$. Recall that $Y = I_1^{(\hat{r})}$ and notice that, since $(X'_j \setminus U_1) \cap I_1^{(t')} = \emptyset$ and $t' > \hat{r}$, it holds that $X'_j \setminus U_1 \subseteq R \setminus Y$.

Let $Y'_j \subseteq [\ell_j]$ be the set of the indices of the vertices of X'_j in U_1. Given that $U_1 = X'_j \cap (I_1^{(t'-\hat{r}+1)} \setminus X_{\text{in}})$ and X'_j is (ℓ_j, r_j)-scattered in $\mathfrak{B}$, and $\mathfrak{B} \models \bigwedge_{x\in X'_j} \psi_j(x)$, we get that U_1 is $(|Y'_j|, r_j)$-scattered in $\mathfrak{B}[I_1^{(t')} \setminus X_{\text{in}}]$ and $\mathfrak{B} \models \bigwedge_{x\in U_1} \psi_j(x)$. At this point, observe that, since the formula $\psi_j(\mathrm{x})$ is r_j-local, $U_1 = X'_j \cap (I_1^{(t'-\hat{r}+1)} \setminus X_{\text{in}})$, where $\hat{r} \geq r_j$ and $t' \leq d - \frac{r}{2} - \hat{r}$, for every $x \in U_1$ we have that $N^{(\leq r_j)}_{G_{\mathfrak{B}}}(x) \subseteq I_1^{(t')} \setminus X_{\text{in}} \subseteq I_1^{(d)} \setminus X_{\text{in}}$. The latter implies that

$$\mathfrak{B} \models \bigwedge_{x\in U_1} \psi_j(x) \Leftrightarrow \mathfrak{B}[I_1^{(d)} \setminus X_{\text{in}}] \models \bigwedge_{x\in U_1} \psi_j(x). \tag{12}$$

Also, note that $\mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}')[I_1^{(d)} \setminus X_{\text{in}}] = \mathfrak{B}[I_1^{(d)} \setminus X_{\text{in}}]$.

We will now use the equality of in-signatures. As we mentioned before the statement of Claim 1, $\mathsf{in\text{-}sig}(\mathfrak{R}_1, R_1, d, L, Z, X_{\text{in}}) = \mathsf{in\text{-}sig}(\mathfrak{R}_2, R_2, d, L, Z', X')$. This implies the existence of a set $U_2 \subseteq (I_2^{(t'-\hat{r})} \setminus X') \cap R_2 \subseteq R \setminus Y$ such that U_2 is $(|Y'_j|, r_j)$-scattered in $\mathfrak{B}[I_2^{(t')} \setminus X']$ and

$$\mathfrak{B}[I_1^{(d)} \setminus X_{\text{in}}] \models \bigwedge_{x\in U_1} \psi_j(x) \Leftrightarrow \mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}')[I_2^{(d)} \setminus X'] \models \bigwedge_{x\in U_2} \psi_j(x). \tag{13}$$

By Equations (12) and (13), we derive that

$$\mathfrak{B} \models \bigwedge_{x\in U_1} \psi_j(x) \Leftrightarrow \mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}')[I_2^{(d)} \setminus X'] \models \bigwedge_{x\in U_2} \psi_j(x). \tag{14}$$

Recall that for the set $X_{\text{out}} \cup X'$ it holds that $X' \subseteq I_2^{(d-r)}$ and $X_{\text{out}} \cap I_2^{(w)} = \emptyset$. Since $X_{\text{out}} \cap I_2^{(w)} = \emptyset$, it holds that $I_2^{(t')} \setminus X' \subseteq I_2^{(d)} \setminus X'$ and $I_2^{(d)} \setminus X' \subseteq V(\mathfrak{B}')$.

Since U_2 is $(|Y_j'|, r_j)$-scattered in $\mathfrak{B}[I_2^{(t')} \setminus X']$, where $U_2 \subseteq I_2^{(t'-\hat{r}+1)} \setminus X'$ and $t' < w - \hat{r}$, U_2 is also $(|Y_j'|, r_j)$-scattered in $\mathfrak{B}'$. Moreover, the formula $\psi_j(\mathrm{x})$ is r_j-local, so

$$\mathsf{ap}_{\mathbf{c}}(G, \mathbf{a}')[I_2^{(d)} \setminus X'] \models \bigwedge_{x \in U_2} \psi_j(x) \Leftrightarrow \mathfrak{B}' \models \bigwedge_{x \in U_2} \psi_j(x). \tag{15}$$

Therefore, by Equations (14) and (15), it follows that $\mathfrak{B} \models \bigwedge_{x \in U_1} \psi_j(x) \Leftrightarrow \mathfrak{B}' \models \bigwedge_{x \in U_2} \psi_j(x)$.

Consider the set

$$X_j^\bullet := (X_j' \setminus U_1) \cup U_2.$$

Notice that since $X_j' \setminus U_1 \subseteq V(\mathfrak{B})$ and $X_j' \setminus U_1$ does not intersect neither $I_2^{(d-r+1)}$ (where X' lies), nor $I_1^{(d-r+1)} \subseteq I_1^{(t')}$ (where X_{in} lies), it follows that $X_j' \setminus U_1 \subseteq V(\mathfrak{B}) \cap V(\mathfrak{B}')$. This implies that $X_j' \setminus U_1$ is an $(\ell_j - |Y_j'|, r_j)$-scattered set in $\mathfrak{B}$ and an $(\ell_j - |Y_j'|, r_j)$-scattered set in $\mathfrak{B}'$. Since $U_2 \subseteq I_2^{(t'-\hat{r}+1)} \setminus Z'$, $X_j' \cap I_2^{(t)} = \emptyset$, and $t' \leq t - 2\hat{r}$, we have that for every $x \in X_j' \setminus U_1$ and $x' \in U_2$ it holds that $N_{G_{\mathfrak{B}'}}^{(\leq r_j)}(x) \cap N_{G_{\mathfrak{B}'}}^{(\leq r_j)}(x') = \emptyset$. The latter, together with the fact that $X_j' \setminus U_1$ is an $(\ell_j - |Y_j'|, r_j)$-scattered set in $\mathfrak{B}'$ and U_2 is a $(|Y_j'|, r_j)$-scattered set in $\mathfrak{B}'$, implies that $X_j^\bullet$ is an (ℓ_j, r_j)-scattered set in $\mathfrak{B}'$. Also, notice that $X_j^\bullet \subseteq R_{\mathfrak{B}}' \setminus Y$. Furthermore, since the formula $\psi_j(\mathrm{x})$ is r_j-local, it follows that $\mathfrak{B}' \models \bigwedge_{x \in X_j} \psi_j(x) \Leftrightarrow \mathfrak{B}' \models \bigwedge_{x \in X_j^\bullet} \psi_j(x)$.

Thus, assuming that there is a set $X_j \subseteq R_{\mathfrak{B}}$ that is (ℓ_j, r_j)-scattered in $\mathfrak{B}$ and $\mathfrak{B} \models \bigwedge_{x \in X_j} \psi_j(x)$, we proved that there is a set $X_j^\bullet \subseteq R_{\mathfrak{B}}' \setminus Y \subseteq R \setminus Y$ that is (ℓ_j, r_j)-scattered in $\mathfrak{B}'$ and $\mathfrak{B}' \models \bigwedge_{x \in X_j^\bullet} \psi_j(x)$.

To conclude Case 1, notice that we can prove the inverse implication, i.e., by assuming the existence of a set $X_j^\bullet \subseteq R_{\mathfrak{B}}' \setminus Y \subseteq R \setminus Y$ that is (ℓ_j, r_j)-scattered in $\mathfrak{B}'$ and $\mathfrak{B}' \models \bigwedge_{x \in X_j^\bullet} \psi_j(x)$ and, by using the same arguments as above (replacing $(\tilde{W}_1, \tilde{\mathfrak{R}}_1)$ with $(\tilde{W}_2, \tilde{\mathfrak{R}}_2)$, X_{in} with X' and R with $R \setminus Y$), we can prove the existence of a set $X_j \subseteq R$ that is (ℓ_j, r_j)-scattered in $\mathfrak{B}$ such that $\mathfrak{B} \models \bigwedge_{x \in X_j} \psi_j(x)$.

Case 2. $j \notin J$.

We aim to prove that $\mathsf{ap}_{\mathbf{c}}((G, R, \mathbf{a}')[C]) \models \neg\breve{\zeta}_j \Leftrightarrow \mathsf{ap}_{\mathbf{c}}((G, R \setminus Y, \mathbf{a}')[C']) \models \neg\breve{\zeta}_j$. In other words, we show that for every set $X_j \subseteq R \cap C$ that is (ℓ_j, r_j)-scattered in $\mathfrak{B}$, $\mathfrak{B} \models \neg\psi_j(x)$, for some $x \in X_j$ if and only if for every set $X_j' \subseteq (R \setminus Y) \cap C'$ that is (ℓ_j, r_j)-scattered in $\mathfrak{B}'$, $\mathfrak{B}' \models \neg\psi_j(x)$, for some $x \in X_j'$. In Case 1, we showed that there is a set $X_j \subseteq R_{\mathfrak{B}}$ that is (ℓ_j, r_j)-scattered in $\mathfrak{B}$ and $\mathfrak{B} \models \bigwedge_{x \in X_j} \psi_j(x)$ if and only if there is a set $X_j^\bullet \subseteq R_{\mathfrak{B}}' \setminus Y \subseteq R \setminus Y$ that is (ℓ_j, r_j)-scattered in $\mathfrak{B}'$ and $\mathfrak{B}' \models \bigwedge_{x \in X_j^\bullet} \psi_j(x)$. This directly implies that $(\mathfrak{B}, R_{\mathfrak{B}}) \models \neg\breve{\zeta}_j \Leftrightarrow (\mathfrak{B}', R_{\mathfrak{B}}' \setminus Y) \models \neg\breve{\zeta}_j$. This concludes Case 2 and completes the proof of Claim 2. ⋄

Also, recall that the algorithm `Find_Equiv_FlatPairs` in Section 7.4 outputs a central vertex v of W_1 and the set $Y = V(\mathsf{compass}_{\breve{\mathfrak{R}}'}(\breve{W}'))$, where $(\breve{W}', \breve{\mathfrak{R}}')$ is a $\breve{W}'$-tilt of $(W, \mathfrak{R})$ and $\breve{W}$ is the central j'-subwall of W_1. Finally, recall that $j' = 2\hat{r} + 2$. The definition of a tilt of a flatness pair implies that $v \in Y$. By Claim 1, $(\mathsf{stell}(G, X), X) \models \beta \Leftrightarrow (\mathsf{stell}(G \setminus v, X_{\text{out}} \cup X'), X_{\text{out}} \cup X') \models \beta$.

Recall that all the basic Gaifman variables in $\breve{\zeta}_{\mathsf{R}}$ are contained in R and every $\psi_i(\mathrm{x})$ is r_i-local. The fact that v is a central vertex of W_1, $\breve{W}$ has height $j' = 2\hat{r} + 2$, and $R \cap Y = \emptyset$ implies that none of the local formulas $\psi_i(\mathrm{x})$ is evaluated using v. Therefore, $\mathsf{ap}_{\mathbf{c}}((G, R \setminus Y, \mathbf{a})[V(G) \setminus (X_{\text{out}} \cup X')]) \models \breve{\zeta}_{\mathsf{R}} \Leftrightarrow \mathsf{ap}_{\mathbf{c}}((G \setminus v, R \setminus Y, \mathbf{a})[V(G) \setminus (X_{\text{out}} \cup X')]) \models \breve{\zeta}_{\mathsf{R}}$, and, by Claim 2, $\mathsf{ap}_{\mathbf{c}}((G, R, \mathbf{a})[V(G) \setminus X]) \models \breve{\zeta}_{\mathsf{R}} \Leftrightarrow \mathsf{ap}_{\mathbf{c}}((G \setminus v, R \setminus Y, \mathbf{a}')[V(G) \setminus (X_{\text{out}} \cup X')]) \models \breve{\zeta}_{\mathsf{R}}$. Thus, we get that $(G, R, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}} \Leftrightarrow (G \setminus v, R \setminus Y, \mathbf{a}) \models \theta_{\mathsf{R},\mathbf{c}}$.

8 Limitations and Further Directions

To conclude the article, in Section 8.1 we justify the necessity of the ingredients of our logic $\mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$. Next, in Section 8.2 we present several directions and open problems for further research.

8.1 Natural Limitations

We now wish to comment on why the two basic ingredients of the definition of $\mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ are necessary for the statement and the proof of a meta-algorithmic result such as Theorem 3.

The first ingredient of $\mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ is that the modulator sentences belong in $\mathsf{CMSO}^{\mathbf{tw}}$ which is defined so that the treewidth of $\mathbf{torso}(G, X)$ is bounded. While it is known that bounding the treewidth is necessary for CMSO-model-checking [31, 79], one may ask why it is not enough to just bound the treewidth of $G[X]$. To see why this is unavoidable, consider a graph G and let G' be the graph obtained from G by subdividing each edge once. Then, asking whether G is Hamiltonian, which is a well-known NP-complete problem [54], is equivalent to asking whether G' has a vertex set S' such that $G'[S']$ is a cycle and such that $G' \setminus S'$ is an edgeless graph, that is, a K_2-minor-free graph. Notice that, while $\mathbf{tw}(G'[S']) = 2$, $\mathbf{torso}(G', S') = G$ has unbounded treewidth.

The second ingredient of $\mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ is the FO demand. This is also necessary, as otherwise we may choose some target property σ not definable in FO, such as Hamiltonicity, which is CMSO-definable and NP-complete on planar graphs [54]. Without the restriction that σ needs to be FO-definable, a void modulator would be able to model this NP-complete problem on planar input graphs.

8.2 Further Research

The Minor-Exclusion Framework. The graph-structural horizon in Theorem 3 is delimited by minor-exclusion. This restriction is hard-wired in our proof in the way it combines the Flat Wall theorem with Gaifman's theorem. Recently, several efficient algorithms appeared for modification problems assuming topological minor-freeness (see [3, 50, 68, 97] and the meta-algorithmic result in [97]). For such classes, to achieve efficient model-checking for $\mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$, or some fragment of it, is an interesting open challenge.

Quadratic Time. The proof of Theorem 3 can be seen as a possible "meta-algorithmization" of the irrelevant vertex technique introduced by Robertson and Seymour [98], going further than the two known recent attempts in this direction [45, 59]. The main routine of the algorithm transforms the input of the problem to a simpler graph by detecting territories in it that can be safely discarded, therefore producing a simpler instance. This routine is applied repetitively until the graph has "small" treewidth, so that the problem can be solved in linear time by using Courcelle's theorem. This approach gives an algorithm running in quadratic time. Any improvement of this running time should rely on techniques escaping the above scheme of gradual simplification. The only results in this direction are the cases of making a graph planar by deleting at most k vertices (resp. edges) in [69] (resp. [71]) that run in time $O_k(n)$.

Other Modification Problems. One can observe that the definition of $\mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ readily models a wide variety of modification problems involving edge or vertex removals. Is it possible to extend $\mathsf{CMSO}^{\mathbf{tw}} \triangleright \mathsf{FO}$ so that it can also deal with other (local) operations such as edge contractions, edge additions, or others? This was done in [45] for the case of vertex removals and edge removals/additions/contractions to achieve planarity and an FO-definable property. Other types of (not necessarily local) modification operations where studied in [46, 59].

Acknowledgments

We wish to thank Stavros G. Kolliopoulos and Christophe Paul, as well as some anonymous reviewers, for their valuable remarks on earlier versions of this article.

References

[1] Isolde Adler, Martin Grohe, and Stephan Kreutzer. 2008. Computing excluded minors. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 641–650. Retrieved from http://portal.acm.org/citation.cfm?id=1347082.1347153

[2] Akanksha Agrawal, Lawqueen Kanesh, Daniel Lokshtanov, Fahad Panolan, M. S. Ramanujan, and Saket Saurabh. 2021. Elimination distance to topological-minor-free graphs is FPT. arXiv:2104.09950. Retrieved from https://arxiv.org/abs/2104.09950

[3] Akanksha Agrawal, Lawqueen Kanesh, Daniel Lokshtanov, Fahad Panolan, M. S. Ramanujan, Saket Saurabh, and Meirav Zehavi. 2022. Deleting, eliminating and decomposing to hereditary classes are all FPT-equivalent. In *Proceedings of the 32st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1976–2004. DOI: https://doi.org/10.1137/1.9781611977073.79

[4] Akanksha Agrawal, Lawqueen Kanesh, Fahad Panolan, M. S. Ramanujan, and Saket Saurabh. 2021. An FPT algorithm for elimination distance to bounded degree graphs. In *Proceedings of the 38th International Symposium on Theoretical Aspects of Computer Science (STACS).* Leibniz International Proceedings in Informatics, Vol. 187, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 5:1–5:11. DOI: https://doi.org/10.4230/LIPIcs.STACS.2021.5

[5] Stefan Arnborg, Jens Lagergren, and Detlef Seese. 1991. Easy problems for tree-decomposable graphs. *Journal of Algorithms* 12 (1991), 308–340. DOI: https://doi.org/10.1016/0196-6774(91)90006-K

[6] Julien Baste, Ignasi Sau, and Dimitrios M. Thilikos. 2020. A complexity dichotomy for hitting connected minors on bounded treewidth graphs: The chair and the banner draw the boundary. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 951–970. DOI: https://doi.org/10.1137/1.9781611975994.57

[7] Patrick Bellenbaum and Reinhard Diestel. 2002. Two Short Proofs Concerning Tree-Decompositions. *Combinatorics, Probability & Computing* 11, 6 (2002), 541–547. DOI: https://doi.org/10.1017/S0963548302005369

[8] Hans L. Bodlaender. 1993. A linear time algorithm for finding tree-decompositions of small treewidth. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC)*, 226–234. DOI: https://doi.org/10.1145/167088.167161

[9] Hans L. Bodlaender, Pål Grønås Drange, Markus S. Dregi, Fedor V. Fomin, Daniel Lokshtanov, and Michał Pilipczuk. 2016. A $c^k n$ 5-Approximation algorithm for treewidth. *SIAM Journal on Computing* 45, 2 (2016), 317–378. DOI: https://doi.org/10.1137/130947374

[10] Hans L. Bodlaender, Fedor V. Fomin, Daniel Lokshtanov, Eelko Penninkx, Saket Saurabh, and Dimitrios M. Thilikos. 2016. (Meta) Kernelization. *Journal of the ACM* 63, 5 (2016), 44:1–44:69. DOI: https://doi.org/10.1145/2973749

[11] Mikołaj Bojańczyk, Martin Grohe, and Michał Pilipczuk. 2021. Definable decompositions for graphs of bounded linear cliquewidth. *Logical Methods in Computer Science* 17, 1 (2021), 5:1–5:40. DOI: https://doi.org/10.23638/LMCS-17(1:5)2021

[12] Mikołaj Bojańczyk and Michał Pilipczuk. 2016. Definability equals recognizability for graphs of bounded treewidth. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 407–416. DOI: https://doi.org/10.1145/2933575.2934508

[13] Mikołaj Bojańczyk. 2021. Separator logic and star-free expressions for graphs. arXiv:2107.13953. Retrieved from https://arxiv.org/abs/2107.13953

[14] Édouard Bonnet, Jan Dreier, Jakub Gajarský, Stephan Kreutzer, Nikolas Mählmann, Pierre Simon, and Szymon Toruńczyk. 2022. Model checking on interpretations of classes of bounded local cliquewidth. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer (LICS).* ACM, New York, NY, 54:1–54:13. DOI: https://doi.org/10.1145/3531130.3533367

[15] Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Toruńczyk. 2022. Twin-width IV: Ordered graphs and matrices. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC).* ACM, New York, NY. DOI: https://doi.org/10.1145/3519935.3520037

[16] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. 2022. Twin-width I: Tractable FO model checking. *Journal of the ACM* 69, 1 (2022), 3:1–3:46. DOI: https://doi.org/10.1145/3486655

[17] Édouard Bonnet, Jaroslav Nešetřil, Patrice Ossona de Mendez, Sebastian Siebertz, and Stéphan Thomassé. 2021. Twin-width and permutations. arXiv:2102.06880. Retrieved from https://arxiv.org/abs/2102.06880

[18] Richard B. Borie, R. Gary Parker, and Craig A. Tovey. 1992. Automatic generation of linear-time algorithms from predicate calculus descriptions of problems on recursively constructed graph families. *Algorithmica* 7, 5–6 (1992), 555–581. DOI: https://doi.org/10.1007/BF01758777

[19] Marin Bougeret, Bart M. P. Jansen, and Ignasi Sau. 2022. Bridge-depth characterizes which minor-closed structural parameterizations of vertex cover admit a polynomial kernel. *SIAM Journal on Discrete Mathematics* 36, 4 (2022), 2737–2773. DOI: https://doi.org/10.1137/21M1400766

[20] Nicolas Bousquet, Jean Daligault, and Stéphan Thomassé. 2018. Multicut is FPT. *SIAM Journal on Computing*47, 1 (2018), 166–207. DOI: https://doi.org/10.1137/140961808

[21] Jannis Bulian and Anuj Dawar. 2016. Graph isomorphism parameterized by elimination distance to bounded degree. *Algorithmica* 75, 2 (2016), 363–382. DOI: https://doi.org/10.1007/s00453-015-0045-3

[22] Jannis Bulian and Anuj Dawar. 2017. Fixed-parameter tractable distances to sparse graph classes. *Algorithmica* 79, 1 (2017), 139–158. DOI: https://doi.org/10.1007/s00453-016-0235-7

[23] Leizhen Cai. 1996. Fixed-parameter tractability of graph modification problems for hereditary properties. *Information Processing Letters* 58, 4 (1996), 171–176. DOI: https://doi.org/10.1016/0020-0190(96)00050-6

[24] Jianer Chen, Fedor V. Fomin, Yang Liu, Songjian Lu, and Yngve Villanger. 2008. Improved algorithms for feedback vertex set problems. *The Journal of Computer and System Sciences* 74, 7 (2008), 1188–1198. DOI: https://doi.org/10.1016/j.jcss.2008.05.002

[25] Jianer Chen, Iyad A. Kanj, and Ge Xia. 2006. Improved parameterized upper bounds for vertex cover. In *Proceedings of the 31st International Symposium on Mathematical Foundations of Computer Science (MFCS)*. Lecture Notes in Computer Science, Vol. 4162, 238–249. DOI: https://doi.org/10.1007/11821069_21

[26] Julia Chuzhoy. 2015. Improved bounds for the flat wall theorem. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 256–275. DOI: https://doi.org/10.1137/1.9781611973730.20

[27] Bruno Courcelle. 1990. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and Computation* 85, 1 (1990), 12–75. DOI: https://doi.org/10.1016/0890-5401(90)90043-H

[28] Bruno Courcelle. 1992. The monadic second-order logic of graphs III: Tree-decompositions, minor and complexity issues. *RAIRO - Theoretical Informatics and Applications* 26 (1992), 257–286. DOI: https://doi.org/10.1051/ita/1992260302571

[29] Bruno Courcelle. 1997. The expression of graph properties and graph transformations in monadic second-order logic. In *Handbook of Graph Grammars and Computing by Graph Transformations, Vol. 1: Foundations*. World Scientific, 313–400.

[30] Bruno Courcelle and Joost Engelfriet. 2012. *Graph Structure and Monadic Second-Order Logic – A Language-Theoretic Approach*. Encyclopedia of mathematics and its applications, Vol. 138. Cambridge University Press.

[31] Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. 2000. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory of Computing Systems* 33, 2 (2000), 125–150. DOI: https://doi.org/10.1007/s002249910009

[32] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. 2015. *Parameterized Algorithms*. Springer. DOI: https://doi.org/10.1007/978-3-319-21275-3

[33] Marek Cygan, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. 2019. Minimum bisection is fixed-parameter tractable. *SIAM Journal on Computing* 48, 2 (2019), 417–450. DOI: https://doi.org/10.1137/140988553

[34] Marek Cygan, Dániel Marx, Marcin Pilipczuk, and Michał Pilipczuk. 2017. Hitting forbidden subgraphs in graphs of bounded treewidth. *Information and Computation* 256 (2017), 62–82. DOI: https://doi.org/10.1016/j.ic.2017.04.009

[35] Anuj Dawar, Martin Grohe, and Stephan Kreutzer. 2007. Locally excluding a minor. In *Proceedings of the 21st IEEE Symposium on Logic in Computer Science (LICS)*, 270–279. DOI: https://doi.org/10.1109/LICS.2007.31

[36] Reinhard Diestel. 2017. *Graph Theory* (5th ed.), Vol. 173, Springer-Verlag. DOI: https://doi.org/10.1007/978-3-662-53622-3

[37] Jan Dreier and Peter Rossmanith. 2021. Approximate evaluation of first-order counting queries. In *Proceedings of the 32nd ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Dániel Marx (Ed.), SIAM, 1720–1739. DOI: https://doi.org/10.1137/1.9781611976465.104

[38] Zdeněk Dvořák, Daniel Král, and Robin Thomas. 2013. Testing first-order properties for subclasses of sparse graphs. *The Journal of the ACM* 60, 5 (2013), Article 36, 24 pages. DOI: https://doi.org/10.1145/2499483

[39] Eduard Eiben, Robert Ganian, Thekla Hamm, and O-joung Kwon. 2021. Measuring what matters: A hybrid approach to dynamic programming with treewidth. *The Journal of Computer and System Sciences* 121 (2021), 57–75. DOI: https://doi.org/10.1016/j.jcss.2021.04.005

[40] Kord Eickmeyer, Michael Elberfeld, and Frederik Harwath. 2017. Succinctness of order-invariant logics on depth-bounded structures. *ACM Transactions on Computational Logic* 18, 4 (2017), 33:1–33:25. DOI: https://doi.org/10.1145/3152770

[41] Samuel Fiorini, Nadia Hardy, Bruce A. Reed, and Adrian Vetta. 2008. Planar graph bipartization in linear time. *Discrete Applied Mathematics* 156, 7 (2008), 1175–1180. DOI: https://doi.org/10.1016/j.dam.2007.08.013

[42] Jörg Flum and Martin Grohe. 2001. Fixed-parameter tractability, definability, and model-checking. *The SIAM Journal on Computing* 31, 1 (2001), 113–145. DOI: https://doi.org/10.1137/S0097539799360768

[43] Fedor V. Fomin, Petr A. Golovach, Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. 2021. Compound logics for modification problems. arXiv:2111.02755. Retrieved from https://arxiv.org/abs/2111.02755

[44] Fedor V. Fomin, Petr A. Golovach, Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. 2023. Compound logics for modification problems. In *Proceedings of the 50th International Colloquium on Automata, Languages, and Programming (ICALP).* Leibniz International Proceedings in Informatics, Vol. 261, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 61:1–61:21. DOI: https://doi.org/10.4230/LIPICS.ICALP.2023.61

[45] Fedor V. Fomin, Petr A. Golovach, Giannos Stamoulis, and Dimitrios M. Thilikos. 2020. An algorithmic meta-theorem for graph modification to planarity and FOL. In *Proceedings of the 28th Annual European Symposium on Algorithms (ESA).* Leibniz International Proceedings in Informatics, Vol. 173, 51:1–51:17. DOI: https://doi.org/10.4230/LIPIcs.ESA.2020.51

[46] Fedor V. Fomin, Petr A. Golovach, and Dimitrios M. Thilikos. 2019. Modification to planarity is fixed parameter tractable. In *Proceedings of the 36th International Symposium on Theoretical Aspects of Computer Science (STACS),* 28:1–28:17. DOI: https://doi.org/10.4230/LIPIcs.STACS.2019.28

[47] Fedor V. Fomin, Petr A. Golovach, and Dimitrios M. Thilikos. 2022. Parameterized complexity of elimination distance to first-order logic properties. *ACM Transactions on Computational Logic* 23, 3 (2022), 17:1–17:35. DOI: https://doi.org/10.1145/3517129

[48] Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, Geevarghese Philip, and Saket Saurabh. 2016. Hitting forbidden minors: Approximation and kernelization. *SIAM Journal on Discrete Mathematics* 30, 1 (2016), 383–410. DOI: https://doi.org/10.1137/140997889

[49] Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. 2012. Planar F-deletion: Approximation, kernelization and optimal FPT algorithms. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS),* 470–479. DOI: https://doi.org/10.1109/FOCS.2012.62

[50] Fedor V. Fomin, Daniel Lokshtanov, Fahad Panolan, Saket Saurabh, and Meirav Zehavi. 2020. Hitting topological minors is FPT. In *Proceedings of the 52nd Annual ACM Symposium on Theory of Computing (STOC),* 1317–1326. DOI: https://doi.org/10.1145/3357713.3384318

[51] Markus Frick and Martin Grohe. 2001. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM* 48, 6 (2001), 1184–1206. DOI: https://doi.org/10.1145/504794.504798

[52] Haim Gaifman. 1982. On Local and Non-Local Properties. In *Proceedings of the Herbrand Symposium (Studies in Logic and the Foundations of Mathematics, Vol. 107).* Elsevier, 105–135. DOI: https://doi.org/10.1016/S0049-237X(08)71879-2

[53] Jakub Gajarský, Stephan Kreutzer, Jaroslav Nešetřil, Patrice Ossona de Mendez, Michał Pilipczuk, Sebastian Siebertz, and Szymon Toruńczyk. 2020. First-order interpretations of bounded expansion classes. *ACM Transactions on Computational Logic* 21, 4 (2020), 29:1–29:41. DOI: https://doi.org/10.1145/3382093

[54] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman. Retrieved from https://dl.acm.org/doi/10.5555/574848

[55] Archontia C. Giannopoulou, Michał Pilipczuk, Jean-Florent Raymond, Dimitrios M. Thilikos, and Marcin Wrochna. 2021. Linear kernels for edge deletion problems to immersion-closed graph classes. *SIAM Journal on Discrete Mathematics* 35, 1 (2021), 105–151. DOI: https://doi.org/10.1137/18M1228839

[56] Archontia C. Giannopoulou and Dimitrios M. Thilikos. 2013. Optimizing the graph minors weak structure theorem. *SIAM Journal on Discrete Mathematics* 27, 3 (2013), 1209–1227. DOI: https://doi.org/10.1137/110857027

[57] Petr A. Golovach, Dieter Kratsch, and Daniël Paulusma. 2013. Detecting induced minors in AT-free graphs. *Theoretical Computer Science* 482 (2013), 20–32. DOI: https://doi.org/10.1016/j.tcs.2013.02.029

[58] Petr A. Golovach, Giannos Stamoulis, and Dimitrios M. Thilikos. 2022. Model-checking for first-order logic with disjoint paths predicates in proper minor-closed graph classes. arXiv:2211.01723. Retrieved from https://arxiv.org/abs/2211.01723

[59] Petr A. Golovach, Giannos Stamoulis, and Dimitrios M. Thilikos. 2023. Model-checking for first-order logic with disjoint paths predicates in proper minor-closed graph classes. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA '23).* SIAM, 3684–3699. DOI: https://doi.org/10.1137/1.9781611977554.ch141

[60] Julien Grange. 2021. Successor-invariant first-order logic on classes of bounded degree. *Logical Methods in Computer Science* 17, 3 (2021), 20:1–20:25. DOI: https://doi.org/10.46298/lmcs-17(3:20)2021

[61] Martin Grohe. 2008. Logic, graphs, and algorithms. In *Logic and Automata: History and Perspectives, in Honor of Wolfgang Thomas.* Texts in Logic and Games, Vol. 2, Amsterdam University Press, 357–422. Retrieved from https://eccc.weizmann.ac.il/report/2007/091/

[62] Martin Grohe and Stephan Kreutzer. 2009. Methods for algorithmic meta theorems. In *Model Theoretic Methods in Finite Combinatorics - AMS-ASL Joint Special Session*, Vol. 558, AMS, 181–206. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.395.8282 & rep=rep1 & type=pdf

[63] Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. 2017. Deciding first-order properties of nowhere dense graphs. *Journal of the ACM* 64, 3 (2017), 17:1–17:32. DOI: https://doi.org/10.1145/3051095

[64] Martin Grohe and Nicole Schweikardt. 2018. First-order query evaluation with cardinality conditions. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS).* ACM, New York, NY, 253–266. DOI: https://doi.org/10.1145/3196959.3196970

[65] Martin Grohe and Thomas Schwentick. 2000. Locality of order-invariant first-order formulas. *ACM Transactions on Computational Logic* 1, 1 (2000), 112–130. DOI: https://doi.org/10.1145/343369.343386

[66] Jiong Guo, Falk Hüffner, and Rolf Niedermeier. 2004. A structural view on parameterizing problems: Distance from triviality. In *Proceedings of the 1st International Workshop on Parameterized and Exact Computation (IWPEC),* Lecture Notes in Computer Science, Vol. 3162, 162–173. DOI: https://doi.org/10.1007/978-3-540-28639-4_15

[67] Anupam Gupta, Euiwoong Lee, and Jason Li. 2018. Faster exact and approximate algorithms for k-cut. In *Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science (FOCS),* 113–123. DOI: https://doi.org/10.1109/FOCS.2018.00020

[68] Bart M. P. Jansen, Jari J. H. de Kroon, and Michał Włodarczyk. 2021. Vertex deletion parameterized by elimination distance and even less. In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing (STOC)*, 1757–1769. DOI: https://doi.org/10.1145/3406325.3451068

[69] Bart M. P. Jansen, Daniel Lokshtanov, and Saket Saurabh. 2014. A near-optimal planarization algorithm. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1802–1811. DOI: https://doi.org/10.1137/1.9781611973402.130

[70] Ken-ichi Kawarabayashi. 2009. Planarity allowing few error vertices in linear time. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS),* 639–648. DOI: https://doi.org/10.1109/FOCS.2009.45

[71] Ken-ichi Kawarabayashi and Bruce A. Reed. 2007. Computing crossing number in linear time. In *Proceedings of the 39th annual ACM symposium on Theory of computing (STOC)*, 382–390. DOI: https://doi.org/10.1145/1250790.1250848

[72] Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. 2018. A new proof of the flat wall theorem. *Journal of Combinatorial Theory, Series B* 129 (2018), 204–238. DOI: https://doi.org/10.1016/j.jctb.2017.09.006

[73] Ken-ichi Kawarabayashi and Mikkel Thorup. 2011. The minimum k-way cut of bounded size is fixed-parameter tractable. In *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science (FOCS),* 160–169. DOI: https://doi.org/10.1109/FOCS.2011.53

[74] Eun Jung Kim, Alexander Langer, Christophe Paul, Felix Reidl, Peter Rossmanith, Ignasi Sau, and Somnath Sikdar. 2016. Linear kernels and single-exponential algorithms via protrusion decompositions. *ACM Transactions on Algorithms* 12, 2 (2016), 21:1–21:41. DOI: https://doi.org/10.1145/2797140

[75] Philip N. Klein and Dániel Marx. 2012. Solving planar k-terminal cut in $O(n^{c}k)$ time. In *Proceedings of the 39th International Colloquium of Automata, Languages and Programming (ICALP),* Lecture Notes in Computer Science, Vol. 7391, 569–580. DOI: https://doi.org/10.1007/978-3-642-31594-7_48

[76] Tomasz Kociumaka and Marcin Pilipczuk. 2014. Faster deterministic feedback vertex set. *Inform. Process. Lett.* 114, 10 (2014), 556–560. DOI: https://doi.org/10.1016/j.ipl.2014.05.001

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

[78] Tuukka Korhonen. 2021. A single-exponential time 2-approximation algorithm for treewidth. In *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS),* 184–192. DOI: https://doi.org/10.1109/FOCS52979.2021.00026

[79] Stephan Kreutzer. 2011. Algorithmic meta-theorems. In *Finite and Algorithmic Model Theory*. London Mathematical Society Lecture Note Series, Vol. 379, Cambridge University Press, 177–270. Retrieved from http://www.cs.ox.ac.uk/people/stephan.kreutzer/Publications/amt-survey.pdf

[80] Dietrich Kuske and Nicole Schweikardt. 2017. First-order logic with counting. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE Computer Society, 1–12. Retrieved from https://doi.org/10.1109/LICS.2017.8005133

[81] Dietrich Kuske and Nicole Schweikardt. 2018. Gaifman normal forms for counting extensions of first-order logic. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP)*, , Leibniz International Proceedings in Informatics, Vol. 107, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 133:1–133:14. DOI: https://doi.org/10.4230/LIPIcs.ICALP.2018.133

[82] John M. Lewis and Mihalis Yannakakis. 1980. The node-deletion problem for hereditary properties is NP-complete. *The Journal of Computer and System Sciences* 20, 2 (1980), 219–230. DOI: https://doi.org/10.1016/0022-0000(80)90060-4

[83] Leonid Libkin. 2004. *Elements of Finite Model Theory*. Springer. DOI: https://doi.org/10.1007/978-3-662-07003-1

[84] Alexander Lindermayr, Sebastian Siebertz, and Alexandre Vigny. 2020. Elimination distance to bounded degree on planar graphs. In *Proceedings of the 45th International Symposium on Mathematical Foundations of Computer Science (MFCS),* Leibniz International Proceedings in Informatics, Vol. 170, 65:1–65:12. DOI: https://doi.org/10.4230/LIPIcs.MFCS.2020.65

[85] Daniel Lokshtanov, Saket Saurabh, and Vaishali Surianarayanan. 2020. A parameterized approximation scheme for min k-cut. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS),* 798–809. DOI: https://doi.org/10.1109/FOCS46700.2020.00079

[86] Dániel Marx and Igor Razgon. 2014. Fixed-parameter tractability of multicut parameterized by the size of the cutset. *The SIAM Journal on Computing* 43, 2 (2014), 355–388. DOI: https://doi.org/10.1137/110855247

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

[88] Laure Morelle, Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. 2024. Faster parameterized algorithms for modification problems to minor-closed classes. *TheoretiCS* 3 (2024), 1–75. DOI: https://doi.org/10.46298/theoretics.24.19

[89] Jaroslav Nešetřil and Patrice Ossona de Mendez. 2012. *Sparsity - Graphs, Structures, and Algorithms*. Algorithms and combinatorics, Vol. 28, Springer. DOI: https://doi.org/10.1007/978-3-642-27875-4

[90] Jaroslav Nešetřil and Patrice Ossona de Mendez. 2008. Grad and classes with bounded expansion. I. Decompositions. *European Journal of Combinatorics* 29, 3 (2008), 760–776. DOI: https://doi.org/10.1016/j.ejc.2006.07.013

[91] Jaroslav Nešetřil and Patrice Ossona de Mendez. 2008. Grad and classes with bounded expansion. II. Algorithmic aspects. *European Journal of Combinatorics* 29, 3 (2008), 777–791. DOI: https://doi.org/10.1016/j.ejc.2006.07.014

[92] Jaroslav Nešetřil and Patrice Ossona de Mendez. 2008. Grad and classes with bounded expansion. III. Restricted graph homomorphism dualities. *European Journal of Combinatorics* 29, 4 (2008), 1012–1024. DOI: https://doi.org/10.1016/j.ejc.2007.11.019

[93] Jaroslav Nešetřil, Patrice Ossona de Mendez, Michał Pilipczuk, Roman Rabinovich, and Sebastian Siebertz. 2021. Rankwidth meets stability. In *Proceedings of the 32nd ACM-SIAM Symposium on Discrete Algorithms (SODA).* SIAM, 2014–2033. DOI: https://doi.org/10.1137/1.9781611976465.120

[94] Jaroslav Nešetřil, Patrice Ossona de Mendez, Roman Rabinovich, and Sebastian Siebertz. 2020. Linear rankwidth meets stability. In *Proceedings of the 31st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1180–1199. DOI: https://doi.org/10.1137/1.9781611975994.72

[95] Jaroslav Nešetřil, Patrice Ossona de Mendez, and Sebastian Siebertz. 2022. Structural properties of the first-order transduction quasiorder. In *Proceedings of the 30th EACSL Annual Conference on Computer Science Logic (CSL),* Leibniz International Proceedings in Informatics, Vol. 216, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 31:1–31:16. DOI: https://doi.org/10.4230/LIPIcs.CSL.2022.31

[96] Patrice Ossona de Mendez, Michał Pilipczuk, and Sebastian Siebertz. 2022. Transducing paths in graph classes with unbounded shrubdepth. *European Journal of Combinatorics* (2022), 103660. DOI: https://doi.org/10.1016/j.ejc.2022.103660

[97] Michał Pilipczuk, Nicole Schirrmacher, Sebastian Siebertz, Szymon Toruńczyk, and Alexandre Vigny. 2022. Algorithms and Data Structures for First-Order Logic with Connectivity Under Vertex Failures. In *Proceedings of the 49th International Colloquium on Automata, Languages, and Programming (ICALP),* Leibniz International Proceedings in Informatics, Vol. 229, 102:1–102:18. DOI: https://doi.org/10.4230/LIPIcs.ICALP.2022.102

[98] Neil Robertson and Paul D. Seymour. 1995. Graph minors. XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Series B* 63, 1 (1995), 65–110. DOI: https://doi.org/10.1006/jctb.1995.1006

[99] Neil Robertson and Paul D. Seymour. 2004. Graph Minors. XX. Wagner's conjecture. *Journal of Combinatorial Theory, Series B* 92, 2 (2004), 325–357. DOI: https://doi.org/10.1016/j.jctb.2004.08.001

[100] Ignasi Sau and Uéverton dos Santos Souza. 2021. Hitting forbidden induced subgraphs on bounded treewidth graphs. *Information and Computation* 281 (2021), 104812. DOI: https://doi.org/10.1016/j.ic.2021.104812

[101] Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. 2020. An FPT-algorithm for recognizing k-apices of minor-closed graph classes. In *Proceedings of the 47th International Colloquium on Automata, Languages, and Programming (ICALP),* Leibniz International Proceedings in Informatics, Vol. 168, 95:1–95:20. DOI: https://doi.org/10.4230/LIPIcs.ICALP.2020.95

[102] Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. 2022. k-apices of minor-closed graph classes. II. Parameterized algorithms. *ACM Transactions on Algorithms* 18, 3 (2022), Article 21, 30 pages. DOI: https://doi.org/10.1145/3519028

[103] Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. 2023. k-apices of minor-closed graph classes. I. Bounding the obstructions. *Journal of Combinatorial Theory, Series B* 161 (2023), 180–227. DOI: https://doi.org/10.1016/j.jctb.2023.02.012

[104] Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. 2024. A more accurate view of the flat wall theorem. *Journal of Graph Theory* 107, 2 (2024), 263–297. DOI: https://doi.org/10.1002/jgt.23121

[105] Nicole Schirrmacher, Sebastian Siebertz, Giannos Stamoulis, Dimitrios M. Thilikos, and Alexandre Vigny. 2024. Model checking disjoint-paths logic on topological-minor-free graph classes. In *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS).* ACM, New York, NY, Article 68, 12 pages. DOI: https://doi.org/10.1145/3661814.3662089

[106] Nicole Schirrmacher, Sebastian Siebertz, and Alexandre Vigny. 2023. First-order logic with connectivity operators. *ACM Transactions on Computational Logic* 24, 4 (2023), 30:1–30:23. DOI: https://doi.org/10.1145/3595922

[107] Detlef Seese. 1996. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science* 6, 6 (1996), 505–526. DOI: https://doi.org/10.1017/S0960129500070079

[108] Paul D. Seymour and Robin Thomas. 1993. Graph searching and a min-max theorem for tree-width. *Journal of Combinatorial Theory, Series B* 58, 1 (1993), 22–33. DOI: https://doi.org/10.1006/jctb.1993.1027
[109] Sebastian Siebertz. 2016. *Nowhere Dense Classes of Graphs Characterisations and Algorithmic Meta-Theorems*. Ph.D. Dissertation. Universitätsverlag der TU, Berlin. DOI: https://doi.org/10.14279/depositonce-5011
[110] Jan van den Heuvel, Stephan Kreutzer, Michał Pilipczuk, Daniel A. Quiroz, Roman Rabinovich, and Sebastian Siebertz. 2017. Model-checking for successor-invariant first-order formulas on graph classes of bounded expansion. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE Computer Society, 1–11. DOI: https://doi.org/10.1109/LICS.2017.8005115
[111] Mihalis Yannakakis. 1981. Edge-deletion problems. *The SIAM Journal on Computing* 10, 2 (1981), 297–309. DOI: https://doi.org/10.1137/0210021

Appendices

A Transductions

We refer the reader to [30] for a broader discussion on logical structures and monadic second-order logic, from the viewpoint of graphs (see also [83]).

A.1 Transductions

In this subsection we define (a particular type of) transductions between structures. The definitions presented here are taken from [12] (see also [30]).

Let τ and σ be two vocabularies without constant symbols.[6] We define a *transduction* with input vocabulary τ and output vocabulary σ to be a set of pairs $(\mathfrak{A}, \mathfrak{B})$, where $\mathfrak{A}$ is a τ-structure and $\mathfrak{B}$ is a σ-structure. Given a transduction $\mathcal{I}$ with input vocabulary τ and output vocabulary σ and a τ-structure $\mathfrak{A}$, we denote by $\mathcal{I}(\mathfrak{A})$ the set of all σ-structures $\mathfrak{B}$ such that $(\mathfrak{A}, \mathfrak{B}) \in \mathcal{I}$. Notice that a transduction is a binary relation between structures that is not necessarily a function. All the transductions that we will use in our algorithms, are *deterministic*, in the sense that they are partial functions (up to isomorphism).

MSO-Transductions. We now define MSO-transductions, which are a special case of transductions that can be defined using MSO. We begin by defining three types of transductions:

— *Copying*. Let τ be a vocabulary and $k \in \mathbb{N}_{\geq 0}$. We define *$k$-copying* to be the transduction with input vocabulary τ and output vocabulary $\sigma = \tau \cup \{\mathsf{copy}, \mathsf{layer}_1, \ldots, \mathsf{layer}_k\}$, where copy is a binary relation symbol, $\mathsf{layer}_i, i \in [k]$ is a unary relation symbol, and every τ-structure $\mathfrak{A}$, outputs a σ-structure $\mathfrak{B}$, where
  - $V(\mathfrak{B})$ is the disjoint union of k copies of $V(\mathfrak{A})$,
  - for every $\mathsf{R} \in \tau$ or arity $r \geq 1$, $\mathsf{R}^{\mathfrak{B}}$ is the set of all r-tuples over $V(\mathfrak{B})$ such that all the elements of the tuple are in the same copy of $V(\mathfrak{A})$ and the original elements of the copies are in $\mathsf{R}^{\mathfrak{A}}$,
  - $\mathsf{copy}^{\mathfrak{B}}$ is the set of all pairs of elements in $V(\mathfrak{B})$ that are copies of the same element of $V(\mathfrak{A})$, and
  - for $i \in [k]$, $\mathsf{layer}_i^{\mathfrak{B}}$ is the set of all elements that belong to the i-th copy of $V(\mathfrak{A})$.

— *Coloring*. Let τ be a vocabulary and $\mathsf{C} \notin \tau$ be a unary relation symbol. We define *coloring* to be the transduction with input vocabulary τ and output vocabulary $\sigma = \tau \cup \{\mathsf{C}\}$ that, for every τ-structure $\mathfrak{A}$ and every $S \subseteq V(\mathfrak{A})$, outputs the σ-structure $\mathfrak{B}_S$, where $V(\mathfrak{B}) = V(\mathfrak{A})$, for every $\mathsf{R} \in \tau$, $\mathsf{R}^{\mathfrak{B}} = \mathsf{R}^{\mathfrak{A}}$, and $\mathsf{C}^{\mathfrak{B}} = S$.

— *Interpreting*. Let τ and σ be two vocabularies. We define *interpretation* to be the transduction with input vocabulary τ and output vocabulary σ as follows: We consider a family

[6] In this article, we define transductions between structures without constants. We can extend this definition to transductions between structures with constants with the additional "promise" that these transductions do not change the constants.

of $\mathsf{MSO}[\tau]$ formulas

$$\{\varphi_{\mathsf{dom}}, \varphi_{\mathsf{univ}}\} \cup \{\varphi_{\mathsf{R}}\}_{\mathsf{R}\in\sigma},$$

where the formula φ_{dom} is a sentence (i.e., it has no free variables), the formula φ_{univ} has one free variable, and each formula φ_{R} has as many free variables as the arity of R. The free variables in the above formulas are first-order variables. Given a τ-structure $\mathfrak{A}$ such that $\mathfrak{A} \models \varphi_{\mathsf{dom}}$, the output of the interpretation is the σ-structure $\mathfrak{B}$, where

– $V(\mathfrak{B}) = \{a \in V(\mathfrak{A}) \mid \mathfrak{A} \models \varphi_{\mathsf{univ}}(a)\}$ and
– for every $\mathsf{R} \in \sigma$ of arity $r \geq 1$, $\mathsf{R}^{\mathfrak{B}} = \{(a_1, \ldots, a_r) \in V(\mathfrak{B})^r \mid \mathfrak{A} \models \varphi_{\mathsf{R}}(a_1, \ldots, a_r)\}$.

If $\mathfrak{A} \not\models \varphi_{\mathsf{dom}}$, then the output of the interpretation is not defined. Intuitively, the formula φ_{dom} specifies the domain of the interpretation, by "filtering out" all structures that do not satisfy it. Also, the formula φ_{univ} defines the universe of the structure $\mathfrak{B}$, while the formulas φ_{R} allow us to "interpret" the relation symbols in σ.

A relation $\mathcal{I}$ between τ-structures and σ-structures is called an MSO-*transduction with input vocabulary* τ *and output vocabulary* σ if there exists a $k \in \mathbb{N}_{\geq 1}$ such that $\mathcal{I} = \mathcal{R}_k \circ \ldots \circ \mathcal{R}_1$, where, for every $i \in [k]$, $\mathcal{R}_i$ is a copying/coloring/interpreting between τ_i-structures and σ_i-structures, $\tau_1 = \tau$, and $\sigma_k = \sigma$.

The reason why we call the above relations MSO-transductions is based on the fact that the formulas we use in the definition of interpretation are formulas in $\mathsf{MSO}[\tau]$. We can define FO-*transductions* analogously, by demanding that these formulas are FO-formulas. Notice that since every FO-formula is also an MSO-formula, an FO-transduction is also an MSO-transduction.

Backwards Translation Theorem. The following result allows us to translate a question in one structure to an "equivalent" question in another structure through MSO-transductions. It is known as the *Backwards Translation Theorem* [30, Theorem 1.40] (see also [12, Lemma B.1]). We state it for sentences, i.e., formulas without free variables.

Proposition 17. *Let* $\mathcal{L}$ *be either* MSO *or* FO *and let* τ *and* σ *be vocabularies without constant symbols. Let* $\mathcal{I}$ *be an* $\mathcal{L}$*-transduction with input vocabulary* τ *and output vocabulary* σ*. If* φ *is a sentence in* $\mathcal{L}[\sigma]$*, then there is a sentence* $\psi \in \mathcal{L}[\tau]$ *such that for every* σ*-structure* $\mathfrak{B}$*, if* $\mathfrak{B} \in \mathcal{I}(\mathfrak{A})$ *for some* τ*-structure* $\mathfrak{A}$*, it holds that*

$$\mathfrak{A} \models \psi \Leftrightarrow \mathfrak{B} \models \varphi.$$

We now state the following result. Intuitively, it says that in the case of structures whose Gaifman graphs have bounded Hadwiger number, one can transduce the original structure from its Gaifman graph. This was proved in a more general setting in [17, Lemma 3.1] for the case where the Gaifman graphs have bounded star chromatic number, a property satisfied in classes of bounded expansion such as classes of bounded Hadwiger number.

Proposition 18. *Let* τ *be a vocabulary without constant symbols, let* $\mathsf{E} \notin \tau$ *be a binary relation symbol, let* $c \in \mathbb{N}$*, and let* $C \subseteq \mathbb{STR}[\tau]$*. There is an FO-transduction* $\mathcal{I}$ *from* E*-structures to* τ*-structures such that if all graphs in* $\{G_{\mathfrak{A}} \mid \mathfrak{A} \in C\}$ *have Hadwiger number at most* c*, then, if* $G = G_{\mathfrak{A}}$ *for some* $\mathfrak{A} \in C$*, it holds that* $\mathcal{I}(G) = \mathfrak{A}$*.*

At this point, we should comment that, due to Proposition 18, we can transduce every structure that is a model of a formula in $\mathsf{CMSO}^{\mathsf{tw}} \triangleright \mathsf{FO}$ from its Gaifman graph, given that the latter has bounded has Hadwiger number. This, in turn, together with Proposition 17, the fact that $\mathsf{CMSO}^{\mathsf{tw}} \triangleright \mathsf{FO} \subseteq \mathsf{CMSO}$, and the observation that any FO-transduction is also an MSO-transduction, indicates that the problem of model-checking for $\mathsf{CMSO}^{\mathsf{tw}} \triangleright \mathsf{FO}$ in general structures (whose Gaifman graphs have bounded Hadwiger number) is essentially not more general than in graphs (of bounded Hadwiger number).

A.2 Expressing Stellation and Apex-Projection as Transductions

Lemma 19. *Let τ be a vocabulary, $\mathsf{X} \notin \tau$ be a unary relation symbol, and $\mathsf{E} \notin \tau$ be a binary relation symbol.* $\mathsf{stell}_{\mathsf{X}}$ *is an MSO-transduction from $(\tau \cup \{\mathsf{X}\})$-structures to $(\tau \cup \{\mathsf{X}\})$-structures.*

Proof. We will prove that $\mathsf{stell}_{\mathsf{X}}$ is an MSO-transduction from $(\tau \cup \{\mathsf{X}\})$-structures to $(\tau \cup \{\mathsf{X}\})$-structures. Let $\mathfrak{A}$ be a $(\tau \cup \{\mathsf{X}\})$-structure. To obtain $\mathfrak{B} = \mathsf{stell}_{\mathsf{X}}(\mathfrak{A})$, we first use coloring and add a new unary predicate U in $\mathfrak{A}$ and guess an interpretation U of U in $V(\mathfrak{A})$, which corresponds to a choice of representatives, one for every $C \in \mathsf{cc}(G_{\mathfrak{A}}, \mathsf{X}^{\mathfrak{A}})$. We call $\mathfrak{A}'$ this new $(\tau \cup \{\mathsf{X}, \mathsf{U}\})$-structure. Then, we use interpretation to transform $\mathfrak{A}'$ to $\mathfrak{B}$, by setting φ_{dom} to be always true, $\varphi_{\mathsf{univ}}(\mathsf{x}) = (\mathsf{x} \in \mathsf{X} \vee \mathsf{x} \in \mathsf{U})$, $\varphi_{\mathsf{X}}(\mathsf{x}) = (\mathsf{x} \in \mathsf{U})$, and $\varphi_{\mathsf{E}}(x, y)$ asks whether there is an edge between x and y or x (resp. y) belongs to X and y (resp. x) belongs to U and x (resp. y) is adjacent to a vertex z that is in the same connected component of $G \setminus X$ as y (resp. x). □

Backwards Translating an Apex-Projected Sentence. We now aim to prove that given a vocabulary τ, an $l \in \mathbb{N}$, a collection $\mathbf{c}$ of l constant symbols, and a sentence $\sigma \in \mathsf{FO}[\tau]$, we can find a sentence $\sigma' \in \mathsf{FO}[\tau \cup \mathbf{c}]$ such that for every τ-structure $\mathfrak{A}$ and every apex-tuple $\mathbf{a}$ of $\mathfrak{A}$ of size l, $(\mathfrak{A}, \mathbf{a}) \models \sigma' \Leftrightarrow \mathsf{ap}_{\mathbf{c}}(\mathfrak{A}, \mathbf{a}) \models \sigma^{l}$. For this reason, we first prove that the function $\mathsf{ap}_{\mathbf{c}}$ is an FO-transduction and we then use Proposition 17 to obtain the desired sentence σ' (see Corollary 8). We stress that, in Section A.1, we avoided to define transductions as relations between structures of vocabularies with constant symbols, for the sake of simplicity. In our current case, we slightly abuse the definition of transductions and allow constant symbols, since the function $\mathsf{ap}_{\mathbf{c}}$ leaves the interpretation of $\mathbf{c}$ intact and therfore we can safely extend the definition of transduction and the statement of Proposition 17 to capture this case. We refer the reader to [30, Section 7.1.2] for a discussion on transductions between structures with constants.

Observation 20. *Let τ be a vocabulary, let $l \in \mathbb{N}$, let $\mathbf{c}$ be a collection of l constant symbols, and let $\tau^{\langle \mathbf{c} \rangle}$ be the constant-projection of $\tau \cup \mathbf{c}$. The function that maps every $(\tau \cup \mathbf{c})$-structure $(\mathfrak{A}, \mathbf{a})$ to the $\tau^{\langle \mathbf{c} \rangle}$-structure $ap_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})$ is an FO-transduction. Moreover, there is an FO-transduction from $\tau^{\langle \mathbf{c} \rangle}$ to $\tau \cup \mathbf{c}$ that maps $ap_{\mathbf{c}}(\mathfrak{A}, \mathbf{a})$ to $(\mathfrak{A}, \mathbf{a})$, if $G_{\mathfrak{A}}$ has bounded Hadwiger number.*

B Flat Walls Framework

Here we present the framework on flat walls that was introduced in [104]. In Section B.1 we give some additional basic definitions and in Section B.2 we define walls, subwalls, and other notions related to walls. Next, in Section B.3, we give the definitions of renditions and paintings, that are used in Section B.4 to define flatness pairs. In Section B.4, apart from the definition of flatness pairs, we present notions like influence, regularity, and tilts. Then, in Section B.5, we state Proposition 26 that is a critical ingredient of our algorithm of Theorem 3 in Section 6.2. Finally, in Section B.6, we give the definition of a canonical partition of a wall that is useful in Section 5.2.

B.1 Basic Definitions

Given a graph G, we define the *detail* of G, denoted by $\mathsf{detail}(G)$, to be the maximum among $|E(G)|$ and $|V(G)|$. Given a finite collection $\mathcal{F}$ of graphs, we set $\ell_{\mathcal{F}} = \max\{\mathsf{detail}(H) \mid H \in \mathcal{F}\}$.

Dissolutions and Subdivisions. Given a vertex $v \in V(G)$ of degree two with neighbors u and w, we define the *dissolution* of v to be the operation of deleting v and, if u and w are not adjacent, adding the edge $\{u, w\}$. Given two graphs H, G, we say that H is a *dissolution* of G if H can be obtained from G after dissolving vertices of G. Given an edge $e = \{u, v\} \in E(G)$, we define the *subdivision* of

e to be the operation of deleting e, adding a new vertex w and making it adjacent to u and v. Given two graphs H, G, we say that H is a *subdivision* of G if H can be obtained from G after subdividing edges of G.

Contractions and Minors. A graph G' is a *contraction* of a graph G, if G' can be obtained from G by a sequence of edge contractions. Given two graphs H, G, if H is a minor of G then for every vertex $v \in V(H)$ there is a set of vertices in G that are the endpoints of the edges of G contracted towards creating v. We call this set *model* of v in G.

B.2 Walls and Subwalls

Walls. Let $k, r \in \mathbb{N}$. The $(k \times r)$*-grid* is the graph whose vertex set is $[k] \times [r]$ and two vertices (i, j) and (i', j') are adjacent if and only if $|i - i'| + |j - j'| = 1$. An *elementary r-wall*, for some odd integer $r \geq 3$, is the graph obtained from a $(2r \times r)$-grid with vertices $(x, y) \in [2r] \times [r]$, 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 $\mathbb{R}^2$ 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*. Also, the vertices in the perimeter of an elementary r-wall that have degree two are called *pegs*, while the vertices $(1, 1), (2, r), (2r - 1, 1), (2r, r)$ are called *corners* (notice that the corners are also pegs).

An *r-wall* is any graph W obtained from an elementary r-wall $\bar{W}$ after subdividing edges. A graph W is a *wall* if it is an r-wall for some odd $r \geq 3$ and we refer to r as the *height* of W. Given a graph G, a *wall of G* is a subgraph of G that is a wall. We insist that, for every r-wall, the number r is always odd. See Figure B1 for an example of a 7-wall.

We call the vertices of degree three of a wall W *3-branch vertices.* A cycle of W is a *brick* (resp. the *perimeter*) of W if its 3-branch vertices are the vertices of a brick (resp. the perimeter) of $\bar{W}$. We denote by $\mathcal{C}(W)$ the set of all cycles of W. We use $D(W)$ in order to denote the perimeter of the wall W. A brick of W is *internal* if it is disjoint from $D(W)$.

Subwalls. Given an elementary r-wall $\bar{W}$, some odd $i \in \{1, 3, \ldots, 2r - 1\}$, and $i' = (i + 1)/2$, the *i'-th vertical path* of $\bar{W}$ is the one whose vertices, in order of appearance, are $(i, 1), (i, 2), (i + 1, 2), (i + 1, 3), (i, 3), (i, 4), (i + 1, 4), (i + 1, 5), (i, 5), \ldots, (i, r - 2), (i, r - 1), (i + 1, r - 1), (i + 1, r)$. Also, given some $j \in [2, r - 1]$ the *j-th horizontal path* of $\bar{W}$ is the one whose vertices, in order of appearance, are $(1, j), (2, j), \ldots, (2r, j)$.

A *vertical* (resp. *horizontal*) path of W is one that is a subdivision of a vertical (resp. horizontal) path of $\bar{W}$. Notice that the perimeter of an r-wall W is uniquely defined regardless of the choice of the elementary r-wall $\bar{W}$. A *subwall* of W is any subgraph W' of W that is an r'-wall, with $r' \leq r$, and such the vertical (resp. horizontal) paths of W' are subpaths of the vertical (resp. horizontal) paths of W.

Layers. 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 removing recursively all occurring vertices of degree one. We refer to the $(r - 1)/2$-th layer as the *inner layer* of W. The *central vertices* of an r-wall W are its two 3-branch vertices that do not belong to any of its layers and are connected by a path of W that does not intersect any layers of W. See Figure B1 for an illustration.

Central Walls. Given an r-wall W and an odd $q \in \mathbb{N}_{\geq 3}$ where $q \leq r$, we define the *central q-subwall* of W, denoted by $W^{(q)}$, to be the q-wall obtained from W after removing its first $(r - q)/2$ layers and all occurring vertices of degree one.

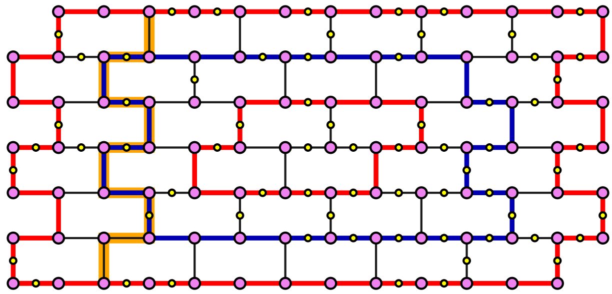

Fig. B1. A 7-wall and its three layers, depicted in alternating red and blue. The perimeter of the wall is the outermost red cycle. The pink-colored vertices of degree three are the 3-branch vertices of the wall. The orange-highlighted path is the second vertical path of the wall.

Tilts. The *interior* of a wall W is the graph obtained from W if we remove from it all edges of $D(W)$ and all vertices of $D(W)$ that have degree two in W. Given two walls W and $\tilde{W}$ of a graph G, we say that $\tilde{W}$ is a *tilt* of W if $\tilde{W}$ and W have identical interiors.

B.3 Paintings and Renditions

In this subsection we present the notions of renditions and paintings, originating in the work of Robertson and Seymour [98]. The definitions presented here were introduced by Kawarabayashi, Thomas, and Wollan [72] (see also [104]).

Paintings. A *closed* (resp. *open*) *disk* is a set homeomorphic to the set $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}$ (resp. $\{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 < 1\}$). Let Δ be a closed disk. Given a subset X of Δ, we denote its closure by $\bar{X}$ and its boundary by $\mathrm{bd}(X)$. A Δ-*painting* is a pair $\Gamma = (U, N)$ where

— N is a finite set of points of Δ,
— $N \subseteq U \subseteq \Delta$, and
— $U \setminus N$ has finitely many arcwise-connected components, called *cells*, where, for every cell c,
  – the closure $\bar{c}$ of c is a closed disk and
  – $|\tilde{c}| \leq 3$, where $\tilde{c} := \mathrm{bd}(c) \cap N$.

We use the notation $U(\Gamma) := U$, $N(\Gamma) := N$ and denote the set of cells of Γ by $C(\Gamma)$. For convenience, we may assume that each cell of Γ is an open disk of Δ. Notice that, given a Δ-painting Γ, the pair $(N(\Gamma), \{\tilde{c} \mid c \in C(\Gamma)\})$ is a hypergraph whose hyperedges have cardinality at most three and Γ can be seen as a plane embedding of this hypergraph in Δ. See Figure B2 for an example of a Δ-painting, where Δ is the disk depicted in white (bounded by the grey area), N corresponds to the blue-colored points of Δ, and the cells of U are the pink-colored regions.

Renditions. Let G be a graph and let Ω be a cyclic permutation of a subset of $V(G)$ that we denote by $V(\Omega)$. By an Ω-*rendition* of G we mean a triple (Γ, σ, π), where

(a) Γ is a Δ-painting for some closed disk Δ,
(b) $\pi : N(\Gamma) \to V(G)$ is an injection, and
(c) σ assigns to each cell $c \in C(\Gamma)$ a subgraph $\sigma(c)$ of G, such that
  (1) $G = \bigcup_{c \in C(\Gamma)} \sigma(c)$,
  (2) for distinct $c, c' \in C(\Gamma)$, $\sigma(c)$ and $\sigma(c')$ are edge-disjoint,
  (3) for every cell $c \in C(\Gamma)$, $\pi(\tilde{c}) \subseteq V(\sigma(c))$,
  (4) for every cell $c \in C(\Gamma)$, $V(\sigma(c)) \cap \bigcup_{c' \in C(\Gamma) \setminus \{c\}} V(\sigma(c')) \subseteq \pi(\tilde{c})$, and

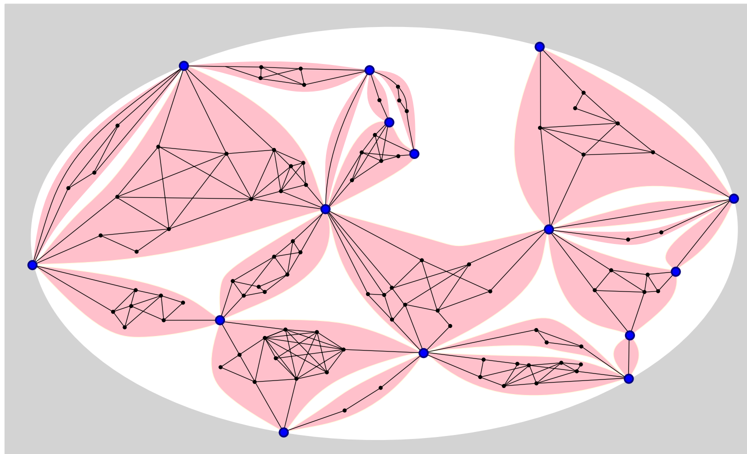

Fig. B2. A graph G together with an Ω-rendition of G.

(5) $\pi(N(\Gamma) \cap \mathrm{bd}(\Delta)) = V(\Omega)$, such that the points in $N(\Gamma) \cap \mathrm{bd}(\Delta)$ appear in $\mathrm{bd}(\Delta)$ in the same ordering as their images, via π, in Ω.

See Figure B2 for an example of an Ω-rendition of a graph G.

B.4 Flatness Pairs

In this subsection we define the notion of a flat wall, originating in the work of Robertson and Seymour [98] and later used in [72]. Here, we define flat walls as in [104].

Flat Walls. Let G be a graph and let W be an r-wall of G, for some odd integer $r \geq 3$. We say that a pair of vertex sets (P, C), where $P, C \subseteq D(W)$, is a *choice of pegs and corners for* W if W is the subdivision of an elementary r-wall $\bar{W}$ where P and C are the pegs and the corners of $\bar{W}$, respectively (clearly, $C \subseteq P$). To get more intuition, notice that a wall W can occur in several ways from the elementary wall $\bar{W}$, depending on the way the vertices in the perimeter of $\bar{W}$ are subdivided. Each of them gives a different selection (P, C) of pegs and corners of W.

We say that W is a *flat r-wall* of G if there is a separation (X, Y) of G and a choice (P, C) of pegs and corners for W such that:

- $V(W) \subseteq Y$,
- $P \subseteq X \cap Y \subseteq V(D(W))$, and
- if Ω is the cyclic ordering of the vertices $X \cap Y$ as they appear in $D(W)$, then there exists an Ω-rendition (Γ, σ, π) of $G[Y]$.

We say that W is a *flat wall* of G if it is a flat r-wall for some odd integer $r \geq 3$.

Flatness Pairs. Given the above, we say that the choice of the 7-tuple $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ *certifies that W is a flat wall of G*. We call the pair $(W, \mathfrak{R})$ a *flatness pair* of G and define the *height* of the pair $(W, \mathfrak{R})$ to be the height of W. We use the term *cell of* $\mathfrak{R}$ in order to refer to the cells of Γ.

We call the graph $G[Y]$ the $\mathfrak{R}$-*compass* of W in G, denoted by $\mathrm{compass}_{\mathfrak{R}}(W)$. It is easy to see that there is a connected component of $\mathrm{compass}_{\mathfrak{R}}(W)$ that contains the wall W as a subgraph. We can assume that $\mathrm{compass}_{\mathfrak{R}}(W)$ is connected, updating $\mathfrak{R}$ by removing from Y the vertices of all the connected components of $\mathrm{compass}_{\mathfrak{R}}(W)$ except of the one that contains W and including them in X (Γ, σ, π can also be easily modified according to the removal of the aforementioned vertices from Y). We define the *flaps* of the wall W in $\mathfrak{R}$ as $\mathsf{flaps}_{\mathfrak{R}}(W) := \{\sigma(c) \mid c \in C(\Gamma)\}$. Given a flap $F \in \mathsf{flaps}_{\mathfrak{R}}(W)$, we define its *base* as $\partial F := V(F) \cap \pi(N(\Gamma))$. A cell c of $\mathfrak{R}$ is *untidy* if $\pi(\tilde{c})$ contains a vertex x of W such that two of the edges of W that are incident to x are edges of $\sigma(c)$. Notice that if c is untidy then $|\tilde{c}| = 3$. A cell c of $\mathfrak{R}$ is *tidy* if it is not untidy. The notion of tidy/untidy cell as well as the notions that we present in the rest of this subsection have been introduced in [104].

Cell Classification. Given a cycle C of $\mathsf{compass}_{\mathfrak{R}}(W)$, we say that C is *$\mathfrak{R}$-normal* if it is not a subgraph of a flap $F \in \mathsf{flaps}_{\mathfrak{R}}(W)$. Given an $\mathfrak{R}$-normal cycle C of $\mathsf{compass}_{\mathfrak{R}}(W)$, we call a cell c of $\mathfrak{R}$ *C-perimetric* if $\sigma(c)$ contains some edge of C. Since every C-perimetric cell c contains some edge of C and $|\partial\sigma(c)| \leq 3$, we observe the following.

Observation 21. *For every pair (C, C') of $\mathfrak{R}$-normal cycles of $\mathsf{compass}_{\mathfrak{R}}(W)$ such that $V(C) \cap V(C') = \emptyset$, there is no cell of $\mathfrak{R}$ that is both C-perimetric and C'-perimetric.*

Notice that if c is C-perimetric, then $\pi(\tilde{c})$ contains two points $p, q \in N(\Gamma)$ such that $\pi(p)$ and $\pi(q)$ are vertices of C where one, say P_c^{in}, of the two $(\pi(p), \pi(q))$-subpaths of C is a subgraph of $\sigma(c)$ and the other, denoted by P_c^{out}, $(\pi(p), \pi(q))$-subpath contains at most one internal vertex of $\sigma(c)$, which should be the (unique) vertex z in $\partial\sigma(c) \setminus \{\pi(p), \pi(q)\}$. We pick a (p, q)-arc A_c in $\hat{c} := c \cup \tilde{c}$ such that $\pi^{-1}(z) \in A_c$ if and only if P_c^{in} contains the vertex z as an internal vertex.

We consider the circle $K_C = \bigcup\{A_c \mid c \text{ is a } C\text{- perimetric cell of } \mathfrak{R}\}$ and we denote by Δ_C the closed disk bounded by K_C that is contained in Δ. A cell c of $\mathfrak{R}$ is called *C-internal* if $c \subseteq \Delta_C$ and is called *C-external* if $\Delta_C \cap c = \emptyset$. Notice that the cells of $\mathfrak{R}$ are partitioned into C-internal, C-perimetric, and C-external cells.

Let c be a tidy C-perimetric cell of $\mathfrak{R}$ where $|\tilde{c}| = 3$. Notice that $c \setminus A_c$ has two arcwise-connected components and one of them is an open disk D_c that is a subset of Δ_C. If the closure $\overline{D}_c$ of D_c contains only two points of $\tilde{c}$ then we call the cell c *C-marginal.* See Figure B3 for a figure illustrating the above notions. We refer the reader to [104] for more figures.

Influence. For every $\mathfrak{R}$-normal cycle C of $\mathsf{compass}_{\mathfrak{R}}(W)$ we define the set

$$\mathsf{influence}_{\mathfrak{R}}(C) = \{\sigma(c) \mid \text{c is a cell of } \mathfrak{R} \text{ that is not C-external}\}.$$

A wall W' of $\mathsf{compass}_{\mathfrak{R}}(W)$ is *$\mathfrak{R}$-normal* if $D(W')$ is $\mathfrak{R}$-normal. Notice that every wall of W (and hence every subwall of W) is an $\mathfrak{R}$-normal wall of $\mathsf{compass}_{\mathfrak{R}}(W)$. We denote by $\mathcal{S}_{\mathfrak{R}}(W)$ the set of all $\mathfrak{R}$-normal walls of $\mathsf{compass}_{\mathfrak{R}}(W)$. Given a wall $W' \in \mathcal{S}_{\mathfrak{R}}(W)$ and a cell c of $\mathfrak{R}$, we say that c is *W'-perimetric/internal/external/marginal* if c is $D(W')$-perimetric/internal/external/marginal, respectively. We also use $K_{W'}$, $\Delta_{W'}$, $\mathsf{influence}_{\mathfrak{R}}(W')$ as shortcuts for $K_{D(W')}$, $\Delta_{D(W')}$, $\mathsf{influence}_{\mathfrak{R}}(D(W'))$, respectively.

Regular Flatness Pairs. We call a flatness pair $(W, \mathfrak{R})$ of a graph G *regular* if none of its cells is W-external, W-marginal, or untidy.

Tilts of Flatness Pairs. Let $(W, \mathfrak{R})$ and $(\tilde{W}', \tilde{\mathfrak{R}}')$ be two flatness pairs of a graph G and let $W' \in \mathcal{S}_{\mathfrak{R}}(W)$. We assume that $\mathfrak{R} = (X, Y, P, C, \Gamma, \sigma, \pi)$ and $\tilde{\mathfrak{R}}' = (X', Y', P', C', \Gamma', \sigma', \pi')$. We say that $(\tilde{W}', \tilde{\mathfrak{R}}')$ is a *W'-tilt* of $(W, \mathfrak{R})$ if

- $\tilde{\mathfrak{R}}'$ does not have $\tilde{W}'$-external cells,
- $\tilde{W}'$ is a tilt of W',
- the set of $\tilde{W}'$-internal cells of $\tilde{\mathfrak{R}}'$ is the same as the set of W'-internal cells of $\mathfrak{R}$ and their images via σ' and σ are also the same,
- $\mathsf{compass}_{\tilde{\mathfrak{R}}'}(\tilde{W}')$ is a subgraph of $\bigcup\mathsf{influence}_{\mathfrak{R}}(W')$, and
- if c is a cell in $C(\Gamma') \setminus C(\Gamma)$, then $|\tilde{c}| \leq 2$.

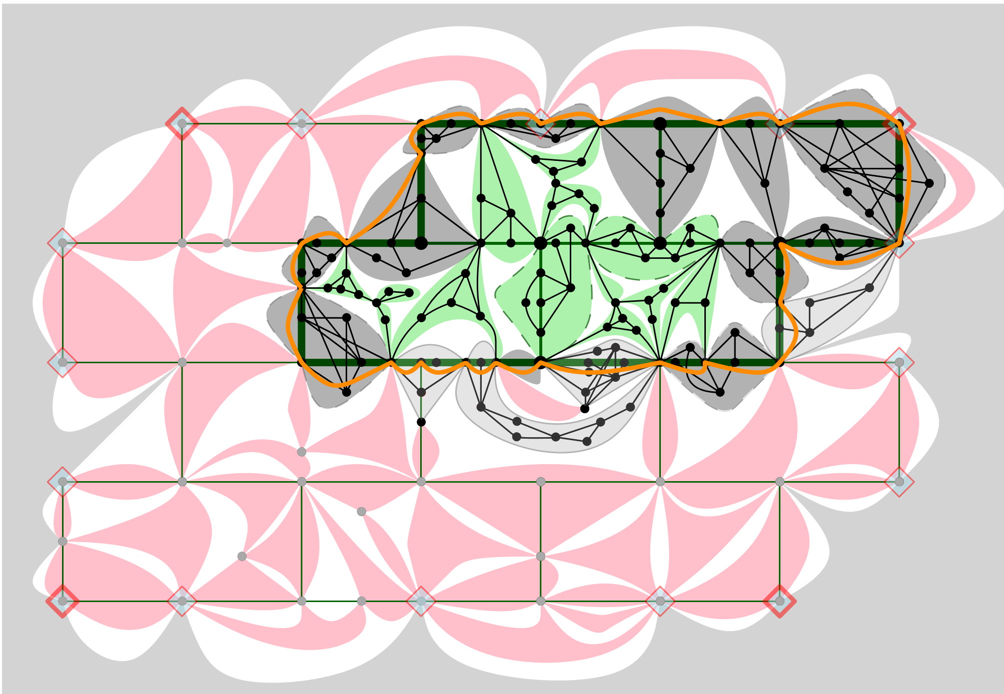

Fig. B3. This picture is taken from [104]. It depicts a flat wall W in a graph G, the painting of a rendition $\mathfrak{R}$ certifying its flatness, a subwall W' of W, of height three, which is $\mathfrak{R}$-normal, and the $\mathfrak{R}$-flaps of W, that correspond to either W'-perimetric (depicted in grey) or W'-internal cells (depicted in green). The circle $K_{W'}$ is the fat orange cycle. The W'-marginal cells are depicted in light grey and the untidy cells are those with dashed boundary.

The next observation follows from the third item above and the fact that the cells corresponding to flaps containing a central vertex of W' are all internal (recall that the height of a wall is always at least three).

Observation 22. *Let $(W, \mathfrak{R})$ be a flatness pair of a graph G and $W' \in \mathcal{S}_{\mathfrak{R}}(W)$. For every W'-tilt $(\tilde{W}', \tilde{\mathfrak{R}}')$ of $(W, \mathfrak{R})$, the central vertices of W' belong to the vertex set of $\text{compass}_{\tilde{\mathfrak{R}}'}(\tilde{W}')$.*

Also, given a regular flatness pair $(W, \mathfrak{R})$ of a graph G and a $W' \in \mathcal{S}_{\mathfrak{R}}(W)$, for every W'-tilt $(\tilde{W}', \tilde{\mathfrak{R}}')$ of $(W, \mathfrak{R})$, by definition, none of its cells is $\tilde{W}'$-external, $\tilde{W}'$-marginal, or untidy – thus, $(\tilde{W}', \tilde{\mathfrak{R}}')$ is regular. Therefore, regularity of a flatness pair is a property that its tilts "inherit."

Observation 23. *If $(W, \mathfrak{R})$ is a regular flatness pair of a graph G, then for every $W' \in \mathcal{S}_{\mathfrak{R}}(W)$, every W'-tilt of $(W, \mathfrak{R})$ is also regular.*

We next present one of the two main results of [104] (see [104, Theorem 5]).

Proposition 24. *There exists an algorithm that given a graph G, a flatness pair $(W, \mathfrak{R})$ of G, and a wall $W' \in \mathcal{S}_{\mathfrak{R}}(W)$, outputs a W'-tilt of $(W, \mathfrak{R})$ in time $O(n + m)$.*

We conclude this subsection with the Flat Wall theorem and, in particular, the version proved by Chuzhoy [26], restated in our framework (see [104, Proposition 7]).

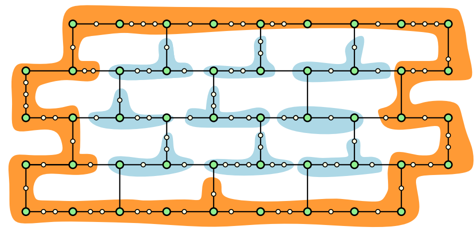

Fig. B4. A 5-wall and its canonical partition $\mathcal{Q}$. The orange bag is the external bag Q_{ext}.

Proposition 25. *There exist two functions $f_5 : \mathbb{N} \to \mathbb{N}$ and $f_6 : \mathbb{N} \to \mathbb{N}$, where the images of f_5 are odd numbers, such that if $r \in \mathbb{N}_{\geq 3}$ is an odd integer, $t \in \mathbb{N}_{\geq 1}$, G is a graph that does not contain K_t as a minor, and W is an $f_5(t) \cdot r$-wall of G, then there is a set $A \subseteq V(G)$ with $|A| \leq f_6(t)$ and a flatness pair $(\tilde{W}', \tilde{\mathfrak{R}}')$ of $G \setminus A$ of height r. Moreover, $f_5(t) = O(t^2)$ and $f_6(t) = t - 5$.*

B.5 Flat Walls with Compasses of Bounded Treewidth

The following result was proved in [104, Theorem 8]. It is a version of the Flat Wall theorem, originally proved in [98]. The proof in [104, Theorem 8] is strongly based on the proof of an improved version of the Flat Wall theorem given by of Kawarabayashi, Thomas, and Wollan [72] (see also [26, 56]).

Proposition 26. *There is a function $f_7 : \mathbb{N} \to \mathbb{N}$ and an algorithm that receives as input a graph G, an odd integer $r \geq 3$, and a $t \in \mathbb{N}_{\geq 1}$, and outputs, in time $2^{O_t(r^2)} \cdot n$, one of the following:*

- *a report that K_t is a minor of G,*
- *a tree decomposition of G of width at most $f_7(t) \cdot r$, or*
- *a set $A \subseteq V(G)$, where $|A| \leq f_6(t)$, a regular flatness pair $(W, \mathfrak{R})$ of $G \setminus A$ of height r, and a tree decomposition of the $\mathfrak{R}$-compass of W of width at most $f_7(t) \cdot r$. (Here $f_6(t)$ is the function of Proposition 25 and $f_7(t) = 2^{O(t^2 \log t)}$.)*

B.6 Canonical Partitions

Canonical Partitions. Let $r \geq 3$ be an odd integer, let W be an r-wall, and let $P_1, \ldots, P_r$ (resp. $L_1, \ldots, L_r$) be its vertical (resp. horizontal) paths. For every even (resp. odd) $i \in [2, r-1]$ and every $j \in [2, r-1]$, we define $A^{(i,j)}$ to be the subpath of P_i that starts from a vertex of $P_i \cap L_j$ and finishes at a neighbor of a vertex in L_{j+1} (resp. L_{j-1}), such that $P_i \cap L_j \subseteq A^{(i,j)}$ and $A^{(i,j)}$ does not intersect L_{j+1} (resp. L_{j-1}). Similarly, for every $i, j \in [2, r-1]$, we define $B^{(i,j)}$ to be the subpath of L_j that starts from a vertex of $P_i \cap L_j$ and finishes at a neighbor of a vertex in P_{i-1}, such that $P_i \cap L_j \subseteq A^{(i,j)}$ and $A^{(i,j)}$ does not intersect P_{i-1}.

For every $i, j \in [2, r-1]$, we denote by $Q^{(i,j)}$ the graph $A^{(i,j)} \cup B^{(i,j)}$ and by Q_{ext} the graph $W \setminus \bigcup_{i,j \in [2,r-1]} Q_{i,j}$. Now consider the collection $\mathcal{Q} = \{Q_{\text{ext}}\} \cup \{Q_{i,j} \mid i, j \in [2, r-1]\}$ and observe that the graphs in $\mathcal{Q}$ are connected subgraphs of W and their vertex sets form a partition of $V(W)$. We call $\mathcal{Q}$ the *canonical partition* of W. Also, we call every $Q_{i,j}$, for $i, j \in [2, r-1]$, an *internal bag* of $\mathcal{Q}$, while we refer to Q_{ext} as the *external bag* of $\mathcal{Q}$. See Figure B4 for an illustration of the notions defined above.

Let $(W, \mathfrak{R})$ be a flatness pair of a graph G. Consider the canonical partition $\mathcal{Q}$ of W. We enhance the graphs of $\mathcal{Q}$ so to include in them all the vertices of G by applying the following procedure. We set $\tilde{\mathcal{Q}} := \mathcal{Q}$ and, as long as there is a vertex $x \in V(\text{compass}_{\mathfrak{R}}(W)) \setminus V(\bigcup \tilde{\mathcal{Q}})$ that is adjacent to a vertex of a graph $Q \in \tilde{\mathcal{Q}}$, update $\tilde{\mathcal{Q}} := \tilde{\mathcal{Q}} \setminus \{Q\} \cup \{\tilde{Q}\}$, where $\tilde{Q} = \text{compass}_{\mathfrak{R}}(W)[\{x\} \cup V(Q)]$.

Since $\text{compass}_{\mathfrak{R}}(W)$ is a connected graph, in this way we define a partition of the vertices of $\text{compass}_{\mathfrak{R}}(W)$ into subsets inducing connected graphs. We call the $\tilde{Q} \in \tilde{\mathcal{Q}}$ that contains Q_{ext} as a subgraph the *external bag* of $\tilde{\mathcal{Q}}$, and we denote it by $\tilde{Q}_{\text{ext}}$, while we call *internal bags* of $\tilde{\mathcal{Q}}$ all graphs in $\tilde{\mathcal{Q}} \setminus \{\tilde{Q}_{\text{ext}}\}$. Moreover, we enhance $\tilde{\mathcal{Q}}$ by adding all vertices of $G \setminus V(\text{compass}_{\mathfrak{R}}(W))$ in its external bag, i.e., by updating $\tilde{Q}_{\text{ext}} := G[V(\tilde{Q}_{\text{ext}}) \cup V(G \setminus V(\text{compass}_{\mathfrak{R}}(W)))]$. We call such a partition $\tilde{\mathcal{Q}}$ a $(W, \mathfrak{R})$*-canonical partition of* G. Notice that a $(W, \mathfrak{R})$-canonical partition of G is not unique, since the graphs in $\mathcal{Q}$ can be "expanded" arbitrarily when introducing vertex x. We stress that every internal bag of a $(W, \mathfrak{R})$-canonical partition of G contains vertices of at most three bricks of W.

Received 25 May 2023; revised 21 June 2024; accepted 21 August 2024