First-Order Model Checking on Monadically Dependent Graph Classes*

Jan Dreier
HPI Potsdam, Germany

Nikolas Mählmann
University of Warsaw, Poland

Szymon Toruńczyk
University of Warsaw, Poland

Abstract

Monadically dependent graph classes are those in which first-order logic cannot encode all graphs. They include all nowhere dense classes and all classes of bounded twin-width.

We show that first-order model checking is fixed-parameter tractable on every monadically dependent graph class. Together with the known hardness result for hereditary classes that are not monadically dependent, this gives a complete characterization of the hereditary graph classes with fixed-parameter tractable first-order model checking. This resolves a central open problem in algorithmic model theory.

Our proof combines ideas from structural graph theory (graph decompositions), finite model theory (locality of first-order logic), and model theory (counting types). In particular, we show that in monadically dependent classes, tuples of fixed length have only $|U|^{O(1)}$ types of fixed quantifier rank over any set of parameters U.

*JD: Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – project number 563448077. NM and ST received funding from the European Research Council (ERC) with grant agreement No. 101126229 – BUKA. NM also received funding from the European Union through an ERA Fellowship with grant agreement No. 101334340 – LoCoMoDe.

Contents

1 Introduction 1

2 Technical overview 5

3 Preliminaries 10

4 Proof of the main theorem 12

I Number of types 17

5 Closures 17

6 Number of types over a set of parameters 26

II Definable construction sequences 28

7 Merge-width with definable stages 30

8 Construction sequences with locally definable partial graphs 34

III Model checking with type trees 38

9 Preliminaries 38

10 Distance logic 38

11 Scatter types and far trees 43

12 Lifting type trees through merges and resolves 46

13 The model checking algorithm 50

AI methodology 52

References 52

Appendix 57

A Forbidden patterns 57

B Corollary of Gaifman's locality theorem 60

C Locally definable construction sequences in nowhere dense classes 62

1 Introduction

Instead of solving a single algorithmic problem, *algorithmic meta-theorems* establish the tractability of whole classes of problems, in a uniform way. The targeted problem class is commonly defined using logic, and hence many algorithmic meta-theorems fit the following template.

> *Every problem expressible in* ⟨**logic**⟩ *can be solved in time* ⟨**runtime**⟩ *on every input satisfying* ⟨**tameness condition**⟩.

A prototypical example is Courcelle's Theorem, which states that every problem definable in *monadic second-order logic* is solvable in *linear time* on graphs with constant *tree-width* [Cou90]. Similarly, a result of Frick and Grohe implies that every graph problem definable in *first-order logic* (FO) can be decided in *linear time* on all *planar graphs* [FG01b]. In particular, the *(first-order) model checking problem* – does a given graph G satisfy a given FO sentence φ? – is *fixed-parameter tractable* (fpt) on planar graphs. That is, it is solvable[1] in time $f(|\varphi|) \cdot |V(G)|^c$, for some function f and constant c (with $c = 1$ in this specific case). Many natural graph problems are expressible in FO. For example, the sentence

$$\exists x_1 \dots \exists x_k \forall y \bigvee_i \big(E(y, x_i) \vee (y = x_i)\big)$$

expresses k-Dominating-Set: the existence of k vertices such that every vertex is either adjacent or equal to at least one of them. Further expressible problems include k-Independent-Set, k-Clique, and H-Subgraph-Isomorphism. With no further restrictions made on the input graphs, these problems appear to be highly intractable. Indeed, model checking on general graphs is believed not to be fpt; this is equivalent to the FPT≠AW[⋆] conjecture. This raises the question – first posed by Grohe [Gro07, Section 8.2] – on which graph classes[2] the problem becomes fpt. After decades of research, *monadic dependence* has crystallized as the conjectured tractability limit for *hereditary*[3] classes (see, e.g., [ALS16; Gaj+20a; Bon+22c; DMS23]).

Model Checking Conjecture. *A hereditary graph class admits fixed-parameter tractable first-order model checking if and only if it is monadically dependent (assuming* FPT≠AW[⋆]*).*

Monadic dependence is a model-theoretic tameness notion that originates in Shelah's classification theory [BS85]. It subsumes many tameness notions from structural graph theory, such as planarity, excluding a minor, bounded degree, bounded expansion, nowhere denseness, and bounded clique-width, twin-width, and merge-width. See Figure 1 for a hierarchy of included graph classes. For hereditary graph classes, monadic dependence is characterized as follows [DMT24]. A hereditary graph class $\mathscr{C}$ is *monadically independent* if there is some formula $\varphi(x, y)$ with the following property: for every graph H there is some graph $G \in \mathscr{C}$ such that H is an induced subgraph of the graph with vertices $V(G)$ and edges $\{uv : \varphi(u, v)$ or $\varphi(v, u)$ holds in $G\}$. Intuitively, this means that every graph can be encoded by φ in $\mathscr{C}$. Then $\mathscr{C}$ is *monadically dependent* if it is not monadically independent.

In a previous paper [DMT24] we resolved the hardness direction of the Model Checking Conjecture, by showing that model checking is AW[⋆]-hard on every hereditary graph class that is not monadically dependent.

[1] Throughout this paper, we use the uniform – but not strongly uniform – version of fpt, and of fpt reductions. That means that the function f need not be computable, but there is a single algorithm which solves the problem, independent of the parameter (φ in the case of model checking).

[2] A *graph class* is a (usually infinite) set of graphs.

[3] A graph class is *hereditary* if it is closed under taking induced subgraphs, i.e., deleting vertices.

1.1 Our results

In this paper, we confirm the Model Checking Conjecture by proving its tractability direction.

Theorem 1.1. *There is an algorithm which solves the model checking problem for first-order logic on graphs – does a given graph G satisfy a given first-order sentence φ? – with the following guarantee. For every monadically dependent graph class $\mathscr{C}$ and real $\varepsilon > 0$, the algorithm runs in time $O_{\mathscr{C},\varphi,\varepsilon}(n^{3+\varepsilon})$ for n-vertex graphs $G \in \mathscr{C}$.*

In particular, this gives fpt first-order model checking on all classes of bounded twin-width and bounded merge-width, without requiring a decomposition as part of the input.

We remark that our algorithm is uniform: it takes both G and φ as input, and does not depend on the graph class $\mathscr{C}$, which only appears in the running-time guarantee. Our result lifts from graphs to monadically dependent classes of binary structures – structures over a fixed signature consisting of binary relation symbols – using a standard encoding (see, e.g., Lemma 19.11 in the full version of [DMT24]).

Our algorithm is based on dynamic programming on graph decompositions called *construction sequences* of *almost bounded merge-width*. In a companion paper with Michał Pilipczuk [Dre+26], we prove that they exist and can be computed efficiently on monadically dependent graph classes.

Existing model checking algorithms for classes of bounded twin- or merge-width have a tower-of-exponentials dependence on the width. This prevents their direct application to decompositions of *almost* bounded merge-width, where the width is only bounded by a subpolynomial function $n^{o(1)}$ of the input size n, rather than a constant. To overcome this obstacle, we make the running time depend on the number of *types* that are realized in the graph. We control this number using a type-counting result for monadically dependent classes, which we believe to be of independent interest.

This result provides polynomial bounds on the number of quantifier rank q types $\mathrm{tp}_q^G(\bar{v}/U)$ of k-tuples $\bar{v}$ over a set U of parameters. The *type* $\mathrm{tp}_q^G(\bar{v}/U)$ is the set of formulas of quantifier rank at most q, with parameters from U, which hold of the tuple $\bar{v}$ in G (see Section 3 for the precise definition). For $q = 0$, the type $\mathrm{tp}_0^G(v/U)$ of a vertex $v \notin U$ represents its neighborhood $N(v) \cap U$ in U. In general graphs, there could be $2^{|U|}$ many such neighborhoods, but the fundamental Sauer–Shelah lemma [Sau72; She72] bounds their number by $(|U|+1)^d$, where d is the VC-dimension of the graph G. The VC-dimension d is bounded uniformly in every monadically dependent class. For $q = 1$, the type $\mathrm{tp}_1^G(v/U)$ of a vertex v determines $\{N(w) \cap U : w \in N(v)\}$, the set of neighborhoods in U of the neighbors of v. In general graphs there could be $2^{2^{|U|}}$ many such sets, and so on for higher q: the number of q-types over a set U can reach a tower of height $q+1$. The following result, which is the cornerstone of our approach, can thus be seen as an extension of the Sauer–Shelah lemma to higher quantifier ranks, in monadically dependent graph classes.

Theorem 1.2 (Simplified form of Theorem 4.4)**.** *Let $\mathscr{C}$ be a monadically dependent graph class and $k, q \in \mathbb{N}$. There is some $d \in \mathbb{N}$ such that for every (possibly infinite) graph $G \in \mathscr{C}$, and finite subset $U \subseteq V(G)$ with $|U| > 1$ we have*

$$\left|\{\mathrm{tp}_q^G(\bar{v}/U) : \bar{v} \in V(G)^k\}\right| \leqslant |U|^d.$$

Theorem 1.2 is a finitary analogue of a fundamental result in stability theory due to Shelah [She69], which bounds the number of complete types (of unbounded quantifier rank) over

an infinite set U, by $|U|^\kappa$ in every *stable* class of structures, for some cardinal κ depending on the class. Its proof blends ideas from [PST18] and [Gaj+23].

A technical overview of the proofs of the main results is presented in Section 2. All our results are formalized in Lean. At `https://laxarchive.org/lax-185574/`, each definition, lemma, and theorem is annotated with its proven Lean counterpart.

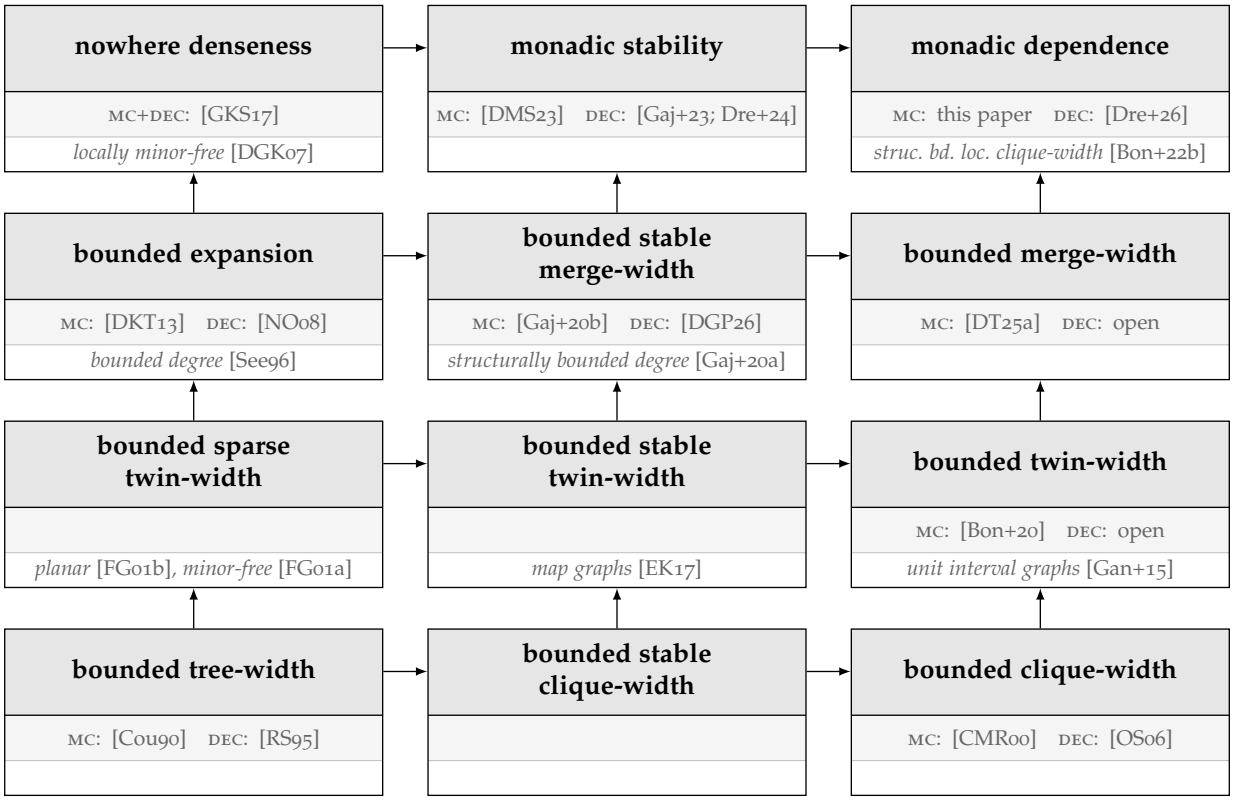


Figure 1: A hierarchy of graph class properties that admit fpt first-order model checking. Arrows represent implications between properties. In particular, *bounded tree-width* is the most restricted class property in this hierarchy and *monadic dependence* is the most general one. In each box where applicable, MC cites the fpt model checking result which assumes a suitable decomposition is given, while DEC cites the result providing such decompositions efficiently; DEC: open means that no algorithm for computing such a decomposition is known. Example graph classes are listed in *italics*, each under the most restricted property that subsumes it, together with a citation to an fpt model checking algorithm for that class.

1.2 Related work

Theorem 1.1 completes a long line of research on first-order model checking on hereditary graph classes. An overview of previous model checking results is given in Figure 1. In particular, the Model Checking Conjecture was previously known to hold in the following three settings, which we discuss below:

- in *monotone*[4] graph classes, [GKS17]
  where monadic dependence coincides with *nowhere denseness*; [AA14; PZ78]

[4] A graph class is *monotone* if it is closed under taking subgraphs, i.e., deleting vertices and edges.

- in *edge-stable*[5] graph classes, [DMS23; Dre+24]
  where monadic dependence coincides with *monadic stability*; [Neš+21]
- in classes of *ordered graphs*[6], [Bon+22c; Bon+20]
  where monadic dependence coincides with *bounded twin-width*. [Bon+22c]

Sparsity. The first model checking result in the sparse regime is the seminal Courcelle's Theorem, which established tractability for more expressive monadic second-order logic on all classes of bounded tree-width [Cou90]. Subsequently, Seese initiated the study of first-order model checking by showing tractability on all classes of bounded degree [See96]. This research direction gained momentum through the work of Grohe and his coauthors [FG01b; FG01a; DGK07]; see Grohe's 2007 survey [Gro07] for an account of these early developments. This line of research (see also [DK09a; DKT13; GK11]) culminated in the model checking algorithm of Grohe, Kreutzer, and Siebertz [GKS17] for *nowhere dense* classes. This notion is one of the cornerstones of the *sparsity theory* by Nešetřil and Ossona de Mendez [NO11; NM12] that shaped the field. For monotone classes, this result is tight: first-order model checking is AW[⋆]-hard on every monotone class that is not nowhere dense [DKT13; DK09b].

With the sparse classes resolved, the research focus has shifted towards possibly dense classes, and two parallel research directions emerged.

Stability. One research direction aims to lift the theory of nowhere dense classes to dense graphs. Here a key notion is that of *monadic stability*, a restriction of monadic dependence that also originates in model theory [BS85]. The study of the model checking problem has revealed multiple new characterizations of monadic stability, such as *flip-flatness* [Dre+23], *flipper games* [Gaj+23], and *low shrub-depth covers* [Bra+22]. They subsequently led to an fpt algorithm for first-order model checking for monadically stable classes [DMS23; Dre+24]. This result is tight in hereditary classes that are additionally edge-stable [Dre+24; Neš+21]. The ideas from this line of research have been further developed into combinatorial characterizations of monadically dependent graph classes such as *flip-breakability* [DMT24] and *flip-separability* [Bon+25]. This also includes the characterization by forbidden induced subgraphs, which settled the hardness part of the Model Checking Conjecture for all hereditary graph classes [DMT24].

Clique-width, twin-width, and merge-width. A second, parallel line of research has pushed the tractability frontier of first-order model checking through the development of new graph decompositions, often accompanied by new graph width parameters. Mirroring tree-width in the sparse setting, it was shown early on by Courcelle, Makowsky, and Rotics that also classes of bounded *clique-width* admit fpt monadic second-order model checking [CMR00].

Twenty years later, inspired by the work of Guillemot and Marx on the permutation pattern problem [GM14], Bonnet, Kim, Thomassé, and Watrigant introduced the more general classes of bounded *twin-width*, and showed that they admit fpt first-order model checking [Bon+20]. Their algorithm is conditional in the sense that it requires as an additional input a decomposition witnessing the twin-width bound. As of today it is not known how to efficiently compute or

[5] A graph class is *edge-stable* if it excludes some semi-induced half-graph, i.e., there is $k \in \mathbb{N}$ such that no graph in the class contains vertices $a_1, \dots, a_k, b_1, \dots, b_k$ where a_i and b_j are adjacent if and only if $i \leqslant j$. Equivalently, the edge relation is a stable formula on the class in the sense of model theory.

[6] An *ordered graph* is a graph equipped with a total order on its vertex set, which formulas can access via an additional binary relation symbol $<$.

approximate these decompositions in general graphs. However, they can be efficiently approximated in many specific cases of interest [Bon+22a], in particular, in hereditary classes of ordered graphs [Bon+22c]. There, bounded twin-width forms the exact tractability limit for first-order model checking.

Inspired by twin-width and the related flip-width parameters [Tor23], Dreier and Toruńczyk introduced *merge-width*, a family of graph parameters based on decompositions called *construction sequences* [DT25a]. A class has *bounded merge-width* if for every radius r, every graph in the class has a construction sequence of radius-r width bounded by a constant. Again, those classes admit fpt first-order model checking under the condition that a construction sequence is given. Crucial to our paper is the following relaxation of merge-width, called *almost bounded merge-width*, which replaces the constant bound on the radius-r width by $n^{o(1)}$, where n is the vertex count of the graph. Dreier and Toruńczyk proved that hereditary classes of almost bounded merge-width are monadically dependent, and conjectured the converse [DT25a]. The converse was first confirmed for radius one [Dre+27], and then in full [Dre+26]. The proof of [Dre+26] is algorithmic: given a graph from a monadically dependent class, it computes in polynomial time a construction sequence of radius-r width $n^{o(1)}$, for every fixed r. Construction sequences of almost bounded merge-width form the basis for our model checking algorithm.

Note that it remains an open question how to approximate decompositions of constant width for classes with constant twin- and merge-width. Thus, up until now and even in the light of [Dre+26], the previous model checking results for classes of bounded twin- and merge-width have remained conditional. Therefore, our new algorithm marks the first time that unconditional model checking for all classes of bounded twin- and merge-width is established.

2 Technical overview

In the related work section, we surveyed two approaches to the model checking problem in previous results. Our algorithm combines ideas from both approaches:

- decomposition-based methods, to do dynamic programming along a construction sequence with almost bounded merge-width; and
- stability-theoretic bounds on the *number of types*, allowing us to bound the runtime of the algorithm.

In this section, we give a high-level overview of the algorithm and the interplay between the two components. We start by introducing the relevant definition.

Merge-width. The merge-width parameters are defined through decompositions called *construction sequences* [DT25a]. We can define them as follows.

- A *partial graph* H consists of a finite vertex set $V = V(H)$, two disjoint sets $E(H), N(H) \subseteq \binom{V}{2}$ of *edges* and *non-edges*, and a partition of V whose parts we treat as colors. The pairs in $R(H) \coloneqq E(H) \cup N(H)$ are *resolved*, and all other pairs are *unresolved*. We view H as a logical structure with binary relations E and N and a unary predicate for every color. Distances dist^H and balls $\mathrm{Ball}^H_r(v)$ are measured in the Gaifman graph $(V, R(H))$, the graph of resolved pairs.
- For two distinct parts P, Q of H, the *merge* of P and Q replaces them by the single part $P \cup Q$. For two parts P, Q of H (possibly $P = Q$), the *positive resolve* of P and Q adds all

unresolved pairs uv with $u \in P$ and $v \in Q$ to $E(H)$, while the *negative resolve* adds them to $N(H)$.

- A *construction sequence* of a graph $G = (V, E)$ is a sequence $G_1, \dots, G_\lambda$ of partial graphs with vertex set V, such that G_1 has no resolved pairs and only singleton parts, every G_{t+1} is obtained from G_t by one merge or one resolve, and, finally, G_λ has a single part, $E(G_\lambda) = E$, and $N(G_\lambda) = \binom{V}{2} - E$.
- The *radius-r width* of the sequence is the maximum, over all $t \in [\lambda]$ and $v \in V$, of the number of parts of G_t that intersect $\mathrm{Ball}_r^{G_t}(v)$. The *radius-r merge-width* of G is the minimum radius-r width over all construction sequences of G.

In a companion paper, together with Michał Pilipczuk, we prove that monadically dependent classes have *almost bounded merge-width* [Dre+26], and that witnessing construction sequences can be computed efficiently. These construction sequences form the backbone of our model checking algorithm.

Theorem 2.1 ([Dre+26]). *Let $\mathscr{C}$ be a monadically dependent graph class and $r \in \mathbb{N}$. Every n-vertex graph $G \in \mathscr{C}$ has radius-r merge-width $n^{o(1)}$, and a witnessing construction sequence can be computed in time $O_{\mathscr{C},r}(n^{3+o(1)})$.*

Types along decompositions. Many model checking algorithms (including ours) process a decomposition of the input graph step by step, and maintain *types*: formulas of bounded quantifier rank satisfied by the intermediate graphs processed so far. For example, Courcelle's theorem [Cou90] maintains, for every node t of a tree decomposition with bag $\bar{b} = (b_1, \dots, b_k)$, the type

$$\{\psi(x_1, \dots, x_k) \,:\, \psi \text{ is an MSO formula of quantifier rank at most } q, \text{ such that } \psi(\bar{b}) \text{ holds in } G_t\},$$

where G_t is the subgraph induced by the bags in the subtree of t. Here q is the quantifier rank of the input formula φ. The type above is computed from the types of the children of t. Similarly, the algorithms for bounded twin-width maintain FO types along contraction sequences [Bon+20; Gaj+22], and the algorithm for bounded merge-width along construction sequences [DT25b].

We call a first-order formula *local* if, roughly speaking, every quantifier has the form

$$\exists y\,(\mathrm{dist}(\bar{x}, y) \leqslant r \wedge \chi(\bar{x}y)) \quad \text{or} \quad \forall y\,(\mathrm{dist}(\bar{x}, y) \leqslant r \to \chi(\bar{x}y))$$

ranging only over vertices close to some free or previously quantified variables $\bar{x}$. The merge-width model checking algorithm [DT25b] maintains, at every step H of the construction sequence and for every vertex v, its *local type* of quantifier rank q, for some constant q depending on φ:

$$\{\psi(x) \,:\, \psi \text{ is a local formula of quantifier rank at most } q, \text{ such that } \psi(v) \text{ holds in } H\}.$$

In the first step of the construction sequence, nothing is resolved, and the local types are trivial. In the last step, all pairs are resolved, and the local types determine whether the input sentence φ holds in G. In between, the local types have to be updated after every merge and every resolve. Merges are handled easily, by simply rewriting color names; the main difficulty in [DT25b] is to lift the local types through resolves.

Combinatorial explosion. For all these type-based model checking algorithms, the number of possible types grows non-elementarily with the width of the decomposition. In particular, in each step of a construction sequence of radius-r width w, each vertex locally reaches at most w colors, and the number of local formulas over w colors is a tower of exponentials in w. This is a constant if the width w is treated as a constant (the radius r here is fixed and depends on the quantifier rank q). However, monadically dependent classes may have unbounded merge-width, and we can only guarantee a merge-width bound of $w = n^{o(1)}$. In this case the number of all possible local types explodes.

Our algorithm. The tower above counts the types that are *possible* with w colors. Our starting point is the observation that in monadically dependent classes, only few of them are *realized*, that is, actually occur as the type of some tuple of vertices of the input graph. We keep the overall structure of the merge-width algorithm, but change it so that its running time depends only on the realized types. This has three parts, which we describe in the rest of this section.

1. We represent local types by *type trees*, whose size depends only on the realized types.
2. We show that only $n^{o(1)}$ types are realized in a ball, which bounds the size of type trees (Lemma 2.3).
3. We lift type trees through merges and resolves in time polynomial in n and the number of realized types (Lemma 2.4).

Type trees. Type trees (defined in Section 10.2) are a variant of Hintikka types [Sco79]. Simplifying for the sake of presentation, up to the precise choice of radii $r_1 < r_2 < \ldots$, the type tree of a tuple $\bar{a}$ is

$$\mathrm{tree}_q(\bar{a}) := \begin{cases} \mathrm{atp}(\bar{a}) & \text{if } q = 0, \\ \big(\mathrm{atp}(\bar{a}),\ \{(i, \mathrm{tree}_{q-1}(\bar{a}b)) : i \in [|\bar{a}|],\ b \in V,\ \mathrm{dist}(a_i, b) \leqslant r_q\}\big) & \text{if } q > 0, \end{cases}$$

where $\mathrm{atp}(\bar{a})$ is the set of atomic formulas satisfied by $\bar{a}$, including atoms that compare distances between the entries of $\bar{a}$ and the colors. The index i records which entry of $\bar{a}$ is close to b. Two tuples have the same type tree if and only if they have the same local type (Corollary 10.3), so type trees are just another representation of local types. Type trees resemble the morphism-trees of [Bon+20] and serve a similar purpose. Figure 2 shows an example.

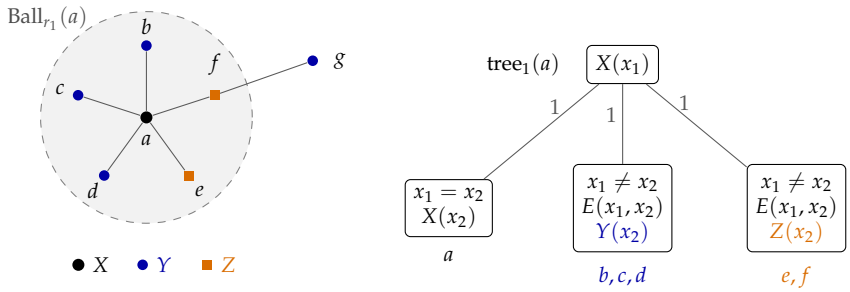


Figure 2: *Left:* A partial graph with 7 vertices and colors X, Y, Z. *Right:* $\mathrm{tree}_q(a)$ for $q = 1$ and $r_1 = 1$, drawn as a tree of depth q whose nodes are the sets $\mathrm{atp}(a)$ and $\mathrm{atp}(a\,v)$. The edge label 1 records that v is close to the first entry of the tuple (a). The lists of vertex labels below the boxes show which vertices of $\mathrm{Ball}_{r_1}(a)$ give each child of the root node. The vertex g outside the locality radius $r_1 = 1$ does not occur in the tree. Note that the size of the tree adapts to the number of atomic types $\mathrm{atp}(a\,v)$ realized in the ball.

Type trees carry the same information as local types, but they may have a much smaller representation. The type tree of a tuple $\bar{a}$ has one subtree for every type tree realized by an extension $\bar{a}b$ with b close to some component a_i of $\bar{a}$, so its size is governed by the nearby realized types rather than by the possible ones. Roughly, this is like storing a small subset of a huge universe as the list of its elements, whereas naive type-based algorithms store it as its characteristic vector. Showing that type trees have size $n^{o(1)}$ is a key step of our algorithm. For this, we first leave the algorithm aside for a moment.

Counting types over parameters. We turn to a purely quantitative result about monadically dependent classes, which counts types over a set of parameters and may be of independent interest. Let H be a graph or partial graph, $U \subseteq V(H)$, and $k, q \in \mathbb{N}$. The *type over* U of a tuple $\bar{a} \in V(H)^k$ is the set of first-order formulas $\psi(x_1, \dots, x_k)$ of quantifier rank at most q which are satisfied by $\bar{a}$ in H, where ψ may use any number of vertices of U as constants, called *parameters*. We denote this type by $\mathrm{tp}_q^H(\bar{a}/U)$. For $X \subseteq V(H)$, we write $\mathrm{Types}_{k,q}^H(X/U)$ for the set of types over U of tuples from X^k,

$$\mathrm{Types}_{k,q}^H(X/U) = \{\mathrm{tp}_q^H(\bar{a}/U) \,:\, \bar{a} \in X^k\}.$$

This set has at most $|X|^k$ elements, but can be much smaller, as two tuples have the same type if no formula distinguishes them. We write $\mathrm{Types}_{k,q}^H(X)$ if $U = \emptyset$. One of the main contributions of this paper is the following bound on the number of types over parameters in monadically dependent graph classes.

Theorem 2.2 (Simplified form of Theorem 4.4)**.** *Let $\mathscr{C}$ be a monadically dependent graph class and $k, q \in \mathbb{N}$. There is some $d \in \mathbb{N}$ such that for every $G \in \mathscr{C}$ and every set $U \subseteq V(G)$ with $|U| > 1$,*

$$|\,\mathrm{Types}_{k,q}^G(V(G)/U)| \leqslant |U|^d.$$

The proof of Theorem 2.2 proceeds as follows. We fix some radius r depending on the quantifier rank q. We then show (in Theorem 5.1) that there is a set $\widehat{U}$ of size polynomial in $|U|$, such that for every vertex $v \in V(G)$, there is some *separator* $S \subseteq \widehat{U}$ of size bounded by some constant s, such that S *separates* v from U *at a distance* r. The notion of separation is technical, but it ensures that the following holds for all $v \in V(G)$ which are r-separated from U by S:

$$\text{The type } \mathrm{tp}_q^G(v/U) \text{ is determined by the type } \mathrm{tp}_{q'}^G(v/S).$$

Here q' depends only on q. Since S has size bounded by s, and both s and q' are constants (depending only on $\mathscr{C}$ and q), the number of types $\mathrm{tp}_{q'}^G(v/S)$ is bounded by a constant c (depending on $\mathscr{C}$ and q). Since there are at most $|\widehat{U}|^s$ considered separators S, it follows that altogether there are at most $c \cdot |\widehat{U}|^s$ types over U of single vertices $v \in V(G)$, of quantifier rank q. This gives the required polynomial bound on the number of types of single vertices v over a set U. The argument for k-tuples $\bar{v}$ proceeds similarly.

The size of type trees. We now return to type trees and bound their sizes using the bound on the number of types. The following lemma makes precise that type trees store only the types realized nearby.

Lemma 2.3 (informal version of Lemma 13.1)**.** *Let $q \in \mathbb{N}$. There is $d \in \mathbb{N}$ depending only on q such that for every partial graph H and vertex v, the type tree $\mathrm{tree}_q(v)$ in H has size polynomial in $K := |\,\mathrm{Types}_{d,d}^H(\mathrm{Ball}_d^H(v))|$.*

We bound K in Lemma 2.3 using the bound on the number of types (Theorem 2.2). The two count different things. K counts types without parameters in a partial graph G_t, whose colors are not present in G, and resolved pairs a priori form an arbitrary subset of vertex pairs. On the other hand, the type bound counts types in G itself, over a set of parameters. While G comes from a monadically dependent graph class, G_t a priori does not, making it difficult to enforce structural properties on G_t.

We bridge this gap by computing construction sequences that are *locally definable* in G: in every local ball, the colors and resolved pairs of G_t can be expressed by first-order formulas in G that mention only $n^{o(1)}$ vertices as parameters. A challenge in encoding G_t into G is that G_t accumulates many resolves over time, which cannot be modeled nicely by parameters. To overcome this, we turn to *merge sequences*, which group many merges and resolves into a single step, and can therefore be much shorter than construction sequences. For every $\varepsilon > 0$, [Dre+26] provides a chain of sets $V(G) = S_1 \supseteq \cdots \supseteq S_\ell = \emptyset$, a closely related decomposition, with only $\ell = O(1/\varepsilon)$ steps, which we convert into a merge sequence of width $n^{O(\varepsilon)}$ with the same number of steps. Each step can be encoded in G by $O(1)$ unary predicates, and when we refine the merge sequence into a construction sequence, every G_t stays locally definable using these $O(1/\varepsilon)$ predicates and few parameters. The bound on the number of types (Theorem 2.2) extends to such expansions of G by boundedly many unary predicates, so it bounds K by $n^{O(\varepsilon)}$, and by Lemma 2.3, every type tree has size $n^{O(\varepsilon)}$ as well.

Lifting type trees through merges and resolves. Now that we know that type trees are small, it remains to compute them efficiently. Our algorithm processes the construction sequence step by step and maintains the type tree of every vertex. So after every merge and every resolve, it has to compute the new type trees from the old ones. Merges are easy. Merging two parts P and Q into $P \cup Q$ does not change the resolved pairs, and hence not the distances. Given a type tree from before the merge, we can update it as follows: in every node, remove all atoms that mention P or Q, and add an atom mentioning $P \cup Q$ whenever the same atom with P or with Q in its place was present. The difficulty lies in resolves, which create new short paths and thereby bring new vertices close to a vertex v (see Figure 3).

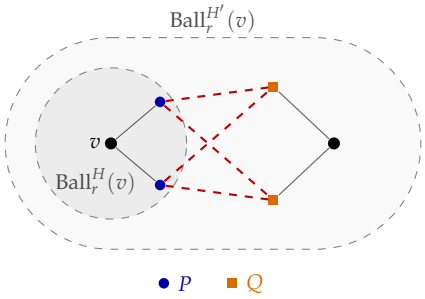


Figure 3: A resolve of the parts P and Q, positive or negative, resolves all pairs between P and Q (red). This creates short paths from v to vertices that were far from v before, so the ball of v grows from $\mathrm{Ball}_r^H(v)$ to $\mathrm{Ball}_r^{H'}(v)$. The old type tree of v contains no information about the new vertices in the ball. The new type tree of v therefore has to be assembled from the old type trees of v and of these vertices.

In [DT25b], resolves are handled by a rank-preserving locality theorem [DT26]. Fix a vertex v. For every formula that could *potentially* belong to the local type of v after the resolve, the theorem expresses whether v satisfies it in terms of the local type of v before the resolve (and some global information about the partial graph). So the new local type is assembled backwards,

one potential formula at a time. For width $n^{o(1)}$, there are far too many of them, even if only few types are realized.

We instead work forwards. We use the same ideas as in the proof of the locality theorem [DT26], but turned around: we construct the new type tree directly from the old type trees, and spend work only on type trees that are actually realized. This is summarized by the following lemma.

Lemma 2.4 (informal version of Lemma 12.2)**.** *Let $q \in \mathbb{N}$. There is $d \in \mathbb{N}$ depending only on q such that the following holds. Given a partial graph H, a resolve producing H', and the type trees of all vertices of H, one can compute the type trees of all vertices of H' in time polynomial in n and $K := \max_v |\operatorname{Types}^H_{d,d}(\operatorname{Ball}^{H'}_d(v))|$.*

A route not taken: games and covers. The algorithms for nowhere dense and monadically stable classes face the same problem of an unbounded number of local types. They sidestep it by a bounded-depth recursion into local neighborhoods, guided by the splitter game [GKS17] or the flipper game [Gaj+23; DMS23; Dre+24], with neighborhood covers keeping the recursion tree small. We tried and failed to adapt the approach to the monadically dependent setting, before settling on our current approach.

Organization. After the preliminaries in Section 3, Section 4 states the ingredients formally and derives Theorem 1.1 from them. The remaining parts prove the ingredients: Part I proves the bound on the number of types (Theorem 2.2), Part II computes the locally definable construction sequences, and Part III gives the model checking algorithm based on type trees (Lemmas 2.3 and 2.4).

3 Preliminaries

$\mathbb{N}$ denotes the set of non-negative integers and $[n] := \{1, \dots, n\}$ for $n \in \mathbb{N}$. For a set V, we write $\binom{V}{2}$ for the set of unordered pairs uv of distinct elements of V, and for $X, Y \subseteq V$ we put $XY := \{xy : x \in X,\ y \in Y,\ x \neq y\}$. A graph G consists of a finite set V of *vertices* and a set $E \subseteq \binom{V}{2}$ of *edges*, denoted by $V(G)$ and $E(G)$. The (open) *neighborhood* of a vertex $v \in V(G)$ is the set $N^G(v) := \{w \in V(G) : vw \in E(G)\}$. For $r \in \mathbb{N}$, denote $\operatorname{Ball}^G_r(v) := \{w \in V(G) : \operatorname{dist}^G(v, w) \leqslant r\}$, where $\operatorname{dist}^G(\cdot, \cdot)$ is the shortest-path distance in G. These notations extend to sets $A, B \subseteq V(G)$ with $\operatorname{Ball}^G_r(A) := \bigcup_{a \in A} \operatorname{Ball}^G_r(a)$ and $\operatorname{dist}^G(A, B) := \min_{a \in A, b \in B} \operatorname{dist}^G(a, b)$, and to tuples of vertices as expected.

Let G be a graph with vertex set V and edge set E. By $\overline{G}$ we denote the edge-complement of G, which has vertex set V and edge set $\binom{V}{2} - E$. The subgraph of G *induced* by a set $X \subseteq V$ is the graph $G[X]$ with vertex set X and edge set $E \cap \binom{X}{2}$. For two disjoint sets $X, Y \subseteq V$, by $G[X, Y]$ we denote the bipartite graph with parts X and Y and edge set $XY \cap E$.

A pair of sets $X, Y \subseteq V(G)$ is *complete* if $XY \subseteq E(G)$ and *anticomplete* if $XY \subseteq E(\overline{G})$. The pair is *homogeneous* if it is complete or anticomplete.

A *graph class* is a set $\mathscr{C}$ of graphs.

Partitions. For two partitions $\mathcal{P}, \mathcal{Q}$ of the same set we write $\mathcal{P} \preccurlyeq \mathcal{Q}$ to denote that $\mathcal{P}$ refines $\mathcal{Q}$, i.e., each part of $\mathcal{P}$ is contained in some part of $\mathcal{Q}$. If $\mathcal{P}$ is a partition of V and $X \subseteq V$, denote $X/\mathcal{P} := \{P \in \mathcal{P} : P \cap X \neq \emptyset\}$. The part of $\mathcal{P}$ containing an element v is denoted by $\mathcal{P}(v)$. Further, $\equiv_{\mathcal{P}}$ is the equivalence relation on V whose classes are the parts of $\mathcal{P}$.

Asymptotic notation. For a possibly empty list of parameters k, by $O_k(\cdot)$ we denote an anonymous function $f\colon \mathbb{R}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$ (depending only on the parameters k) which is upper bounded by $x \mapsto a \cdot x + b$, for some $a, b \in \mathbb{N}$.

Similarly, we write $\mathrm{poly}_k(\cdot)$ for an anonymous function $f\colon \mathbb{R}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$ which is upper bounded by some polynomial; in other words, $g(n) \leqslant \mathrm{poly}_k(n)$ has the same implications as $g(n) \leqslant (n+2)^{O_k(1)}$. We use $\widehat{O}_k(n^c)$ for a bound $O_{k,\varepsilon}(n^{c+\varepsilon})$ for every $\varepsilon > 0$; thus it hides a subpolynomial factor that may depend on the parameters k.

Monadic dependence. We view graphs as relational structures in the signature with one (symmetric, irreflexive) binary relation E. For $k \in \mathbb{N}$, a *k-colored graph* is a graph equipped with k unary predicates $U_1, \ldots, U_k$. A *k-coloring* G^+ of a graph G assigns interpretations of these predicates on $V(G)$; the predicates need not be disjoint. A *coloring* of G is a k-coloring for some $k \in \mathbb{N}$.

A graph class $\mathscr{C}$ is *monadically independent* if there are $k \in \mathbb{N}$ and a first-order formula $\varphi(x, y)$ in the signature of k-colored graphs such that every finite graph H can be obtained from some graph $G \in \mathscr{C}$ as follows: first choose a k-coloring G^+ of G, then consider the graph $\varphi(G^+)$ with vertex set $V(G)$ and edges $uv \in \binom{V(G)}{2}$ such that $G^+ \models \varphi(u, v) \vee \varphi(v, u)$; finally take some induced subgraph of $\varphi(G^+)$, obtaining a graph isomorphic to H. A graph class $\mathscr{C}$ is *monadically dependent* if it is not monadically independent.

Neighborhood complexity. The *neighborhood complexity* of a graph G is the function $\pi_G\colon \mathbb{N} \to \mathbb{N}$ defined by:

$$\pi_G(k) \coloneqq \max_{\substack{S \subseteq V(G) \\ |S| \leqslant k}} \left| \{ N^G(v) \cap S : v \in V(G) \} \right|,$$

which counts the maximum number of distinct neighborhoods of vertices in G restricted to a set of size at most k.

The *VC-dimension* of a graph G is the maximal size of a set $A \subseteq V(G)$ such that $\{N(v) \cap A : v \in V(G)\} = 2^A$. The Sauer–Shelah lemma [Sau72; She72] implies the following.

Lemma 3.1. *If G has VC-dimension d then*

$$\pi_G(k) \leqslant (k+1)^d \quad \text{for all } k \in \mathbb{N}.$$

The following is easy to see from the definition of monadic dependence; see [Dre+26, Observation 3.2] for a proof.

Observation 3.2. *For every monadically dependent graph class $\mathscr{C}$ there is some $d \in \mathbb{N}$ such that every graph $G \in \mathscr{C}$ has VC-dimension at most d.*

Types and parameters. For a structure A and a set W of its elements, let A_W denote the structure obtained by expanding A with constants representing each element of W. A formula $\varphi(x_1, \ldots, x_k)$ *with parameters* from W is a formula in the signature of the structure A_W. For $\bar{v} \in V(A)^k$ we write $A \models \varphi(\bar{v})$ to denote $A_W \models \varphi(\bar{v})$.

For $k, q \in \mathbb{N}$ and a tuple $\bar{v} \in V(A)^k$, define *the type of $\bar{v}$ of quantifier rank q over W*, denoted

$$\mathrm{tp}_q^A(\bar{v}/W),$$

as the set of all formulas $\varphi(x_1, \dots, x_k)$ with parameters from W and of quantifier rank at most q such that $A \models \varphi(\bar{v})$.

For two sets $X, Y \subseteq V(A)$, denote

$$\mathrm{Types}^A_{k,q}(X/Y) := \{\mathrm{tp}^A_q(\bar{v}/Y) \,:\, \bar{v} \in X^k\}.$$

We omit $/Y$ from the notation if $Y = \emptyset$. The *atomic type* of $\bar{v}$ over Y is defined as $\mathrm{atp}^A(\bar{v}/Y) := \mathrm{tp}^A_0(\bar{v}/Y)$ and we set $\mathrm{ATypes}^A(X/Y) := \mathrm{Types}^A_{1,0}(X/Y)$.

Note that for a graph G and vertex set $S \subseteq V(G)$, two vertices $u, v \in V(G)$ have the same atomic type over S if and only if either $u = v$, or $u, v \notin S$ and $N(u) \cap S = N(v) \cap S$. The *partition into atomic types*, denoted $\mathcal{P}_S$, is the corresponding partition of $V(G)$, which partitions S into singletons, and vertices in $V(G) - S$ according to their neighborhood in S.

Gaifman graph, distances, balls. For a structure A with domain V equipped with unary and binary relations, the *Gaifman graph* of A is the graph with vertex set V and edge set consisting of all pairs uv such that (u, v) or (v, u) is in some binary relation of A. We write $\mathrm{dist}^A(u, v)$ to denote the distance between two vertices in the Gaifman graph of A, and $\mathrm{Ball}^A_r(v)$ for the set of vertices at distance at most r from v in the Gaifman graph of A.

Flips and flip metrics. For a graph G and partition $\mathcal{P}$ of $V := V(G)$, a *$\mathcal{P}$-flip* of G is a graph H with $V(H) = V$ such that for every pair $X, Y \in \mathcal{P}$ (possibly $X = Y$), either $H[X, Y] = G[X, Y]$ or $H[X, Y] = \overline{G}[X, Y]$. If $|\mathcal{P}| \leqslant s$, then we call H an *s-flip of G*.

Define the associated *flip metric* and its balls by:

$$\mathrm{dist}_{\mathcal{P}}(u, v) := \max\{\mathrm{dist}^H(u, v) : H \text{ is a } \mathcal{P}\text{-flip of } G\},$$
$$\mathrm{Ball}^{\mathcal{P}}_r(v) := \{u : \mathrm{dist}_{\mathcal{P}}(u, v) \leqslant r\} = \bigcap\{\mathrm{Ball}^H_r(v) \,:\, H \text{ is a } \mathcal{P}\text{-flip of } G\}.$$

As a shorthand, for a graph G and a set $S \subseteq V(G)$, a $\mathcal{P}_S$-flip of G is called an *S-flip*, and we set $\mathrm{dist}^G_S := \mathrm{dist}_{\mathcal{P}_S}$ and $\mathrm{Ball}^S_r := \mathrm{Ball}^{\mathcal{P}_S}_r$.

Partial graphs and merge-width. Partial graphs, construction sequences, their radius-r width, and the radius-r merge-width of a graph are defined in Section 2 (see also [DT25a]). Moreover, for a partial graph H with partition $\mathcal{P}$ and a set $W \subseteq V(H)$, the induced partial graph $H[W]$ is the partial graph with vertex set W, edges $E(H) \cap \binom{W}{2}$, non-edges $N(H) \cap \binom{W}{2}$, and partition $\{P \cap W \,:\, P \in \mathcal{P}\} - \{\emptyset\}$.

4 Proof of the main theorem

The upcoming Parts I, II, and III prove three independent results. This section combines them into the main result, Theorem 1.1. It does so via two theorems that are connected by the following type-based width notion.

Definition 4.1. For $r, q \in \mathbb{N}$, the *radius-r depth-q width* of a construction sequence $G_1, \dots, G_\lambda$ is defined as:

$$\max_{t \in [\lambda-1]} \max_{v \in V(G)} |\mathrm{Types}^{G_t}_{q,q}(\mathrm{Ball}^{G_{t+1}}_r(v))|.$$

The offset between G_t and G_{t+1} mirrors the offset between $\mathcal{P}_{i-1}$ and R_i in the width definition of chains and of merge sequences (see Definitions 6.2 and 6.5).

Theorem 4.2. *Let $\mathscr{C}$ be a monadically dependent graph class, $r, q \in \mathbb{N}$, and let $\varepsilon > 0$ be a rational. There is a constant $\alpha \geqslant 1$ depending only on $\mathscr{C}$ and q such that every n-vertex graph $G \in \mathscr{C}$ has a construction sequence of radius-r depth-q width $O_{\mathscr{C},r,q,\varepsilon}(n^{\alpha\varepsilon})$ and length $O_{\mathscr{C},r,q,\varepsilon}(n^{1+\alpha\varepsilon})$. Moreover, there is an algorithm (independent of $\mathscr{C}$) which, given $r, q \in \mathbb{N}$, a rational $\varepsilon > 0$, and an n-vertex graph $G \in \mathscr{C}$, computes such a construction sequence in time $\widehat{O}_{\mathscr{C},r,q,\varepsilon}(n^3)$.*

It is important that α does not depend on ε, as this allows us to make the exponent $\alpha\varepsilon$ of the width as small as we like. The upcoming Section 4.1 assembles Theorem 4.2 from the key results of Parts I and II.

Theorem 4.3. *There is an algorithm which, given a sentence φ and an n-vertex graph G together with its construction sequence of length λ, determines whether $G \models \varphi$ in time $\lambda \cdot n^2 \cdot \mathrm{poly}_\varphi(K)$, where K is the radius-r depth-q width of the construction sequence, for some r and q which are computable from φ.*

Proof. See Part III. ∎

From these two theorems, we can derive Theorem 1.1 without much effort. The only subtlety is that the algorithm does not know $\mathscr{C}$, so it cannot choose ε itself.

Theorem 1.1. *There is an algorithm which solves the model checking problem for first-order logic on graphs – does a given graph G satisfy a given first-order sentence φ? – with the following guarantee. For every monadically dependent graph class $\mathscr{C}$ and real $\varepsilon > 0$, the algorithm runs in time $O_{\mathscr{C},\varphi,\varepsilon}(n^{3+\varepsilon})$ for n-vertex graphs $G \in \mathscr{C}$.*

Proof of Theorem 1.1. Let φ be the input sentence, and let r, q be the numbers from Theorem 4.3. For $j \geqslant 1$, let $\mathcal{A}_j$ be the algorithm that computes a construction sequence of G by Theorem 4.2 with r, q, and $\varepsilon := 1/j$, and then decides whether $G \models \varphi$ by Theorem 4.3.

Fix a monadically dependent graph class $\mathscr{C}$, and let $\alpha \geqslant 1$ be the constant from Theorem 4.2 for $\mathscr{C}$ and q. Let $c \geqslant 1$ be such that the polynomial in Theorem 4.3 is $O_\varphi(K^c)$. On n-vertex graphs $G \in \mathscr{C}$, the sequence computed by $\mathcal{A}_j$ has radius-r depth-q width $K \leqslant O_{\mathscr{C},\varphi,j}(n^{\alpha/j})$ and length $\lambda \leqslant O_{\mathscr{C},\varphi,j}(n^{1+\alpha/j})$. So $\mathcal{A}_j$ runs in time

$$\widehat{O}_{\mathscr{C},\varphi,j}(n^3) + O_\varphi(\lambda \cdot n^2 \cdot K^c) \leqslant O_{\mathscr{C},\varphi,j}(n^{3+\beta/j}),$$

where $\beta := (c+1)\alpha$. Since $\beta/j > 0$, the subpolynomial factor hidden by $\widehat{O}$ is absorbed by $n^{\beta/j}$. The number β depends on $\mathscr{C}$ and φ, but not on j.

The algorithm does not know $\mathscr{C}$, so it cannot choose j. Instead, it runs all algorithms $\mathcal{A}_j$ in parallel, giving $\mathcal{A}_j$ a 2^{-j} share of the computation, and returns the first answer. Its answer is correct, as every $\mathcal{A}_j$ correctly solves the model checking problem on every input graph.

For the running time, let $\varepsilon > 0$ and $j := \lceil 2\beta/\varepsilon \rceil$. If $\mathcal{A}_j$ halts within time T, then it completes within the first $2^j T$ simulated steps, so the algorithm halts within time

$$O\big(2^j T \log(2^j T + 2)\big) = O_j\big(T \log(T+2)\big).$$

Hence, on n-vertex graphs in $\mathscr{C}$, the algorithm runs in time

$$O_{\mathscr{C},\varphi,\varepsilon}\big(n^{3+\beta/j} \log(n+2)\big) \leqslant O_{\mathscr{C},\varphi,\varepsilon}(n^{3+\varepsilon}),$$

since $\beta/j \leqslant \varepsilon/2$ and $\log(n+2) = O_\varepsilon(n^{\varepsilon/2})$. ∎

4.1 Proving Theorem 4.2 from Parts I and II

We first state the key results of the two parts, and derive Theorem 4.2 afterwards.

4.1.1 The key result of Part I

The key ingredient towards proving Theorem 4.2 is the following type bound derived in Part I.

Theorem 4.4. *Let $\mathscr{C}$ be a monadically dependent graph class and $k, q, c \in \mathbb{N}$. For every graph $G \in \mathscr{C}$, expansion G^+ of G by c unary predicates, and subset $U \subseteq V(G)$,*

$$\left|\mathrm{Types}_{k,q}^{G^+}(V(G)/U)\right| \leqslant O_{\mathscr{C},q,k,c}(1) \cdot \mathrm{poly}_{\mathscr{C},q,k}(|U|).$$

The number c of unary predicates affects only the constant factor, not the degree of the polynomial. This matters, as our construction sequences use $O(r/\varepsilon)$ unary predicates, while the constant α in Theorem 4.2 must not depend on ε.

4.1.2 The key result of Part II

The type bound counts types in G, while the width (Definition 4.1) counts types in the partial graphs G_t. To bridge this gap, Part II produces a construction sequence in which local balls of G_t are defined by formulas in G.

Definition 4.5. Let G be a graph with vertex set V and let H be a partial graph with vertex set $W \subseteq V$. For $q, c, w \in \mathbb{N}$, we say that H is *definable* in G by formulas of quantifier rank q using c unary predicates and w parameters, if there is an expansion G^+ of G by c unary predicates and a set $S \subseteq V(G)$ of parameters with $|S| \leqslant w$, such that for every relation symbol Q of H (when viewed as a relational structure) of arity k there is some formula $\varphi_Q(x_1, \ldots, x_k)$ of quantifier rank q with parameters from S such that

$$G^+ \models \varphi_Q(\bar{v}) \quad \iff \quad \bar{v} \in W^k \text{ and } H \models Q(\bar{v}) \qquad \text{for all } \bar{v} \in V^k.$$

The vertex set W of H is the union of its parts, so it is defined by the disjunction of the formulas defining the parts.

Definition 4.6. Let G be a graph and $r, c, s, q_0 \in \mathbb{N}$. A construction sequence $G_1, \ldots, G_\lambda$ of G is *radius-r definable* with c unary predicates, s parameters, and quantifier rank q_0 if for all $t \in [\lambda - 1]$ and $v \in V(G)$, the induced partial graph

$$G_t[\mathrm{Ball}_r^{G_{t+1}}(v)]$$

is definable in G by formulas of quantifier rank q_0 using c unary predicates and s parameters.

Note that the expansion and the parameters may depend on t and v.

The key result of Part II is the following theorem.

Theorem 4.7. *Let $r \geqslant 1$ be an integer, $\varepsilon > 0$ be a rational, and $\mathscr{C}$ be a monadically dependent graph class. There is a constant $\delta \geqslant 1$ depending only on $\mathscr{C}$ such that every n-vertex graph $G \in \mathscr{C}$ has a construction sequence of G which is radius-r definable with $O(r/\varepsilon)$ unary predicates, $O_{\mathscr{C},r,\varepsilon}(n^{\delta\varepsilon})$ parameters, and quantifier rank 7, and has length $O_{\mathscr{C},r,\varepsilon}(n^{1+\delta\varepsilon})$. Moreover, there is an algorithm (independent of $\mathscr{C}$) which, given r, ε, and G, computes such a sequence, and runs in time $\widehat{O}_{\mathscr{C},r,\varepsilon}(n^3)$ whenever $G \in \mathscr{C}$.*

For nowhere dense classes, the following stronger bounds follow easily from known results. The proof is given in Section C.

Theorem 4.8. *Let $\mathscr{C}$ be a nowhere dense graph class, $r \in \mathbb{N}$, and $\varepsilon > 0$. Every n-vertex graph $G \in \mathscr{C}$ has a radius-r definable construction sequence with one unary predicate, $O_{\mathscr{C},r,\varepsilon}(n^{\varepsilon})$ parameters, quantifier rank 0, and length $O_{\mathscr{C},r,\varepsilon}(n^{1+\varepsilon})$. Such a sequence can be computed in time $O_{\mathscr{C},r,\varepsilon}(n^{1+\varepsilon})$.*

In particular, Theorem 4.8, together with the results of Parts I and III, gives an alternative proof of the result of Grohe, Kreutzer, and Siebertz [GKS17] that first-order model checking is fixed-parameter tractable on nowhere dense classes. (Alternatively, in nowhere dense classes one can replace the type bound from Part I by a stronger bound deducible from [PST18]).

4.1.3 Assembling the parts

We use the following consequence of Gaifman locality, proved in Section B.

Lemma 4.9 (Corollary of Gaifman locality)**.** *For every $k, q, \alpha \in \mathbb{N}$, there are numbers $h = h(q)$ and $q' = q'(k, q, \alpha)$ computable from k, q, α, with the following property. Fix a finite relational signature σ of arity at most α and a σ-structure A. Then, for all $\bar{a} \in V(A)^k$ and $B \supseteq \mathrm{Ball}_h^A(\bar{a})$, $\mathrm{tp}_q^A(\bar{a})$ depends only on $\mathrm{tp}_{q'}^{A[B]}(\bar{a})$.*

Lemma 4.10. *For every $q \in \mathbb{N}$, there is $h \in \mathbb{N}$, computable from q, such that the following holds for every monadically dependent graph class $\mathscr{C}$ and all $r, c, s, q_0 \in \mathbb{N}$. If a construction sequence of a graph $G \in \mathscr{C}$ is radius-$(r+h)$ definable with c unary predicates, s parameters, and rank q_0, then its radius-r depth-q width is at most $O_{\mathscr{C},q,q_0,c}(1) \cdot \mathrm{poly}_{\mathscr{C},q,q_0}(s)$.*

Proof. We will apply Lemma 4.9 to partial graphs from a construction sequence of G. These partial graphs will have a finite signature σ with up to $|V(G)|$ unary predicates marking the partition, but the arity is fixed to $\alpha = 2$. Hence, we can choose $h := h(q)$ and $q' := q'(q, q, 2)$ from Lemma 4.9. Let $G_1, \dots, G_\lambda$ be a construction sequence of $G \in \mathscr{C}$ which is radius-$(r+h)$ definable with c unary predicates, s parameters, and rank q_0. Fix $t \in [\lambda - 1]$ and $v \in V(G)$, and let $B := \mathrm{Ball}_{r}^{G_{t+1}}(v)$ and $W := \mathrm{Ball}_{r+h}^{G_{t+1}}(v)$. By Definition 4.1, we have to bound $\left|\mathrm{Types}_{q,q}^{G_t}(B)\right|$.

First, we move from G_t to the partial graph $G_t[W]$. As G_{t+1} has all resolved pairs of G_t, every tuple $\bar{u} \in B^q$ satisfies $\mathrm{Ball}_h^{G_t}(\bar{u}) \subseteq W$. So by Lemma 4.9, $\mathrm{tp}_q^{G_t}(\bar{u})$ is determined by $\mathrm{tp}_{q'}^{G_t[W]}(\bar{u})$, and hence

$$\left|\mathrm{Types}_{q,q}^{G_t}(B)\right| \leqslant \left|\mathrm{Types}_{q,q'}^{G_t[W]}(W)\right|.$$

Next, we move from $G_t[W]$ to G. By assumption, $G_t[W]$ is definable in an expansion G^+ of G by c unary predicates, by formulas of quantifier rank q_0 with parameters from a set S of at most s vertices. Replace every atom of a formula over $G_t[W]$ by its defining formula, and restrict every quantifier to the vertex set of $G_t[W]$. This turns a formula of quantifier rank q' over $G_t[W]$ into a formula of quantifier rank at most $q' + q_0$ over G^+ with parameters from S, which holds for the same tuples from W. So for every $\bar{u} \in W^q$, $\mathrm{tp}_{q'}^{G_t[W]}(\bar{u})$ is determined by $\mathrm{tp}_{q'+q_0}^{G^+}(\bar{u}/S)$.

Together with Theorem 4.4, we obtain

$$\left|\mathrm{Types}_{q,q}^{G_t}(B)\right| \leqslant \left|\mathrm{Types}_{q,q'}^{G_t[W]}(W)\right| \leqslant \left|\mathrm{Types}_{q,q'+q_0}^{G^+}(V(G)/S)\right| \leqslant O_{\mathscr{C},q,q'+q_0,c}(1) \cdot \mathrm{poly}_{\mathscr{C},q,q'+q_0}(s).$$

As q' depends only on q, this is the claimed bound $O_{\mathscr{C},q,q_0,c}(1) \cdot \mathrm{poly}_{\mathscr{C},q,q_0}(s)$. ∎

Proof of Theorem 4.2. Let $\mathscr{C}$ be a monadically dependent graph class containing G. Let h be the number given by Lemma 4.10 for depth q. We run the algorithm of Theorem 4.7 with radius $r+h$ and ε, which takes time $\widehat{O}_{\mathscr{C},r,q,\varepsilon}(n^3)$.

By Theorem 4.7, there is a constant $\delta \geqslant 1$ depending only on $\mathscr{C}$ such that the sequence is radius-$(r+h)$ definable with $c \leqslant O((r+h)/\varepsilon) \leqslant O_{r,q,\varepsilon}(1)$ unary predicates, $s \leqslant O_{\mathscr{C},r,q,\varepsilon}(n^{\delta\varepsilon})$ parameters, and rank $q_0 := 7$, and it has length $O_{\mathscr{C},r,q,\varepsilon}(n^{1+\delta\varepsilon})$. By Lemma 4.10, there is an integer $D \geqslant 1$ such that its radius-r depth-q width is at most

$$O_{\mathscr{C},q,q_0,c}(1) \cdot \mathrm{poly}_{\mathscr{C},q,q_0}(s) \leqslant O_{\mathscr{C},r,q,\varepsilon}(n^{D\delta\varepsilon}),$$

where we used $q_0 = 7$ and $c \leqslant O_{r,q,\varepsilon}(1)$. As D depends only on $\mathscr{C}$ and q, we can choose $\alpha := D\delta$ to satisfy the statement. Since $D \geqslant 1$ and $\delta \geqslant 1$, we have $\alpha \geqslant \delta \geqslant 1$, so this choice also bounds the sequence length by $O_{\mathscr{C},r,q,\varepsilon}(n^{1+\alpha\varepsilon})$. ■

Part I
Number of types

In this part we prove the following result. As we believe it to be of independent interest also in the context of infinite graphs, all considered graphs may be finite or infinite in Part I, and we track the necessary finiteness assumptions in the arguments below. Everywhere else in the paper, graphs are assumed to be finite.

Theorem 4.4. *Let $\mathscr{C}$ be a monadically dependent graph class and $k, q, c \in \mathbb{N}$. For every (possibly infinite) graph $G \in \mathscr{C}$, expansion G^+ of G by c unary predicates, and finite subset $U \subseteq V(G)$,*

$$\left|\mathrm{Types}_{k,q}^{G^+}(V(G)/U)\right| \leqslant O_{\mathscr{C},q,k,c}(1) \cdot \mathrm{poly}_{\mathscr{C},q,k}(|U|).$$

We make two remarks, which are of independent interest in the context of model theory.

Remark 4.11. Conversely, the conclusion of Theorem 4.4 implies monadic dependence. Indeed, let $c \in \mathbb{N}$, and let $\varphi(x,y)$ be any formula in the signature of c-colored graphs, and let q denote its quantifier rank. Let N be a number and H be the bipartite graph with parts $A = \{a_1, \dots, a_N\}$ and $B = \{b_I : I \subseteq [N]\}$, where a_i is adjacent to b_I if and only if $i \in I$. Suppose there is some $G \in \mathscr{C}$ and c-coloring G^+ of G, such that H is an induced subgraph of $\varphi(G^+)$, the graph with vertex set $V(G)$ and edge set $\{uv \in \binom{V(G)}{2} : G^+ \models \varphi(u,v) \vee \varphi(v,u)\}$. Then $\mathrm{tp}^q(b_I/A)$ is different for every $I \subseteq [N]$, so the number of types of quantifier rank q over A is at least 2^N. On the other hand, the bound in Theorem 4.4 is at most polynomial in $|A| = N$, which implies that N is bounded by a constant depending on $\mathscr{C}$ and φ. Hence, not arbitrarily large graphs can be encoded in $\mathscr{C}$ by φ.

Remark 4.12. An inspection of our proof of Theorem 4.4 shows that it only uses the assumption that $\mathscr{C}$ is a *dp-minimal* class of graphs (see Section A for a definition). Thus, the conclusion of Theorem 4.4 holds under the assumption of dp-minimality.

Monadic dependence implies dp-minimality, as shown in Lemma A.2. For general classes of structures, however, dp-minimality is a strictly weaker assumption than monadic dependence[7]. With the previous remark, we conclude that for classes $\mathscr{C}$ of graphs, the notions of dp-minimality and monadic dependence coincide. This extends to classes of structures over a finite signature consisting of binary relations, via the standard encoding of such structures as graphs (see [DMT24, full version, Section 19]).

5 Closures

The goal of this section is to prove the following theorem, which can be seen as a finitary analogue of [Gaj+23, Theorem 7.1 in the full version].

Theorem 5.1 (Closure Theorem)**.** *For every monadically dependent graph class $\mathscr{C}$ and radius $r \in \mathbb{N}$ there are some $c, s \in \mathbb{N}$ with the following property.*

For every $G \in \mathscr{C}$ and finite set $A \subseteq V(G)$ there is a set $A^* \subseteq V(G)$ of size $O_{\mathscr{C},r}(|A|^c)$ such that for every vertex $v \notin A$ there is a set $S(v) \subseteq A^*$ of size at most s and an $S(v)$-flip H with

$$\mathrm{dist}^H(v, A) > r.$$

[7] For example, $(\mathbb{Z}, +)$ is dp-minimal [Asc+16, paragraph before Proposition 6.6] but not monadically dependent.

We also extend the statement to tuples. For simplicity, and in contrast to Theorem 5.1, the conclusion uses non-definable flips. Definability can easily be restored using the methods from [Bon+22b] (see also [Bon+25]).

Theorem 5.2 (Closure Theorem for Tuples)**.** *For every monadically dependent graph class $\mathscr{C}$, radius $r \in \mathbb{N}$, and tuple size $k \in \mathbb{N}$, there are some $c, s \in \mathbb{N}$ with the following property.*

For every $G \in \mathscr{C}$ and finite set $A \subseteq V(G)$ there is a set $\mathcal{H}$ of $O_{\mathscr{C},r,k}(|A|^c)$ many s-flips of G such that for every tuple $\bar{v} \in (V(G) - A)^k$ there is a flip $H \in \mathcal{H}$ with

$$\operatorname{dist}^H(\bar{v}, A) > r.$$

5.1 Preliminaries

Definable flips. For $s \in \mathbb{N}$, let $\Pi = \{(x = y), (x \neq y) \wedge E(x,y), (x \neq y) \wedge \neg E(x,y)\}$ denote the set of three possible atomic types of a pair of vertices in a graph. For an s-tuple $\bar{p} = (p_1, \dots, p_s)$ of vertices in a graph G and a vertex v, we identify $\operatorname{atp}(v/\bar{p})$ with an s-tuple in Π^s, whose ith component is the atomic type of the pair (v, p_i). For a symmetric relation $R \subseteq \Pi^s \times \Pi^s$, we denote by $G \oplus (\bar{p}, R)$ the $\{p_1, \dots, p_s\}$-flip H of G obtained by partitioning $V(G)$ into atomic types over $\bar{p}$ and then flipping according to R. More precisely, for $uv \in \binom{V}{2}$, we have that $uv \in E(H) \triangle E(G)$ (where $\triangle$ denotes the symmetric difference) if and only if $(\operatorname{atp}^G(u/\bar{p}), \operatorname{atp}^G(v/\bar{p})) \in R$. We call such a flip *$s$-definable*. Equivalently, an s-definable flip is an S-flip for some $S \subseteq V(G)$ of size at most s. Note that an s-definable flip is also s'-definable for every $s' \geqslant s$.

Observation 5.3. *We can define the edges of a definable flip in the original graph. More precisely, for every $s \in \mathbb{N}$ and symmetric relation $R \subseteq \Pi^s \times \Pi^s$ there exists a quantifier-free formula $E_{\oplus R}(x, y, \bar{z})$ such that for every graph G, tuple $\bar{p} \in V(G)^s$, and vertices $u, v \in V(G)$ we have*

$$G \oplus (\bar{p}, R) \models E(u,v) \iff G \models E_{\oplus R}(u, v, \bar{p}).$$

The following is an easy consequence of the characterization of monadic dependence through transductions (see, e.g., [DMT24]). It also holds in the greater generality of s-flips (rather than s-definable flips).

Fact 5.4. *Let $\mathscr{C}$ be a monadically dependent graph class and $s \in \mathbb{N}$. The class consisting of all s-definable flips of graphs from $\mathscr{C}$ is monadically dependent.*

The following is well-known, see, e.g., Lemma 2.6 in the full version of [Gaj+23].

Fact 5.5 (Transitivity of definable flips)**.** *Let G_1 be an s_1-definable flip of G_0 and G_2 be an s_2-definable flip of G_1. Then G_2 is an $(s_1 + s_2)$-definable flip of G_0.*

5.2 The main argument

The argument is a finitary analogue of the reasoning in [Gaj+23, Section 7]. We use two lemmas. The first enlarges a given set A to a set A^* such that balls avoiding A^* have few types over A. The second uses this bound to increase the separation radius by one.

Lemma 5.6. *For every monadically dependent graph class $\mathscr{C}$ and $r, s \in \mathbb{N}$, there are some $k, t \in \mathbb{N}$ with the following property. For every $G \in \mathscr{C}$ and finite set $A \subseteq V(G)$, there is $A^* \supseteq A$ of size $O_{\mathscr{C},r,s}(|A|^k)$ such that, for every s-definable flip H of G and every vertex v,*

$$\operatorname{Ball}_r^H(v) \cap A^* = \emptyset \quad \Longrightarrow \quad |\operatorname{ATypes}^H(\operatorname{Ball}_r^H(v)/A)| \leqslant t.$$

The inequality above can be equivalently stated as follows: $\mathrm{Ball}_r^H(v)$ intersects at most t parts of the partition of $V(G)$ into atomic types over A in H. The proof of Lemma 5.6 is deferred to Section 5.3.

Lemma 5.7. *For all $d, t \in \mathbb{N}$, there is some $z \in \mathbb{N}$ with the following property.*
For every graph G of VC-dimension at most d, radius $r \in \mathbb{N}$, finite set $A \subseteq V(G)$, and arbitrary set $C \subseteq V(G)$ satisfying

$$\mathrm{Ball}_r^G(C) \cap A = \emptyset \qquad \text{and} \qquad |\mathrm{ATypes}^G(\mathrm{Ball}_r^G(C)/A)| \leqslant t,$$

there exists a z-definable flip H of G in which $\mathrm{Ball}_{r+1}^H(C) \cap A = \emptyset$.

The proof is deferred to Section 5.4. We now conclude Theorem 5.1 from the two lemmas.

Theorem 5.1 (Closure Theorem)**.** *For every monadically dependent graph class $\mathscr{C}$ and radius $r \in \mathbb{N}$ there are some $c, s \in \mathbb{N}$ with the following property.*
For every $G \in \mathscr{C}$ and finite set $A \subseteq V(G)$ there is a set $A^* \subseteq V(G)$ of size $O_{\mathscr{C},r}(|A|^c)$ such that for every vertex $v \notin A$ there is a set $S(v) \subseteq A^*$ of size at most s and an $S(v)$-flip H with

$$\mathrm{dist}^H(v, A) > r.$$

For brevity, we say a flip H of G *r-separates* two sets $A, B \subseteq V(G)$ if $\mathrm{dist}^H(A, B) > r$. This terminology extends from sets to vertices as expected.

Proof of Theorem 5.1. We induct on r, with $A^* := A$, $S(v) := \emptyset$, and $H := G$ proving the base case $r = 0$. Suppose the theorem holds at radius r, with parameter bound s. By Fact 5.4, and since monadically dependent classes have bounded VC-dimension (Observation 3.2), it follows that every s-definable flip of a graph in $\mathscr{C}$ has VC-dimension at most some $d \leqslant O_{\mathscr{C},r}(1)$.

Fix $G \in \mathscr{C}$ and a finite set $A \subseteq V(G)$. Assume A is nonempty, as otherwise the statement is trivial. Applying Lemma 5.6 to r, s, G, A gives a polynomial-size superset $A^* \supseteq A$ and a constant $t \leqslant O_{\mathscr{C},r}(1)$ such that every r-ball in any s-definable flip H of G that avoids A^* realizes at most t types over A in H. (Crucially, to avoid cyclic dependencies, note that s is an input, not an output, of Lemma 5.6.)

Apply the induction hypothesis to A^*, obtaining a vertex set B of size $\mathrm{poly}_{\mathscr{C},r}(|A^*|) \leqslant \mathrm{poly}_{\mathscr{C},r}(|A|)$ such that every vertex outside A^* is r-separated from A^* by a flip defined over a subset of B of size at most s. We may assume that B is nonempty.

For each $\bar{p} \in B^s$ and each symmetric relation $R \subseteq \Pi^s \times \Pi^s$, put $H := G \oplus (\bar{p}, R)$. Let V_H be the set of vertices v that are r-separated from A^* by H, that is, $\mathrm{Ball}_r^H(v) \cap A^* = \emptyset$.

We assign to each $v \in V_H$ its *profile*

$$\mathrm{profile}_H(v) := \mathrm{ATypes}^H(\mathrm{Ball}_r^H(v)/A).$$

Equivalently, if $\mathcal{P}_A$ denotes the partition of $V(G)$ into atomic types over A in H, then $\mathrm{profile}_H(v) = \mathrm{Ball}_r^H(v)/\mathcal{P}_A$ is the set of equivalence classes that are intersected by $\mathrm{Ball}_r^H(v)$.

Then $\mathrm{profile}_H(v)$ is a subset of $\mathrm{ATypes}^H(V(G)/A)$ of size at most t, by Lemma 5.6. By the Sauer–Shelah lemma (Lemma 3.1), we have that

$$|\mathrm{ATypes}^H(V(G)/A)| \leqslant (|A|+1)^d + |A|.$$

It follows that

$$|\{\mathrm{profile}_H(v) \,:\, v \in V_H\}| \leqslant |\mathrm{ATypes}^H(V(G)/A)|^t \leqslant \left((|A|+1)^d + |A|\right)^t \leqslant \mathrm{poly}_{\mathscr{C},r}(|A|).$$

Group the vertices of V_H by their profiles. For each group C, the ball $\mathrm{Ball}_r^H(C)$:

- avoids A, because each $\mathrm{Ball}_r^H(v)$ with $v \in V_H$ does, and
- realizes at most t types over A in H, because all $v \in C$ have the same profile of size $\leqslant t$.

Applying Lemma 5.7 yields a number $z \leqslant O_{d,t}(1) \leqslant O_{\mathscr{C},r}(1)$ and a further z-definable flip H_C of H that $(r+1)$-separates C from A. By transitivity of definable flips (Fact 5.5), H_C is an $(s+z)$-definable flip of G. Choose a set $U_{\bar{p},R,C} \subseteq V(G)$ of size at most $s+z$ such that H_C is a $U_{\bar{p},R,C}$-flip of G.

We finish the proof by setting

$$A^{**} := A^* \cup \bigcup_{\bar{p},R,C} U_{\bar{p},R,C}.$$

There are at most $|B|^s$ choices of $\bar{p}$, $O_s(1)$ choices of R for each $\bar{p}$, and $\mathrm{poly}_{\mathscr{C},r}(|A|)$ groups C for each choice of $\bar{p}, R$. Since each $U_{\bar{p},R,C}$ has size at most $s+z$, we have $|A^{**}| \leqslant \mathrm{poly}_{\mathscr{C},r}(|A|)$.

For each $v \notin A^*$, the induction hypothesis gives a tuple $\bar{p} \in B^s$ and a $\bar{p}$-flip H that r-separates v from A^*. Let R be its flip relation, so that $H = G \oplus (\bar{p}, R)$, and let C be the group of V_H containing v. Then we may take $S(v) := U_{\bar{p},R,C} \subseteq A^{**}$, since H_C $(r+1)$-separates v from A. For $v \in A^* - A$, take $S(v) := \{v\} \subseteq A^{**}$ and the flip that isolates v. Hence, the theorem holds at radius $r+1$ with parameter set A^{**} and parameter bound $s+z$. ∎

5.3 Few types in balls

In this subsection we prove Lemma 5.6. We use the following Ramsey-type result for set systems due to Ding, Oporowski, Oxley, and Vertigan [Din+96].

Definition 5.8. For a graph G and tuples $\bar{a}, \bar{b} \in V(G)^t$, we say $\bar{a} = a_1 \dots a_t$ and $\bar{b} = b_1 \dots b_t$ are *atomically matched* if there is a relation $\sim \in \{=, \neq, \geqslant\}$ such that $G \models E(a_i, b_j)$ if and only if $i \sim j$ for all $i, j \in [t]$. In particular, the bipartite graph semi-induced between the tuples $\bar{a}$ and $\bar{b}$ is a matching, a co-matching, or a half-graph.

Theorem 5.9 ([Din+96, Corollary 2.4], see also [Gra+04, Theorem 2])**.** *There exists a computable function $Q : \mathbb{N} \to \mathbb{N}$ such that for every $\ell \in \mathbb{N}$ and for every bipartite graph $G = (L, R, E)$ with $|\mathrm{ATypes}(L/R)| > Q(\ell)$, there are two tuples $\bar{a} \in L^\ell$ and $\bar{b} \in R^\ell$ that are atomically matched in G.*

We use the above result to construct the following obstructions to monadic dependence.

Definition 5.10. Fix first-order formulas $\varphi(\bar{x}, z)$ and $\psi(\bar{y}, z)$. A (φ, ψ)-*pattern* of order $s \times t$ in a graph G consists of tuples $\bar{a}_1, \dots, \bar{a}_s \in V(G)^{|\bar{x}|}$, $\bar{b}_1, \dots, \bar{b}_t \in V(G)^{|\bar{y}|}$, and vertices $(c_{i,j} : i \in [s], j \in [t])$ such that the following holds for all $i, i' \in [s]$ and $j, j' \in [t]$:

1. $G \models \varphi(\bar{a}_i, c_{i,j})$ and $G \not\models \varphi(\bar{a}_{i'}, c_{i,j})$ if $i' > i$,
2. $G \models \psi(\bar{b}_j, c_{i,j})$ and $G \not\models \psi(\bar{b}_{j'}, c_{i,j})$ if $j' > j$.

We call $\{c_{i,j} : i \in [s], j \in [t]\}$ the *grid-points* of the pattern, and s its *height*. See Figure 4.

The following lemma is a corollary of [ST21, Lemma C.1]. We include a self-contained proof in Section A for completeness.

Lemma 5.11. *For every monadically dependent graph class $\mathscr{C}$ and formulas $\varphi(\bar{x}, z), \psi(\bar{y}, z)$, there is some $t \in \mathbb{N}$ such that no graph $G \in \mathscr{C}$ contains a (φ, ψ)-pattern of order $t \times t$.*

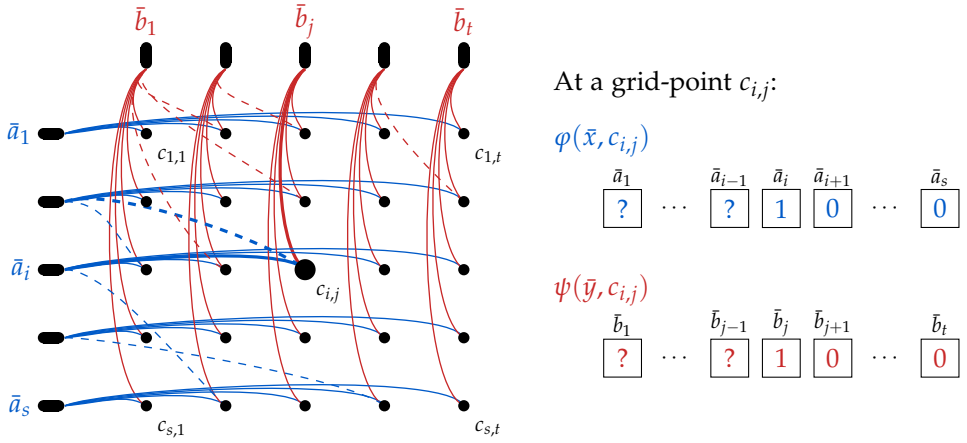


Figure 4: A (φ, ψ)-pattern of order $s \times t$ for $s = t = 5$. Edges indicate $\varphi(\bar{a}_{i'}, c_{i,j})$ or $\psi(\bar{b}_{j'}, c_{i,j})$. Each grid-point $c_{i,j}$ is related via φ to the tuple $\bar{a}_i$ in its row and possibly (indicated by dashed edges) to some earlier row tuples, and analogously for ψ and the column tuples. The profiles on the right illustrate these conditions at a grid-point $c_{i,j}$.

Definition 5.12. For a first-order formula $\varphi(\bar{x}, z)$, a *φ-pattern* is a (φ, ψ)-pattern with $\psi(y, z) = E(y, z)$ or $\psi(y, z) = \neg E(y, z)$. We call those a *$(\varphi, E)$-pattern* and a *$(\varphi, \neg E)$-pattern*, respectively. Such a pattern consists of tuples $\bar{a}_1, \dots, \bar{a}_s \in V(G)^{|\bar{x}|}$, vertices $b_1, \dots, b_t \in V(G)$ and $(c_{i,j} : i \in [s], j \in [t])$. We call $\bar{b} = (b_1, \dots, b_t)$ the *base* of the pattern.

Corollary 5.13. *For every monadically dependent graph class $\mathscr{C}$ and formula $\varphi(\bar{x}, z)$, there is some $t \in \mathbb{N}$ such that no graph $G \in \mathscr{C}$ contains a φ-pattern of order $t \times t$.*

We now prove Lemma 5.6, which we repeat for convenience.

Lemma 5.6. *For every monadically dependent graph class $\mathscr{C}$ and $r, s \in \mathbb{N}$, there are some $k, t \in \mathbb{N}$ with the following property. For every $G \in \mathscr{C}$ and finite set $A \subseteq V(G)$, there is $A^* \supseteq A$ of size $O_{\mathscr{C},r,s}(|A|^k)$ such that, for every s-definable flip H of G and every vertex v,*

$$\mathrm{Ball}_r^H(v) \cap A^* = \emptyset \quad \Longrightarrow \quad |\mathrm{ATypes}^H(\mathrm{Ball}_r^H(v)/A)| \leqslant t.$$

Proof. For each symmetric relation $R \subseteq \Pi^s \times \Pi^s$, let $\varphi_R(\bar{z}x, y)$ be a first-order formula expressing that x and y have distance at most r in the flip $G \oplus (\bar{z}, R)$, where $\bar{z}$ is an s-tuple. Its existence follows from Observation 5.3.

Let $k \geqslant 1$ be the bound from Corollary 5.13, such that no $G \in \mathscr{C}$ contains a φ_R-pattern of order $k \times k$, for any R. Let $Q\colon \mathbb{N} \to \mathbb{N}$ be as in Theorem 5.9, and set

$$t := Q\big(k \cdot |\Pi^s|^2\big) = Q(k \cdot 9^s).$$

Fix $G \in \mathscr{C}$ and a finite set $A \subseteq V(G)$. For every R as above and $N \in \{E, \neg E\}$, and every $\bar{b} \in A^k$, choose a maximum-height (φ_R, N)-pattern with base $\bar{b}$ in G. Each of them is an $\ell \times k$ pattern for some $\ell < k$ (possibly $\ell = 0$).

Let A^* consist of A and all grid-points of the chosen patterns. There are $O_s(1)$ many choices of R and $|A|^k$ many choices of $\bar{b}$, and each chosen pattern contributes at most k^2 vertices, so

$$|A^*| \leqslant |A| + O_s(1) \cdot |A|^k \cdot k^2 = O_{\mathscr{C},r,s}(|A|^k).$$

Fix a symmetric relation $R \subseteq \Pi^s \times \Pi^s$, tuple $\bar{p} \in V(G)^s$, and vertex $v \in V(G)$. Let $H = G \oplus (\bar{p}, R)$ and suppose $B := \mathrm{Ball}_r^H(v)$ satisfies $B \cap A^* = \emptyset$. We show that B realizes at most t atomic types over A in H, thus proving the lemma.

Suppose the contrary. By Theorem 5.9, there are tuples $\bar{u}'$ in B and $\bar{w}'$ in A, each of length $k \cdot |\Pi^s|^2$, that are atomically matched in H.

Let $\mathcal{P}$ be the partition of $V(G)$ into atomic types over $\bar{p}$; then H is a $\mathcal{P}$-flip of G and $|\mathcal{P}| \leqslant |\Pi^s|$. By the pigeonhole principle, we can pick corresponding subsequences $\bar{u}$ of $\bar{u}'$ and $\bar{w}$ of $\bar{w}'$ of length k, such that $\bar{u}$ is contained in a single part $X \in \mathcal{P}$ and $\bar{w}$ is contained in a single part $Y \in \mathcal{P}$. Write $\bar{u} = (u_1, \ldots, u_k)$ and $\bar{w} = (w_1, \ldots, w_k)$. Since H is a $\mathcal{P}$-flip of G, adjacency between X and Y is either preserved or complemented. Together with the matching condition in H, this gives a choice of $N \in \{E, \neg E\}$ such that

$$\begin{aligned} G &\models N(w_j, u_j) \qquad \text{for every } j \in [k], \\ G &\not\models N(w_{j'}, u_j) \quad \text{for } 1 \leqslant j < j' \leqslant k. \end{aligned}$$

Consider the maximum-height (φ_R, N)-pattern chosen for R and N and for base $\bar{b} := \bar{w} = (b_1, \ldots, b_k)$. Write $\bar{a}_1, \ldots, \bar{a}_\ell$ for its witnessing tuples and $c_{i,j}$ for its grid-points, where $i \in [\ell]$ and $j \in [k]$ (possibly $\ell = 0$). All these grid-points belong to A^*. Append a new row by setting

$$\bar{a}_{\ell+1} := \bar{p}v, \qquad c_{\ell+1,j} := u_j \quad \text{for every } j \in [k].$$

We verify that this yields a (φ_R, N)-pattern with the same base $\bar{b}$ and height $\ell + 1$, contradicting the maximality of ℓ. The column condition holds by the choice of N, $b_j = w_j$ and $c_{\ell+1,j} = u_j$:

$$\begin{aligned} G &\models N(b_j, c_{\ell+1,j}) \qquad \text{for every } j \in [k], \\ G &\not\models N(b_{j'}, c_{\ell+1,j}) \quad \text{for } 1 \leqslant j < j' \leqslant k. \end{aligned}$$

For every $j \in [k]$, we have $G \models \varphi_R(\bar{a}_{\ell+1}, c_{\ell+1,j})$, since $c_{\ell+1,j} \in B = \mathrm{Ball}_r^H(v)$. Moreover, for every $i \in [\ell]$, we have $G \not\models \varphi_R(\bar{a}_{\ell+1}, c_{i,j})$, since $c_{i,j} \in A^*$ and $B \cap A^* = \emptyset$. This proves the lemma. ■

5.4 Definable flips for balls

The following lemma is a corollary of the (p, q)-Theorem for bounded VC-dimension from [Mato4]. See [Sim15, Corollary 6.13] or [Tor23, full version, Theorem E.2] for self-contained proofs.

For a graph G and sets $S, X \subseteq V(G)$, we say that S *dominates* X if every vertex of $X - S$ has a neighbor in S, and that S *antidominates* X if every vertex of $X - S$ has a nonneighbor in S.

Lemma 5.14 (Bounded-VC duality)**.** *For every $d \in \mathbb{N}$ there is some $k \in \mathbb{N}$ with the following property. Let G be a graph of VC-dimension at most d, and let $X, Y \subseteq V(G)$ be two (possibly overlapping) sets with at least one of X, Y finite. Then there exists a set* $\mathrm{sample}(X, Y)$ *of size at most k such that at least one of the following holds:*

1. $\mathrm{sample}(X, Y)$ *is contained in Y and dominates X;*
2. $\mathrm{sample}(X, Y)$ *is contained in X and antidominates Y.*

We remark that in the sources we have given, the above lemma is stated either for finite set systems or finite binary relations. The translation to finite graphs is straightforward, by either viewing the graph as a set system of neighborhoods or working with the adjacency relation. Moreover, the statement lifts to infinite graphs if at least one of the sets X and Y is finite, as the

other one can be replaced by its finite subset representing all neighborhoods in the finite one. (However, the statement may fail if both X and Y are infinite, as witnessed by the bi-infinite half-graph with vertices $\mathbb{Z} \times \{1,2\}$ and edges between $(x,1)$ and $(y,2)$ if $x \leqslant y$.)

It was observed already in [Bon+22b] that Lemma 5.14 is useful to obtain definable flips. We will use it in a similar way to deduce the remaining result of Section 5.2.

Lemma 5.7. *For all $d, t \in \mathbb{N}$, there is some $z \in \mathbb{N}$ with the following property.*
For every graph G of VC-dimension at most d, radius $r \in \mathbb{N}$, finite set $A \subseteq V(G)$, and arbitrary set $C \subseteq V(G)$ satisfying

$$\mathrm{Ball}_r^G(C) \cap A = \emptyset \qquad \text{and} \qquad |\mathrm{ATypes}^G(\mathrm{Ball}_r^G(C)/A)| \leqslant t,$$

there exists a z-definable flip H of G in which $\mathrm{Ball}_{r+1}^H(C) \cap A = \emptyset$.

Note that, given the assumptions of the lemma, it is easy to find a (non-definable) $(2^t + t + 1)$-flip H' of G in which $\mathrm{Ball}_{r+1}^{H'}(C) \cap A = \emptyset$: group the vertices of A and of $\mathrm{Ball}_r^G(C)$ by their neighborhoods in the other set, which gives at most 2^t and t groups, respectively, and flip every group of A with every group of $\mathrm{Ball}_r^G(C)$ that it is complete to. So the challenge is really to get such a flip in a definable way.

We also remark that Lemma 5.7 can be seen as a finitization of Lemma 7.3 in the full version of [Gaj+23]. That lemma works in a more model-theoretic setting, where A is a (possibly infinite) model instead of a set. It uses a version of Lemma 5.14 which also works for infinite X, Y, under the stronger assumption that G excludes a half-graph of order d.

Proof of Lemma 5.7. Put $B := \mathrm{Ball}_r^G(C)$. Let $\mathcal{A}$ be the partition of A that groups vertices by their neighborhood in B, and let $\mathcal{B}$ be the partition of B that groups vertices by their neighborhood in A. By assumption $|\mathcal{B}| \leqslant t$. The neighborhood in B of a vertex of A is a union of parts of $\mathcal{B}$, so $|\mathcal{A}| \leqslant 2^t$. Moreover, every pair $P \in \mathcal{A}$, $Q \in \mathcal{B}$ is homogeneously connected: is either complete or anticomplete.

Let $k \leqslant O_d(1)$ be given by Lemma 5.14. Pick any set S which contains

- an arbitrary representative $\mathrm{rep}(P) \in P$ for each part $P \in \mathcal{A} \cup \mathcal{B}$, and
- the set $\mathrm{sample}(X, Y)$ supplied by Lemma 5.14, for every $X \in \mathcal{A}$ and $Y \in \mathcal{A} \cup \{V(G) - B\}$ (the lemma is applicable because $X \subseteq A$ is finite).

We can pick such a set with $|S| \leqslant z$ for some $z \leqslant O_{k,t}(1) \leqslant O_{d,t}(1)$. We will construct H as an S-flip of G; in particular, H is z-definable.

Let $\mathcal{P}_S$ be the partition of $V(G)$ into atomic types over S. Write $x \sim y$ when x and y are in the same $\mathcal{P}_S$-part. Observe the following two properties of $\mathcal{P}_S$.

Claim 5.1. *$\mathcal{P}_S$ refines $\mathcal{A}$ on A and $\mathcal{B}$ on B.*

Proof. If two vertices a and a' of A are in different $\mathcal{A}$-parts, then their adjacencies differ on some vertex $b \in B$. Let P be the $\mathcal{B}$-part containing b. The vertex $\mathrm{rep}(P) \in S$ has the same adjacencies to A as b and hence distinguishes a and a' in $\mathcal{P}_S$. Thus, $\mathcal{P}_S$ refines $\mathcal{A}$ on A and symmetrically also $\mathcal{B}$ on B. ◀

Claim 5.2. *Let $X \in \mathcal{A}$ and $Y \in \mathcal{A} \cup \{V(G) - B\}$. There do not exist vertices x, y, x', y' with:*

$$x \in X, \quad y \in Y, \quad x \sim x', \quad y \sim y', \quad Y \subseteq N^{\overline{G}}(x'), \quad X \subseteq N^G(y'). \tag{†}$$

Proof. Suppose that $\text{sample}(X, Y)$ lies in Y and dominates X. Then there is some $s \in \text{sample}(X, Y) \cap Y$ with $sx \in E(G)$. If $Y \subseteq N^{\overline{G}}(x')$, then $sx' \notin E(G)$. Thus, $x \nsim x'$.

The case when $\text{sample}(X, Y)$ lies in X and antidominates Y, is symmetric, by swapping X with Y and G with $\overline{G}$. Both cases are illustrated in Figure 5. ◀

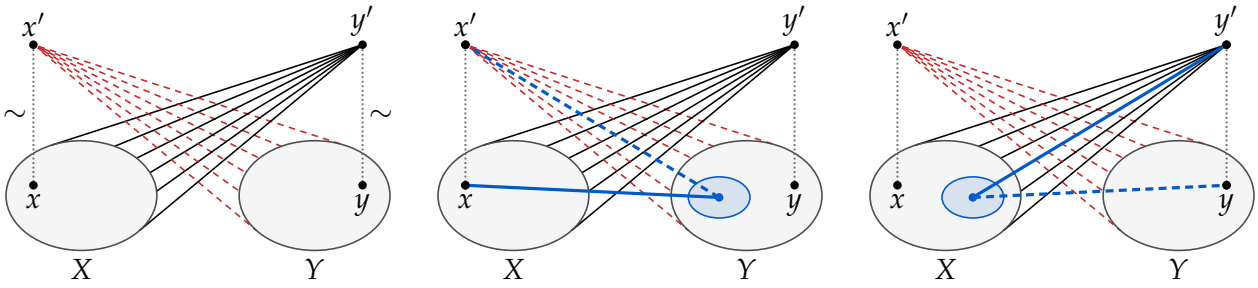


Figure 5: Left: the forbidden configuration (†). Middle: the case of the proof of Claim 5.2 where $\text{sample}(X, Y)$ (marked in blue) lies in Y. Right: the case where $\text{sample}(X, Y)$ lies in X.

Define the S-flip H of G where a pair (P, Q) of parts of $\mathcal{P}_S$ is flipped if and only if there is an edge in G with endpoints in P and in Q and also in A and in B. More formally, $H[P, Q] = G[P, Q]$ whenever $E(G) \cap PQ \cap AB = \emptyset$, and otherwise, $H[P, Q] = \overline{G}[P, Q]$.

The following three properties will yield the conclusion of the lemma:

(P1) A and $\text{Ball}_r^G(C)$ are nonadjacent in H. (proven in Claim 5.3)

(P2) $\text{Ball}_{r-1}^H(C) = \text{Ball}_{r-1}^G(C)$. (proven in Claim 5.5)

(P3) $\text{Ball}_{r-1}^G(C)$ and $V(G) - \text{Ball}_r^G(C)$ are nonadjacent in H. (proven in Claim 5.6)

If $r = 0$, then already P1 alone yields the conclusion. This justifies the implicit assumption $r \geqslant 1$ made in P2 and P3. By combining P2 and P3, we derive:

(P4) $\text{Ball}_r^H(C) \subseteq \text{Ball}_r^G(C)$.

The conclusion of the lemma thus follows:

$$A \cap \text{Ball}_{r+1}^H(C) = A \cap \text{Ball}_1^H\left(\text{Ball}_r^H(C)\right) \underset{\text{P4}}{\subseteq} A \cap \text{Ball}_1^H\left(\text{Ball}_r^G(C)\right) \underset{\text{P1}}{=} A \cap \text{Ball}_r^G(C) \underset{\text{ass.}}{=} \emptyset.$$

It only remains to establish P1–P3.

Claim 5.3 (P1)**.** *A and B are nonadjacent in H.*

Proof. By contrapositive, suppose $ab \in E(H)$ for some $a \in A$ and $b \in B$. By definition of H, $ab \in E(G)$ implies $ab \notin E(H)$, so $ab \notin E(G)$. Since $ab \in E(H)$, the parts of $\mathcal{P}_S$ containing a and b are flipped, that is, there is some $a'b' \in E(G)$ with $a' \in A$ and $b' \in B$, and

$$(a \sim a' \text{ and } b \sim b') \qquad \text{or} \qquad (a \sim b' \text{ and } b \sim a').$$

The first case cannot happen as $\mathcal{P}_S$ refines $\mathcal{A}$ on A and $\mathcal{B}$ on B, but the two sets $\{a, a'\}$ and $\{b, b'\}$ are inhomogeneous as $ab \notin E(G)$ and $a'b' \in E(G)$.

Thus, assume $a \sim b'$ and $b \sim a'$. Let $X, Y \in \mathcal{A}$ contain a', a, respectively. Then $a'b' \in E(G)$ implies that b' is complete to X, whereas $ab \notin E(G)$ implies that b is anticomplete to Y (both by homogeneity). This contradicts (†) with $(x, y, x', y') = (a', a, b, b')$. ◀

For brevity, we define $D := \mathrm{Ball}^G_{r-1}(C)$. Recall that $\mathcal{P}_S(v)$ is the part of $\mathcal{P}_S$ containing v.

Claim 5.4. *If b is adjacent to some $a \in A$ in G then $\mathcal{P}_S(b) \cap D = \emptyset$.*

Proof. The vertices in $B \cap \mathcal{P}_S(b)$ have the same neighborhood on A as b. In particular, they are all adjacent to a, whereas $D \subseteq B$ is anticomplete to A since $\mathrm{Ball}^G_1(D) = \mathrm{Ball}^G_r(C)$ is disjoint from A by assumption. ◀

Claim 5.5 (P2)**.** $\mathrm{Ball}^H_{r-1}(C) = D$.

Proof. The claim is trivially true for $r = 1$, hence assume $r \geqslant 2$. Consider any $ab \in E(G)$ with $a \in A$ and $b \in B$ witnessing a flipped pair of $\mathcal{P}_S$-parts $\mathcal{P}_S(a)$ and $\mathcal{P}_S(b)$.

1. *$\mathcal{P}_S(b)$ is disjoint from $\mathrm{Ball}^G_{r-2}(C)$:*

   This is implied by Claim 5.4, because $\mathrm{Ball}^G_{r-2}(C) \subseteq \mathrm{Ball}^G_{r-1}(C) = D$.

2. *$\mathcal{P}_S(a)$ is disjoint from $\mathrm{Ball}^G_{r-2}(C)$:*

   Let $Q \in \mathcal{B}$ contain b. Every vertex of Q is adjacent to a and hence lies in $B - D$. Thus, $\mathrm{rep}(Q) \in S$ is adjacent to a and nonadjacent to every vertex of $\mathrm{Ball}^G_{r-2}(C)$. Hence, $\mathrm{rep}(Q)$ witnesses that a is in a different $\mathcal{P}_S$-part than all the vertices in $\mathrm{Ball}^G_{r-2}(C)$.

This means that whenever two parts $\mathcal{P}_S(a)$ and $\mathcal{P}_S(b)$ are flipped, then neither of them intersects $\mathrm{Ball}^G_{r-2}(C)$. Hence, the vertices in $\mathrm{Ball}^G_{r-2}(C)$ are incident to the same edges in H as in G, which gives the desired conclusion $\mathrm{Ball}^H_{r-1}(C) = \mathrm{Ball}^G_{r-1}(C)$. ◀

For brevity, we define $\overline{B} := V(G) - B$.

Claim 5.6 (P3)**.** *D and $\overline{B}$ are nonadjacent in H.*

Proof. Otherwise, an edge $dz \in E(H)$ with $d \in D$ and $z \in \overline{B}$ is witnessed by some $ab \in E(G)$ with $a \in A$ and $b \in B$ (because D and $\overline{B}$ are nonadjacent in G). By Claim 5.4, the $\mathcal{P}_S$-part of b is disjoint from D, so necessarily $d \sim a$ and $z \sim b$. Let $X \in \mathcal{A}$ contain a. The vertex b is complete to X, whereas d is anticomplete to $\overline{B}$, contradicting (†) for $(X, \overline{B})$ with $(x, y, x', y') = (a, z, d, b)$. ◀

As described above, establishing P1–P3 finishes the proof. ■

5.5 Extension to tuples

In this subsection we prove Theorem 5.2. Recall that for $s \in \mathbb{N}$, an *s-flip* of a graph G is a $\mathcal{P}$-flip of G for some partition $\mathcal{P}$ of $V(G)$ into at most s parts. For $S \subseteq V(G)$, the partition $\mathcal{P}_S$ into atomic types over S has at most $3^{|S|}$ parts, so every s-definable flip is a 3^s-flip. The flip metric $\mathrm{dist}_{\mathcal{P}}$ associated with a partition $\mathcal{P}$ of $V(G)$ is the maximum of dist^H over all $\mathcal{P}$-flips H of G, and $\mathrm{Ball}^{\mathcal{P}}_r(v) := \{u \in V(G) : \mathrm{dist}_{\mathcal{P}}(u, v) \leqslant r\}$ denotes its balls. We use the same notation for the minimum distance between two tuples or sets.

Observation 5.15. *If $\mathcal{Q}$ refines $\mathcal{P}$, then every $\mathcal{P}$-flip is also a $\mathcal{Q}$-flip, and hence*

$$\mathrm{dist}_{\mathcal{Q}}(u, v) \geqslant \mathrm{dist}_{\mathcal{P}}(u, v).$$

We use the following metric-conversion result.

Lemma 5.16 ([Bon+25, Lemma 8]). *Let G be a graph and let $\mathcal{P}$ be a partition of $V(G)$. There is a refinement $\mathcal{Q}$ of $\mathcal{P}$ with $|\mathcal{Q}| \leqslant |\mathcal{P}| \cdot 2^{|\mathcal{P}|}$ and a $\mathcal{Q}$-flip H of G such that, for every $v \in V(G)$ and $r \in \mathbb{N}$,*

$$\mathrm{Ball}_r^H(v) \subseteq \mathrm{Ball}_{6r}^{\mathcal{P}}(v).$$

We now combine the flips from the Closure Theorem (Theorem 5.1) for the individual entries of a tuple into a single flip.

Theorem 5.2 (Closure Theorem for Tuples). *For every monadically dependent graph class $\mathscr{C}$, radius $r \in \mathbb{N}$, and tuple size $k \in \mathbb{N}$, there are some $c, s \in \mathbb{N}$ with the following property.*
For every $G \in \mathscr{C}$ and finite set $A \subseteq V(G)$ there is a set $\mathcal{H}$ of $O_{\mathscr{C},r,k}(|A|^c)$ many s-flips of G such that for every tuple $\bar{v} \in (V(G) - A)^k$ there is a flip $H \in \mathcal{H}$ with

$$\mathrm{dist}^H(\bar{v}, A) > r.$$

Proof. Let $c_0, s_0 \in \mathbb{N}$ be given by Theorem 5.1 for $\mathscr{C}$ and radius $6r$. Fix $G \in \mathscr{C}$ and a finite set $A \subseteq V(G)$, and let S and the sets $S(v)$ for $v \notin A$ be as in Theorem 5.1. For a tuple $\bar{v} = (v_1, \ldots, v_k) \in (V(G) - A)^k$, put

$$S(\bar{v}) \coloneqq S(v_1) \cup \cdots \cup S(v_k).$$

Then $S(\bar{v}) \subseteq S$ and $|S(\bar{v})| \leqslant k\, s_0$. As $\mathcal{P}_{S(\bar{v})}$ refines every $\mathcal{P}_{S(v_i)}$, Observation 5.15 gives

$$\mathrm{dist}^G_{S(\bar{v})}(v_i, A) \geqslant \mathrm{dist}^G_{S(v_i)}(v_i, A) > 6r \qquad \text{for every } i \in [k]. \tag{1}$$

For every $T \subseteq S$ with $|T| \leqslant k s_0$, the partition $\mathcal{P}_T$ has at most $3^{k s_0}$ parts. So Lemma 5.16 yields an s-flip H_T of G, for some $s \leqslant O_{\mathscr{C},r,k}(1)$ independent of T, such that

$$\mathrm{Ball}_r^{H_T}(x) \subseteq \mathrm{Ball}_{6r}^T(x) \qquad \text{for every } x \in V(G).$$

Let $\mathcal{H} \coloneqq \{H_T : T \subseteq S,\ |T| \leqslant k s_0\}$. Then $|\mathcal{H}| \leqslant (|S|+1)^{k s_0} \leqslant O_{\mathscr{C},r,k}(|A|^c)$ for $c \coloneqq c_0\, k\, s_0$.

Finally, let $\bar{v} = (v_1, \ldots, v_k) \in (V(G) - A)^k$ and $H \coloneqq H_{S(\bar{v})} \in \mathcal{H}$. For every $i \in [k]$, the ball $\mathrm{Ball}_r^H(v_i) \subseteq \mathrm{Ball}_{6r}^{S(\bar{v})}(v_i)$ is disjoint from A by (1). Thus $\mathrm{dist}^H(\bar{v}, A) > r$. ∎

6 Number of types over a set of parameters

In this section we prove Theorem 4.4, which bounds the number of types of k-tuples over a set U of parameters polynomially in $|U|$.

We use a consequence of locality of first-order logic, which bounds the number of types of tuples that are far from U. The following result is an adaptation of [PST18, Corollary 16]. We provide a self-contained proof in Section B.

Lemma 6.1. *For every $q, k, c \in \mathbb{N}$ there are numbers $r = r(q)$ and $T = T(q, k, c)$ with the following property. Let H be a graph expanded by at most c unary predicates, and let $U \subseteq V(H)$ be finite. Fix a tuple $\bar{u}$ of vertices of U of finite length. Let $F \coloneqq \{a \in V(H) : \mathrm{dist}^H(a, U) > r\}$. Then*

$$\left|\{\mathrm{tp}_q^H(\bar{a}\bar{u}/U) : \bar{a} \in F^k\}\right| \leqslant T(q, k, c).$$

We now prove Theorem 4.4 by combining Lemma 6.1 with the Closure Theorem for Tuples (Theorem 5.2).

Theorem 4.4. *Let $\mathscr{C}$ be a monadically dependent graph class and $k, q, c \in \mathbb{N}$. For every (possibly infinite) graph $G \in \mathscr{C}$, expansion G^+ of G by c unary predicates, and finite subset $U \subseteq V(G)$,*

$$\left|\mathrm{Types}_{k,q}^{G^+}(V(G)/U)\right| \leqslant O_{\mathscr{C},q,k,c}(1) \cdot \mathrm{poly}_{\mathscr{C},q,k}(|U|).$$

Proof. Fix $G \in \mathscr{C}$, its expansion G^+ by c unary predicates, and a finite set $U \subseteq V(G)$. Let $r = r(q)$ and T be given by Lemma 6.1, and put $\overline{U} := V(G) - U$. By the Closure Theorem for Tuples (Theorem 5.2) for $\mathscr{C}$, radius r, and k, there is some $s \leqslant O_{\mathscr{C},q,k}(1)$ and some set $\mathcal{H}$ of $\mathrm{poly}_{\mathscr{C},q,k}(|U|)$ many s-flips of G such that every $\bar{v} \in \overline{U}^k$ has some $H(\bar{v}) \in \mathcal{H}$ with $\mathrm{dist}^{H(\bar{v})}(\bar{v}, U) > r$. Note that s and $\mathcal{H}$ do not depend on the unary predicates of G^+. For each $H \in \mathcal{H}$, fix a partition $\mathcal{Q}_H$ of $V(G)$ into at most s parts such that H is a $\mathcal{Q}_H$-flip of G, and let H^+ be the expansion of H by the c unary predicates of G^+ and by at most s unary predicates marking the parts of $\mathcal{Q}_H$. The edge relation of G is quantifier-free definable in H^+, by flipping back the pairs of parts that are flipped in H. Substituting this definition for every edge atom turns each formula over G^+ with parameters from U into a formula over H^+ with the same parameters and quantifier rank that is satisfied by the same tuples. Hence, for every H and tuple $\bar{w}$ of vertices,

$$\mathrm{tp}_q^{G^+}(\bar{w}/U) \text{ depends only on } \mathrm{tp}_q^{H^+}(\bar{w}/U).$$

We want to bound the number of types of tuples $\bar{v} \in V(G)^k$ over U. By padding and permuting $2k$-tuples, up to a multiplicative factor of $O_k(1)$, it suffices to bound the number of types for the slightly longer tuples of the form $\bar{a}\bar{u}$ with $\bar{a} \in \overline{U}^k$ and $\bar{u} \in U^k$.

For $H \in \mathcal{H}$, let $F_H := \{w \in V(G) : \mathrm{dist}^H(w, U) > r\} \subseteq \overline{U}$. As $\bar{a} \in \overline{U}^k$ lies in $F^k_{H(\bar{a})}$ for some $H(\bar{a}) \in \mathcal{H}$, we have

$$\{\mathrm{tp}_q^{G^+}(\bar{a}\bar{u}/U) : \bar{a} \in \overline{U}^k,\ \bar{u} \in U^k\} = \bigcup_{H \in \mathcal{H}} \bigcup_{\bar{u} \in U^k} \{\mathrm{tp}_q^{G^+}(\bar{a}\bar{u}/U) : \bar{a} \in F_H^k\}.$$

It remains to bound each of the $|\mathcal{H}| \cdot |U|^k$ terms on the right. So fix $H \in \mathcal{H}$ and $\bar{u} \in U^k$. As H^+ has at most $c + s$ unary predicates, Lemma 6.1, applied to H^+, U, and $\bar{u}$, gives

$$\left|\{\mathrm{tp}_q^{H^+}(\bar{a}\bar{u}/U) : \bar{a} \in F_H^k\}\right| \leqslant T(q, k, c + s).$$

Since $\mathrm{tp}_q^{G^+}(\bar{a}\bar{u}/U)$ depends only on $\mathrm{tp}_q^{H^+}(\bar{a}\bar{u}/U)$, the same bound holds with G^+ in place of H^+. Hence,

$$\left|\{\mathrm{tp}_q^{G^+}(\bar{a}\bar{u}/U) : \bar{a} \in \overline{U}^k,\ \bar{u} \in U^k\}\right| \leqslant |\mathcal{H}| \cdot |U|^k \cdot T(q, k, c + s) \leqslant O_{\mathscr{C},q,k,c}(1) \cdot \mathrm{poly}_{\mathscr{C},q,k}(|U|),$$

where we use $T(q, k, c + s) \leqslant O_{\mathscr{C},q,k,c}(1)$, as $s \leqslant O_{\mathscr{C},q,k}(1)$, and $|\mathcal{H}| \leqslant \mathrm{poly}_{\mathscr{C},q,k}(|U|)$. ∎

Part II
Definable construction sequences

Recall the following definitions from Section 4.1.2.

Definition 4.5. Let G be a graph with vertex set V and let H be a partial graph with vertex set $W \subseteq V$. For $q, c, w \in \mathbb{N}$, we say that H is *definable* in G by formulas of quantifier rank q using c unary predicates and w parameters, if there is an expansion G^+ of G by c unary predicates and a set $S \subseteq V(G)$ of parameters with $|S| \leqslant w$, such that for every relation symbol Q of H (when viewed as a relational structure) of arity k there is some formula $\varphi_Q(x_1, \dots, x_k)$ of quantifier rank q with parameters from S such that

$$G^+ \models \varphi_Q(\bar{v}) \quad \iff \quad \bar{v} \in W^k \text{ and } H \models Q(\bar{v}) \qquad \text{for all } \bar{v} \in V^k.$$

Definition 4.6. Let G be a graph and $r, c, s, q_0 \in \mathbb{N}$. A construction sequence $G_1, \dots, G_\lambda$ of G is *radius-r definable* with c unary predicates, s parameters, and quantifier rank q_0 if for all $t \in [\lambda - 1]$ and $v \in V(G)$, the induced partial graph

$$G_t[\mathrm{Ball}_r^{G_{t+1}}(v)]$$

is definable in G by formulas of quantifier rank q_0 using c unary predicates and s parameters.

In this part, we establish the following main result.

Theorem 4.7. *Let $r \geqslant 1$ be an integer, $\varepsilon > 0$ be a rational, and $\mathscr{C}$ be a monadically dependent graph class. There is a constant $\delta \geqslant 1$ depending only on $\mathscr{C}$ such that every n-vertex graph $G \in \mathscr{C}$ has a construction sequence of G which is radius-r definable with $O(r/\varepsilon)$ unary predicates, $O_{\mathscr{C},r,\varepsilon}(n^{\delta\varepsilon})$ parameters, and quantifier rank 7, and has length $O_{\mathscr{C},r,\varepsilon}(n^{1+\delta\varepsilon})$. Moreover, there is an algorithm (independent of $\mathscr{C}$) which, given r, ε, and G, computes such a sequence, and runs in time $\widehat{O}_{\mathscr{C},r,\varepsilon}(n^3)$ whenever $G \in \mathscr{C}$.*

We will use the following variation on the definition of merge-width, originating in [Dra+26a].

Definition 6.2. A *chain* in a graph G is a decreasing sequence $V(G) = S_1 \supseteq \cdots \supseteq S_\ell = \emptyset$ of vertex sets. Its *radius-r width*, for $r \in \mathbb{N}$, is defined as

$$\max_{2 \leqslant i \leqslant \ell} \max_{v \in V(G)} \left| \mathrm{ATypes}^G(\mathrm{Ball}_r^{S_i}(v)/S_{i-1}) \right| = \max_{2 \leqslant i \leqslant \ell} \max_{v \in V(G)} \left| \mathrm{Ball}_r^{S_i}(v)/\mathcal{P}_{S_{i-1}} \right|,$$

where $\mathrm{Ball}_r^S(v)$ is defined above via S-flips and $\mathcal{P}_S$ is the partition into atomic types over S.

The offset between S_{i-1} and S_i in the definition of the radius-r width ensures that the definition is robust. For instance, a chain can be completed into one in which two consecutive sets differ by one element, by removing the elements of S_{i-1} one by one in any order, until S_i is reached; this transformation preserves the width parameters.

Chains can be converted to construction sequences and vice versa, with a loss in the radius-r width. Intuitively, a chain gives rise to a sequence of ever coarser partitions $\mathcal{P}_{S_1} \preccurlyeq \mathcal{P}_{S_2} \preccurlyeq \dots \preccurlyeq \mathcal{P}_{S_\ell}$, which can then be converted into a sequence of partial graphs (see [Dra+26a] for details).

We use the following key result from [Dre+26]. It says that in monadically dependent classes, chains of constant length and small width can be computed efficiently.

Theorem 6.3 ([Dre+26, Theorem 4.4])**.** *There is an algorithm which, given an n-vertex graph G, radius $r \in \mathbb{N}$, and rational $\varepsilon > 0$, computes a chain of G of length $O(r/\varepsilon)$ and radius-r width $O_{\mathscr{C},r,\varepsilon}(n^\varepsilon)$, and runs in time $\widehat{O}_{\mathscr{C},r}(n^3)$, where $\mathscr{C}$ is any monadically dependent graph class containing G.*

To prove Theorem 4.7, we turn these chains into construction sequences using the following main technical contribution of this part.

Lemma 6.4. *For every n-vertex graph G and its chain of length ℓ there is a construction sequence $G_1, \dots, G_\lambda$ of G such that the following properties hold for all $r \geqslant 1$. Let k_r denote the radius-$(5r+1)$ width of the chain and $w_r := k_r \cdot \pi_G(k_r)$.*

1. *The resulting construction sequence has radius-r width at most w_r and length $\lambda \leqslant O(nw_1)$.*
2. *For every $t \in [\lambda - 1]$ and $v \in V(G)$, the induced partial graph*

$$G_t[\mathrm{Ball}_r^{G_{t+1}}(v)]$$

   is definable in some expansion of G with $2\ell + 1$ unary predicates, using formulas of quantifier rank 7 with w_r parameters.
3. *The construction sequence can be computed in time $O(\ell n^3)$, given G and its chain.*

Proof of Theorem 4.7. Let $\mathscr{C}$ be a monadically dependent graph class containing G, and let d be an upper bound on the VC-dimension of graphs in $\mathscr{C}$, which exists by Observation 3.2. We apply Theorem 6.3 with radius $5r+1$ and ε. This computes a chain of length $\ell \leqslant O((5r+1)/\varepsilon) = O(r/\varepsilon)$ in time $\widehat{O}_{\mathscr{C},r}(n^3)$. We then apply Lemma 6.4 to G and this chain, which takes time $O(\ell n^3)$. So the total running time is

$$\widehat{O}_{\mathscr{C},r}(n^3) + O(\ell n^3) = \widehat{O}_{\mathscr{C},r,\varepsilon}(n^3).$$

The chain has radius-$(5r+1)$ width $k_r \leqslant O_{\mathscr{C},r,\varepsilon}(n^\varepsilon)$, and $\pi_G(k_r) \leqslant (k_r+1)^d$ by the Sauer–Shelah lemma (Lemma 3.1), so

$$w_r = k_r \cdot \pi_G(k_r) \leqslant (k_r+1)^{d+1} \leqslant O_{\mathscr{C},r,\varepsilon}(n^{(d+1)\varepsilon}).$$

We set $\delta := d+1$. The sequence is radius-r definable with $2\ell + 1 \leqslant O(r/\varepsilon)$ unary predicates, w_r parameters, and rank 7. As $k_1 \leqslant k_r$, it has length $O(nw_1) \leqslant O(nw_r) \leqslant O_{\mathscr{C},r,\varepsilon}(n^{1+\delta\varepsilon})$. ∎

Proof strategy for Lemma 6.4 The proof of Lemma 6.4 occupies the remainder of Part II. It combines constructions from [Dra+26a] and [DT25a], together with a careful analysis of definability. We revisit several variants of merge-width studied in those papers. In particular, we use the following variation of the definition of merge-width.

Definition 6.5. A *merge sequence* of a graph G is a sequence $(\mathcal{P}_1, R_1), \dots, (\mathcal{P}_\ell, R_\ell)$ such that:

1. $\mathcal{P}_1$ is the partition into singletons and $\mathcal{P}_\ell = \{V(G)\}$,
2. $\mathcal{P}_1 \preccurlyeq \dots \preccurlyeq \mathcal{P}_\ell$,
3. $R_1 \subseteq \dots \subseteq R_\ell \subseteq \binom{V(G)}{2}$,
4. for every $i \in [\ell]$ and $X, Y \in \mathcal{P}_i$, the set $XY - R_i$ is contained in either $E(G)$ or in $E(\overline{G})$, that is, the pair X, Y is homogeneous once pairs in R_i are ignored.

The *radius-r width* of a merge sequence as above is

$$\max_{2 \leqslant i \leqslant \ell} \max_{v \in V(G)} \left| \mathrm{Ball}_r^{R_i}(v) / \mathcal{P}_{i-1} \right|. \tag{2}$$

Here the offset in the indices of R_i and $\mathcal{P}_{i-1}$ is intentional. Every construction sequence naturally induces a merge sequence of the same radius-r width for all $r \in \mathbb{N}$, and vice versa (see [DT25b; DT25a] for details).

7 Merge-width with definable stages

In this section, we prove the following lemma.

Lemma 7.1. *For every graph G and chain of length ℓ there is a merge sequence $(\mathcal{Q}_1, R_1), \dots, (\mathcal{Q}_\ell, R_\ell)$ with the following properties.*

1. *For all $r \geqslant 1$, the radius-r width of the merge sequence is bounded by $k \cdot \pi_G(k)$, where k is the radius-$(5r+1)$ width of the given chain.*
2. *For each $i \in [\ell]$, the relations $\equiv_{\mathcal{Q}_i}$ and R_i are definable in some expansion G^i of G with two unary predicates, using two fixed first-order formulas $\varphi_\equiv(x,y)$ and $\varphi_R(x,y)$ of quantifier rank at most 5.*
3. *The merge sequence can be computed in time $O(\ell n^3)$, given G and its chain.*

7.1 Preliminaries

The proof of Lemma 7.1 uses ideas developed in [Dra+26a]. We first recall the required notions and lemmas from that paper.

Lemma 7.2 (folklore)**.** *For every graph G, at least one of G and $\overline{G}$ has diameter at most 3.*

Proof. Suppose G has diameter greater than 2. Then there are non-adjacent vertices u, v with no common neighbor in G. In $\overline{G}$, uv is an edge and all vertices are neighbors of either u or v. Thus $\overline{G}$ has diameter at most 3. ∎

If B is a bipartite graph with parts X, Y and edge set E, its *bipartite complement*, denoted $\overline{B}$, is the bipartite graph with parts X, Y and edge set $XY - E$.

Lemma 7.3 ([Dra+26a, Lemma A.3])**.** *Let $B = (X, Y, E)$ be a bipartite graph with bipartite complement $\overline{B}$, and pick arbitrary vertices $x \in X$, $y \in Y$. Partition X according to adjacency to y, i.e. $X = X_+ \uplus X_-$ with $X_+ = N(y)$ and $X_- = X - N(y)$, and similarly $Y = Y_+ \uplus Y_-$ according to x. Then there is a bipartite graph $B' = (X, Y, E')$ obtained from B by some $\{X_+, X_-, Y_+, Y_-\}$-flip satisfying:*

$$\text{for all } uv \in E', \quad \operatorname{dist}^B(u,v) \leqslant 5 \quad \text{and} \quad \operatorname{dist}^{\overline{B}}(u,v) \leqslant 5.$$

Note that [Dra+26a, Lemma A.3] states the above lemma with 6 instead of 5; however, the two statements are equivalent, since if $uv \in E'$ then u and v lie in opposite parts of the bipartite graphs B', B, and $\overline{B}$, and thus $\operatorname{dist}(u,v) \leqslant 6$ is equivalent to $\operatorname{dist}(u,v) \leqslant 5$ for parity reasons.

Definition 7.4. Fix a graph G and a partition $\mathcal{P}$ of $V(G)$. We define $\mathcal{P} \wedge T$ to be the refinement of $\mathcal{P}$ where each part $P \in \mathcal{P}$ is split into atomic types over $T - P$ (i.e. into sets of the form $P \cap Q$, where Q is the set of all vertices of G with a given neighborhood in $T - P$).

In this definition, we choose not to divide a part P into types over $T \cap P$. This is for technical convenience only, to ensure that $\mathcal{P} \wedge T = \mathcal{P}$ whenever $\mathcal{P}$ is the partition with only one part.

Lemma 7.5 ([Dra+26a, Corollary 5.4])**.** *Let G be a graph, $\mathcal{P} \preccurlyeq \mathcal{Q}$ partitions of $V(G)$, and T a transversal of $\mathcal{P}$. Suppose $|\operatorname{Ball}^{\mathcal{Q}}_{r+1}(v)/\mathcal{P}| \leqslant k$ for every $v \in V(G)$. Then for every $v \in V(G)$, we have*

$$\left|\operatorname{Ball}^{\mathcal{Q}}_r(v)/(\mathcal{P} \wedge T)\right| \leqslant k \cdot \pi_G(k).$$

7.2 Constructing the merge sequence

Recall that for a graph G and disjoint vertex sets $X, Y \subseteq V(G)$, by $G[X,Y]$ we denote the bipartite graph with parts X and Y, and edges connecting pairs of vertices which are adjacent in G. Furthermore, $G[X,X]$ denotes the bipartite graph whose parts are two copies of X, and edges connect two vertices in opposite parts if the corresponding vertices in G are either adjacent or equal. We call the edges of $G[X,X]$ connecting the two copies of each vertex in X *diagonal edges*. Thus, $\overline{G}[X,X]$ is the bipartite complement $\overline{G[X,X]}$ of $G[X,X]$ extended additionally with the diagonal edges.

Definition 7.6. Let G be a graph and $\mathcal{P}$ a partition of $V = V(G)$. Define $R = R(G,\mathcal{P}) \subseteq \binom{V}{2}$ so that for all parts $X, Y \in \mathcal{P}$ (possibly $X = Y$) and pairs uv with $u \in X$, $v \in Y$, and $u \neq v$, we have $uv \in R$ if and only if

$$\operatorname{dist}^{G[X,Y]}(u,v) \leqslant 5 \quad \text{and} \quad \operatorname{dist}^{\overline{G}[X,Y]}(u,v) \leqslant 5.$$

We first describe the construction of the merge sequence that underlies the proof of Lemma 7.1. Consider a graph G and its chain $V(G) = S_1 \supseteq \cdots \supseteq S_\ell = \emptyset$. Let $\mathcal{P}_i$ be the partition into atomic types over S_i. Fix an arbitrary linear order $<$ on $V(G)$ giving rise to representative sets $T_i := \{\min_< P : P \in \mathcal{P}_i\}$. Then $T_1 \supseteq T_2 \supseteq \cdots \supseteq T_\ell$. For $i = 1, \ldots, \ell$ define

$$R_i := R(G,\mathcal{P}_i) \qquad \text{and} \qquad \mathcal{Q}_i := \mathcal{P}_i \wedge T_i.$$

The following lemmas establish that

$$(\mathcal{Q}_1, R_1), \ldots, (\mathcal{Q}_\ell, R_\ell)$$

is a merge sequence satisfying the conditions stated in Lemma 7.1.

Lemma 7.7. *Suppose G is a graph and $\mathcal{P}$ is the partition of $V(G)$ into atomic types over some set $S \subseteq V(G)$. Then the relation $R := R(G,\mathcal{P})$ is definable in G expanded with one unary predicate marking S, using a fixed first-order formula $\varphi_R(x,y)$ of quantifier rank at most 5.*

Proof. The following formula $\theta_S(x,y)$, using a unary predicate marking S and with one quantifier, defines the equivalence relation $\equiv_{\mathcal{P}}$:

$$\theta_S(x,y) := x = y \vee \big(\neg S(x) \wedge \neg S(y) \wedge \forall z\,(S(z) \to (E(x,z) \leftrightarrow E(y,z)))\big). \tag{3}$$

The following formula $\eta_S(x,y;a,b)$ says that x and y are in opposite parts of the bipartite graph $G[X,Y]$, where X and Y are the parts of $\mathcal{P}$ containing a and b:

$$\eta_S(x,y;a,b) := (\theta_S(x,a) \wedge \theta_S(y,b)) \vee (\theta_S(y,a) \wedge \theta_S(x,b)).$$

The following formulas $e_S^+(x,y;a,b)$ and $e_S^-(x,y;a,b)$ say that x and y are adjacent in the bipartite graph $G[X,Y]$ or $\overline{G}[X,Y]$, respectively, where X and Y are the parts of $\mathcal{P}$ containing a and b:

$$\begin{aligned} e_S^+(x,y;a,b) &:= \eta_S(x,y;a,b) \wedge (x = y \vee E(x,y)),\\ e_S^-(x,y;a,b) &:= \eta_S(x,y;a,b) \wedge (x = y \vee \neg E(x,y)). \end{aligned}$$

For $\sigma \in \{+,-\}$, define

$$\rho_S^\sigma(x,y) := \bigvee_{p \in \{1,3,5\}} \exists z_1 \cdots \exists z_{p-1} \bigwedge_{i=0}^{p-1} e_S^\sigma(z_i, z_{i+1}; x, y),$$

where $z_0 := x$ and $z_p := y$. This formula expresses the existence of a walk of length 1, 3, or 5 between opposite copies of x and y in $G[X,Y]$ for $\sigma = +$, and in $\overline{G}[X,Y]$ for $\sigma = -$, where $X := \mathcal{P}(x)$ and $Y := \mathcal{P}(y)$. Each step switches sides, also when $X = Y$; an equality step uses a diagonal edge. Since every path between opposite sides of a bipartite graph has odd length, these formulas express precisely the two distance bounds defining R. Consequently,

$$\varphi_R(x,y) := x \neq y \wedge \rho_S^+(x,y) \wedge \rho_S^-(x,y)$$

defines $R(G,\mathcal{P})$. Each e_S^σ has quantifier rank 1, and each disjunct in ρ_S^σ quantifies at most four intermediate vertices. Thus φ_R is a fixed first-order formula of quantifier rank at most 5. ■

Lemma 7.8. *Fix a graph G and partitions $\mathcal{P}$ and $\mathcal{P}'$ of $V(G)$ such that $\mathcal{P}$ refines $\mathcal{P}'$. Then*

$$R(G,\mathcal{P}) \subseteq R(G,\mathcal{P}').$$

Proof. Choose any $uv \in R(G,\mathcal{P})$. Let

$$X := \mathcal{P}(u), \quad X' := \mathcal{P}'(u), \quad Y := \mathcal{P}(v), \quad Y' := \mathcal{P}'(v)$$

be the parts of $\mathcal{P}$ and $\mathcal{P}'$ that contain u and v. As $\mathcal{P}$ refines $\mathcal{P}'$, we have $X \subseteq X'$ and $Y \subseteq Y'$. Hence $G[X,Y]$ and $\overline{G}[X,Y]$ are induced subgraphs of $G[X',Y']$ and $\overline{G}[X',Y']$, respectively. Thus, by the definition of R,

$$\begin{aligned}\operatorname{dist}^{G[X',Y']}(u,v) &\leqslant \operatorname{dist}^{G[X,Y]}(u,v) \leqslant 5,\\ \operatorname{dist}^{\overline{G}[X',Y']}(u,v) &\leqslant \operatorname{dist}^{\overline{G}[X,Y]}(u,v) \leqslant 5.\end{aligned}$$

This proves $uv \in R(G,\mathcal{P}')$, as required. ■

Lemma 7.9. *Fix a graph G and a partition $\mathcal{P}$ of $V(G)$ and let $R := R(G,\mathcal{P})$. For every radius $r \in \mathbb{N}$ and vertex $v \in V(G)$ we have*

$$\operatorname{Ball}_r^R(v) \subseteq \operatorname{Ball}_{5r}^{\mathcal{P}}(v).$$

Proof. It suffices to show that for every $uv \in R$ we have $\operatorname{dist}_{\mathcal{P}}(u,v) \leqslant 5$. Let $X, Y \in \mathcal{P}$ be the parts containing u and v, respectively. Fix an arbitrary $\mathcal{P}$-flip H of G. Then the bipartite graph $H[X,Y]$ is either $G[X,Y]$ or $\overline{G}[X,Y]$. By the definition of R, this graph contains a u–v path of length at most 5. If $X \neq Y$, then this path is also a path in H. If $X = Y$, we can project this path to a length $\leqslant 5$ path in H by contracting diagonal edges. Hence, in both cases, $\operatorname{dist}^H(u,v) \leqslant 5$, which finishes the proof. ■

Lemma 7.10. *Fix a graph G and a partition $\mathcal{P}$ of $V(G)$ and let T be a transversal of $\mathcal{P}$ and $R := R(G,\mathcal{P})$. Then $\mathcal{Q} := \mathcal{P} \wedge T$ is* homogeneous modulo R, *i.e., for all $X, Y \in \mathcal{Q}$ either*

$$XY - R \subseteq E(G) \quad \text{or} \quad XY - R \subseteq E(\overline{G}).$$

Proof. We construct a $\mathcal{Q}$-flip H of G with $E(H) \subseteq R$. This will yield the conclusion. Indeed, let $X, Y \in \mathcal{Q}$. Since H is a $\mathcal{Q}$-flip of G, either

$$E(H) \cap XY = E(G) \cap XY \quad \text{or} \quad E(H) \cap XY = E(\overline{G}) \cap XY.$$

Since $E(H) \subseteq R$, no pair in $XY - R$ is an edge of H. Hence $XY - R \subseteq E(\overline{G})$ in the first case, and $XY - R \subseteq E(G)$ in the second. It remains to define H. We separately describe how to flip a part X with itself, and how to flip two distinct parts X and Y with each other.

- For every part $X \in \mathcal{P}$, by Lemma 7.2 we have that either $G[X]$ or $\overline{G}[X]$ has diameter at most 3. If $G[X]$ has diameter at most 3, set $H[X] := \overline{G}[X]$, otherwise set $H[X] := G[X]$.

  Suppose now that $uv \in E(H[X])$ for some $u, v \in X$. We consider the case when $G[X]$ has diameter at most 3 and $H[X] = \overline{G}[X]$, the other case being symmetric. Then $\mathrm{dist}^{G[X]}(u,v) \leqslant 3$ and $\mathrm{dist}^{\overline{G}[X]}(u,v) = 1$. The path of length at most 3 in $G[X]$ lifts to a path of length at most 3 between opposite copies in $G[X,X]$ after adding at most one diagonal step to make its length odd. Similarly, the edge uv in $\overline{G}[X]$ lifts to a path of length 1 between opposite copies in $\overline{G}[X,X]$. This shows that $uv \in R$, and thus $E(H[X]) \subseteq R$.

- For every pair of distinct parts $X, Y \in \mathcal{P}$, let x, y be their representatives in T and apply Lemma 7.3 to $G[X,Y]$. Use the resulting bipartite graph as $H[X,Y]$. The statement of the lemma implies that $E(H[X,Y]) \subseteq R$.

  The partition $\mathcal{Q}$ restricted to $X \cup Y$, denote it by $\mathcal{Q}[X,Y]$, refines the neighborhood partitions of X over y and of Y over x, so $H[X,Y]$ is a $\mathcal{Q}[X,Y]$-flip of $G[X,Y]$. Altogether, H is a $\mathcal{Q}$-flip of G with $E(H) \subseteq R$. ■

We now prove Lemma 7.1.

Lemma 7.1. *For every graph G and chain of length ℓ there is a merge sequence $(\mathcal{Q}_1, R_1), \ldots, (\mathcal{Q}_\ell, R_\ell)$ with the following properties.*

1. *For all $r \geqslant 1$, the radius-r width of the merge sequence is bounded by $k \cdot \pi_G(k)$, where k is the radius-$(5r+1)$ width of the given chain.*
2. *For each $i \in [\ell]$, the relations $\equiv_{\mathcal{Q}_i}$ and R_i are definable in some expansion G^i of G with two unary predicates, using two fixed first-order formulas $\varphi_{\equiv}(x,y)$ and $\varphi_R(x,y)$ of quantifier rank at most 5.*
3. *The merge sequence can be computed in time $O(\ell n^3)$, given G and its chain.*

Proof. Let $V(G) = S_1 \supseteq \cdots \supseteq S_\ell = \emptyset$ be the given chain. Let $\mathcal{P}_i$ be the partition into atomic types over S_i. Fix an arbitrary linear order $<$ on $V(G)$ giving rise to representative sets $T_i := \{\min_< P : P \in \mathcal{P}_i\}$. Then $T_1 \supseteq T_2 \supseteq \cdots \supseteq T_\ell$. Define

$$R_i := R(G, \mathcal{P}_i) \qquad \text{and} \qquad \mathcal{Q}_i := \mathcal{P}_i \wedge T_i.$$

We claim that

$$(\mathcal{Q}_1, R_1), \ldots, (\mathcal{Q}_\ell, R_\ell)$$

is a merge sequence. Indeed:

1. Since $\mathcal{P}_1$ is the partition into singletons, so is $\mathcal{Q}_1$. It follows from $\mathcal{P}_\ell = \{V(G)\}$ and the definition of $\mathcal{P}_\ell \wedge T_\ell$ that also $\mathcal{Q}_\ell = \{V(G)\}$.
2. For $i < j$ we have that $T_i \supseteq T_j$ and $\mathcal{P}_i \preccurlyeq \mathcal{P}_j$, which implies that $\mathcal{P}_i \wedge T_i \preccurlyeq \mathcal{P}_j \wedge T_j$. Hence,

   $$\mathcal{Q}_1 \preccurlyeq \mathcal{Q}_2 \preccurlyeq \ldots \preccurlyeq \mathcal{Q}_\ell$$

   is a sequence of ever coarser partitions.
3. By Lemma 7.8, we have that

   $$R_1 \subseteq R_2 \subseteq \cdots \subseteq R_\ell.$$

4. By Lemma 7.10, we have that for all $P, Q \in \mathcal{Q}_i$ either

$$PQ - R_i \subseteq E(G) \quad \text{or} \quad PQ - R_i \subseteq E(\overline{G}).$$

Fix $r \in \mathbb{N}$, and let k be the radius-$(5r+1)$ width of the given chain. For every $i \in \{2, \dots, \ell\}$ and $v \in V(G)$, the definition of chain width gives

$$|\operatorname{Ball}^{\mathcal{P}_i}_{5r+1}(v)/\mathcal{P}_{i-1}| \leqslant k.$$

Applying Lemma 7.5 to $\mathcal{P}_{i-1} \preccurlyeq \mathcal{P}_i$ and its transversal T_{i-1} yields

$$|\operatorname{Ball}^{\mathcal{P}_i}_{5r}(v)/\mathcal{Q}_{i-1}| = |\operatorname{Ball}^{\mathcal{P}_i}_{5r}(v)/(\mathcal{P}_{i-1} \wedge T_{i-1})| \leqslant k \cdot \pi_G(k). \tag{4}$$

By Lemma 7.9, we have

$$\operatorname{Ball}^{R_i}_{r}(v) \subseteq \operatorname{Ball}^{\mathcal{P}_i}_{5r}(v). \tag{5}$$

Combining (4) and (5) yields the desired bound on the radius-r width of the merge sequence:

$$|\operatorname{Ball}^{R_i}_{r}(v)/\mathcal{Q}_{i-1}| \leqslant k \cdot \pi_G(k).$$

Finally, consider the graph G equipped with two unary predicates marking S_i and T_i. The relation R_i is definable in this expansion by Lemma 7.7 by a formula of quantifier rank 5. The relation $\equiv_{\mathcal{Q}_i}$ is defined easily using S_i and T_i by a formula of quantifier rank 2.

We now analyze the running time of a straightforward algorithm computing the resulting merge sequence. Put $n := |V(G)|$. For each $i \in [\ell]$, keep every vertex of S_i in a singleton part and group the remaining vertices by their adjacency vectors over S_i. This computes $\mathcal{P}_i$ in $O(n(|S_i| + 1)) \leqslant O(n^2)$ time. Scanning the resulting partition computes T_i in $O(n)$ time. Similarly, compute the partition $\mathcal{Q}_i$ in time $O(n(|T_i| + 1)) \leqslant O(n^2)$. Thus computing and outputting $\mathcal{P}_i$, $\mathcal{Q}_i$, S_i, and T_i takes $O(n^2)$ time for a fixed i.

Compute R_i by running breadth-first searches to depth five from every vertex in each of $G[X, Y]$ and $\overline{G}[X, Y]$, for all pairs $X, Y \in \mathcal{P}_i$, allowing $X = Y$. The cost per pair is $O((|X| + |Y|)|X||Y|)$. We have:

$$\sum_{X,Y \in \mathcal{P}_i} (|X| + |Y|)|X||Y| = 2 \sum_{X,Y \in \mathcal{P}_i} |X||Y|^2 = 2 \sum_{X \in \mathcal{P}_i} |X| \cdot \sum_{Y \in \mathcal{P}_i} |Y|^2 \leqslant 2n \Big(\sum_{Y \in \mathcal{P}_i} |Y| \Big)^2 = 2n^3,$$

thus R_i can be computed in time $O(n^3)$, and the whole merge sequence in time $O(\ell n^3)$. ■

8 Construction sequences with locally definable partial graphs

For the purpose of model checking, we need to work with construction sequences rather than merge sequences. The construction of [DT25b, Lemma 3.3] shows how to convert a merge sequence to a construction sequence. However, a single step in a merge sequence is translated into a potentially long sequence of steps in a construction sequence. When aiming at defining each step of the construction sequence in the original graph G, we need to remember the result of the entire sequence of all the steps performed so far. We cannot hope to obtain construction sequences in which each part is definable by a formula with fewer than $n/3$ parameters, even for the edgeless graph on n vertices. Therefore, in the following lemma, we only aim at proving that *locally* – in each ball of radius r in the graph formed by the resolved vertex pairs – the current set of resolved edges and each part of the current partition are definable in G by formulas of

bounded quantifier rank, involving a number of parameters which is bounded by the radius-r merge-width of the input merge sequence. Recall the precise definition of definability stated in Definition 4.5, and used in the statement below. It is crucial that when considering merge sequences of bounded length, we can retain the unary predicates describing every stage, and thus obtain formulas using only boundedly many unary predicates.

Lemma 8.1. *For every n-vertex graph G and merge sequence $(\mathcal{P}_1, R_1), \ldots, (\mathcal{P}_\ell, R_\ell)$ of G, there is a construction sequence $G_1, \ldots, G_\lambda$ of G such that the following properties hold for all $r \geqslant 1$. Let w_r be the radius-r width of the merge sequence.*

1. *The obtained construction sequence has radius-r width at most w_r, and has length $\lambda \leqslant O(nw_1)$.*
2. *Suppose that for every $i \in [\ell]$, the relations $\equiv_{\mathcal{P}_i}$ and R_i are definable in some expansion of G with 2 unary predicates, using two fixed first-order formulas of quantifier rank at most 5.*

   Then, for every $t \in [\lambda - 1]$ and $v \in V(G)$, the induced partial graph $G_t[\mathrm{Ball}_r^{G_{t+1}}(v)]$ is definable in some expansion of G with $2\ell + 1$ unary predicates, using formulas of quantifier rank 7 with w_r parameters.
3. *The construction sequence can be computed in time $O(\ell n^2)$, given an n-vertex graph G and a merge sequence of length ℓ.*

Proof. We convert the merge sequence to a construction sequence by following the proof of [DT25b, Lemma 3.3]. Set $V := V(G)$, $R_{\ell+1} := \binom{V}{2}$, and $\mathcal{P}_{\ell+1} = \{V\}$.

Start with the singleton partition $\mathcal{P}$ of V and empty set R of resolved vertex pairs. The construction proceeds in *stages* $j = 1, \ldots, \ell$. The invariant that is maintained during stage j is that in each step, the current partition $\mathcal{P}$ and current set R of resolved vertex pairs satisfy:

$$\mathcal{P}_j \preccurlyeq \mathcal{P} \preccurlyeq \mathcal{P}_{j+1} \tag{6}$$

$$R \subseteq R_{j+1} \tag{7}$$

$$XY - R \subseteq E(G) \quad \text{or} \quad XY - R \subseteq E(\overline{G}) \quad \text{for all } X, Y \in \mathcal{P}. \tag{8}$$

We start the process with $\mathcal{P} := \mathcal{P}_1$ being the partition into singletons and $R := \emptyset$. Stage j proceeds as follows. First consider all unordered pairs of parts $X, Y \in \mathcal{P}_j$, allowing $X = Y$, for which $XY \subseteq R_{j+1}$. If there are unresolved vertex pairs in XY (i.e. $XY \not\subseteq R$), resolve the pair X, Y with the common sign – positively if $XY - R \subseteq E(G)$ and negatively if $XY - R \subseteq E(\overline{G})$ – and update the set of resolved vertex pairs $R := R \cup XY$. Then repeatedly merge pairs A, B of parts of the current partition for which $A \cup B$ is contained in one part of $\mathcal{P}_{j+1}$, until the current partition $\mathcal{P}$ becomes $\mathcal{P}_{j+1}$; if $j = \ell$, terminate. This results in a construction sequence of G.

Fix $r \geqslant 1$. We now upper bound the radius-r width of the obtained construction sequence by the radius-r width w_r of the merge sequence. Fix a stage j in the merge sequence and a step of the construction above in that phase, with current partition $\mathcal{P}$ and current set of resolved pairs $R \subseteq \binom{V}{2}$. By (7), for every $v \in V$ we have:

$$\mathrm{Ball}_r^R(v) \subseteq \mathrm{Ball}_r^{R_{j+1}}(v). \tag{9}$$

As $\mathrm{Ball}_r^{R_{j+1}}(v)$ intersects at most w_r parts of $\mathcal{P}_j$ and $\mathcal{P}$ is coarser than $\mathcal{P}_j$ by (6), also $\mathrm{Ball}_r^R(v)$ intersects at most w_r parts of $\mathcal{P}$, proving the required bound.

The obtained construction sequence is *effective*: each resolve operation introduces at least one new resolved vertex pair. By [DT25b, Lemma 3.4] (and the comment before it), the length λ of the construction sequence is at most $n(2w_1 + 1) \leqslant O(nw_1)$.

A straightforward implementation processes each stage in time $O(n^2)$. Namely, for each pair X, Y of parts of $\mathcal{P}_j$, test whether $XY \subseteq R_{j+1}$ by scanning all vertex pairs $uv \in XY$, and record the sign with which X, Y should be resolved. As this processes each pair $uv \in \binom{V}{2}$ once, this takes time $O(n^2)$ in total. Similarly, the merge phase can be implemented in time $O(n)$. Altogether, the construction sequence is computable in time $O(\ell n^2)$. This proves the first and last conditions of the lemma.

We now prove the second condition of the lemma, regarding definability. Fix a step $G_t \to G_{t+1}$ of the obtained construction sequence, and let j be the stage in which this step is performed. Let $\mathcal{P}$ and R denote the partition and the set of resolved vertex pairs in G_t, and let R_+ denote the set of resolved vertex pairs in G_{t+1}; then $R \subseteq R_+ \subseteq R_{j+1}$.

Let $\varphi^i_{\equiv}(x,y)$ and $\varphi^i_R(x,y)$ be formulas defining $\equiv_{\mathcal{P}_i}$ and R_i, respectively, as given by the assumption. For $1 \leqslant j < \ell$, the set

$$\bigcup \{XY \,:\, X, Y \in \mathcal{P}_j, XY \subseteq R_{j+1}\}$$

is defined by the formula

$$\gamma_j(x,y) \coloneqq x \neq y \ \wedge \forall a \forall b \left((\varphi^j_{\equiv}(x,a) \wedge \varphi^j_{\equiv}(y,b) \wedge a \neq b) \ \to \ \varphi^{j+1}_R(a,b) \right),$$

which has quantifier rank at most 7. Consequently, at the start of stage j, the resolved relation, which we denote R'_j, is defined by

$$\Gamma_j(x,y) \coloneqq \bigvee_{i=1}^{j-1} \gamma_i(x,y).$$

This formula uses $2j$ unary predicates and has quantifier rank at most 7.

Choose one parameter $p_X \in X$ for every part $X \in \mathcal{P}_j$. The part X is then defined by $\varphi^j_{\equiv}(x, p_X)$, which uses 2 unary predicates, one parameter, and has quantifier rank at most 5.

Fix a vertex $v \in V$. Denote

$$B \coloneqq \mathrm{Ball}_r^{R_+}(v) = \mathrm{Ball}_r^{G_{t+1}}(v) \subseteq \mathrm{Ball}_r^{R_{j+1}}(v).$$

Mark the set B with a unary predicate U_B.

We will show how to define in G the relation $R \cap \binom{B}{2}$ and the sets $X \cap B$, for each $X \in \mathcal{P}$, using Boolean combinations of the formulas defined above, using the $2j+1$ distinguished unary predicates, and at most w_r parameters. Once $R \cap \binom{B}{2}$ is defined, the set of edges and non-edges of G_t, restricted to $\binom{B}{2}$, can be recovered by further intersecting with $E(G)$ and $E(\overline{G})$, respectively. This will therefore define the induced substructure $G_t[B]$, as required in the second condition of the lemma.

We first focus on defining $R \cap \binom{B}{2}$. Note that R is the union of R'_j and all already processed rectangles XY with $X, Y \in \mathcal{P}_j, XY \subseteq R_{j+1}$. Therefore, R is defined by the disjunction of $\Gamma_j(x,y)$ and disjuncts of the form

$$(x \neq y) \wedge \Big((\varphi^j_{\equiv}(x, p_X) \wedge \varphi^j_{\equiv}(y, p_Y)) \vee (\varphi^j_{\equiv}(x, p_Y) \wedge \varphi^j_{\equiv}(y, p_X)) \Big).$$

To define the intersection of R with $\binom{B}{2}$, add a conjunct $U_B(x) \wedge U_B(y)$, and keep only those disjuncts which involve parts X, Y that intersect B. As there are at most w_r many such parts, this results in a formula using w_r parameters.

Each current part $X \in \mathcal{P}$ is a union of parts of $\mathcal{P}_j$, so its intersection with B is defined by the formula

$$U_B(x) \wedge \bigvee_Y \varphi^j_{\equiv}(x, p_Y)$$

where Y ranges over all parts $Y \in \mathcal{P}_j$ with $Y \cap B \neq \emptyset$ and $Y \subseteq X$, with empty disjunction interpreted as *false*. The resulting formula uses only w_r parameters. This completes the proof of the lemma. ■

We now prove Lemma 6.4.

Lemma 6.4. *For every n-vertex graph G and its chain of length ℓ there is a construction sequence $G_1, \ldots, G_\lambda$ of G such that the following properties hold for all $r \geqslant 1$. Let k_r denote the radius-$(5r+1)$ width of the chain and $w_r := k_r \cdot \pi_G(k_r)$.*

1. *The resulting construction sequence has radius-r width at most w_r and length $\lambda \leqslant O(nw_1)$.*
2. *For every $t \in [\lambda - 1]$ and $v \in V(G)$, the induced partial graph*

$$G_t[\mathrm{Ball}_r^{G_{t+1}}(v)]$$

   is definable in some expansion of G with $2\ell + 1$ unary predicates, using formulas of quantifier rank 7 with w_r parameters.
3. *The construction sequence can be computed in time $O(\ell n^3)$, given G and its chain.*

Proof. The empty graph is immediate, so assume G is nonempty. Apply Lemma 7.1 to the given chain to obtain a merge sequence with ℓ stages, each definable by two unary predicates and fixed formulas of quantifier rank at most 5. Fix $r \geqslant 1$ and let k_r be the radius-$(5r+1)$ width of the chain. Then the radius-r width w_r of this merge sequence is at most $k_r \cdot \pi_G(k_r)$. Apply Lemma 8.1 to obtain a construction sequence of length $\lambda \leqslant O(nw_1)$ and radius-r width at most w_r, with local definitions using at most w_r parameters, $2\ell + 1$ unary predicates, and formulas of quantifier rank at most 7. Both algorithms take $O(\ell n^3)$ time and use only G and the given chain. ■

Part III
Model checking with type trees

In this part we prove:

Theorem 4.3. *There is an algorithm which, given a sentence φ and an n-vertex graph G together with its construction sequence of length λ, determines whether $G \models \varphi$ in time $\lambda \cdot n^2 \cdot \mathrm{poly}_\varphi(K)$, where K is the radius-r depth-d width of the construction sequence, for some r and d which are computable from φ.*

9 Preliminaries

In Part III, unless stated otherwise, G denotes a partial graph. We slightly extend the notions of partial graphs and of construction sequences.

Partial graphs. To track parts over time, we give them names. Fix a countably infinite set of unary symbols, called *colors*. We regard a partial graph G as a relational structure with the two symmetric binary relations E and N and a unary relation Y for every color Y. The vertices of color Y form its color class Y^G, and the nonempty color classes are the parts of G. We write $|G| := |V(G)|$ and $\|G\| := |V(G)| + |R(G)|$ for the size of the representation of a partial graph G. The *distance matrix* of G is the $|G| \times |G|$ matrix D_G with entries $D_G(u,v) := \mathrm{dist}^G(u,v) \in \mathbb{N} \cup \{\infty\}$. Passing it as an additional input to the algorithms below saves time, as distances need not be recomputed repeatedly.

Construction sequences. For colors P, Q and a partial graph G, we define three partial graphs $\mathrm{merge}_{P,Q}(G)$ (for $P \neq Q$), $\mathrm{resolve}^+_{P,Q}(G)$, and $\mathrm{resolve}^-_{P,Q}(G)$, as the partial graph H obtained from G by leaving all relations of G unchanged, apart from redefining the following:

$$\begin{aligned}
\mathrm{merge}_{P,Q}(G)\colon &\quad P^H := P^G \cup Q^G, \qquad Q^H := \emptyset && \text{for } P \neq Q,\\
\mathrm{resolve}^+_{P,Q}(G)\colon &\quad E(H) := E(G) \cup \big(P^G Q^G - N(G)\big),\\
\mathrm{resolve}^-_{P,Q}(G)\colon &\quad N(H) := N(G) \cup \big(P^G Q^G - E(G)\big).
\end{aligned}$$

After either resolve, every pair in $P^G Q^G$ is resolved. In this representation, every step of a construction sequence $G_1, \ldots, G_\lambda$ (Section 2) has the form $G_{t+1} = \mathrm{merge}_{P,Q}(G_t)$, $\mathrm{resolve}^+_{P,Q}(G_t)$, or $\mathrm{resolve}^-_{P,Q}(G_t)$ for some colors P, Q.

10 Distance logic

Following [DT25b; DT25a] (and earlier work by [GKS17]), we define a logic called dist-FO, extending first-order logic by distance atoms, and a suitable notion of rank. We then define compact representations of types tailored to this logic.

10.1 Distance logic and distance rank

First-order distance logic. The logic dist-FO extends first-order logic by additionally allowing the following two types of atoms:

For each radius $r \in \mathbb{N}$, dist-FO introduces a *binary distance atom* $\mathrm{dist}(x,y) \leqslant r$ expressing that the distance between x and y is at most r.

For each radius $r \in \mathbb{N}$ and color Y, dist-FO introduces a *unary distance atom* $\mathrm{dist}(x,Y) < r$ expressing that the distance from x to Y is smaller than r, equivalent to $\exists y\ Y(y) \wedge \mathrm{dist}(x,y) < r$.

In both cases, we call r the *radius* of the distance atom. Distance atoms are expressible in first-order logic, but only with quantifiers. Treating them as atoms lets us measure distances without increasing the quantifier rank, as in [GKS17].

Horizon functions. We fix sufficiently fast-growing horizon functions as in [DT26]:

$$\rho^-(k,q) := 9^{(k+q+1)q} \qquad \text{and} \qquad \rho^+(k,q) := 9^{(k+q)(q+1)}. \tag{10}$$

They satisfy the two identities

$$\rho^+(k,q) = 9^k \rho^-(k,q) \quad \text{for all } k,q \in \mathbb{N}, \qquad \rho^-(k,q) = 9^q \rho^+(k+1,q-1) \quad \text{for } q \geqslant 1. \tag{11}$$

In particular, as $9^q \geqslant 2$,

$$\rho^+(k,q) \geqslant \rho^-(k,q) \geqslant 2\rho^+(k+1,q-1) \quad \text{for all } k \in \mathbb{N} \text{ and } q \geqslant 1. \tag{12}$$

Distance rank. Let $k,q \in \mathbb{N}$. A dist-FO formula has *distance rank* (k,q) if it is a Boolean combination of

- first-order atoms: relational symbols in the signature or equality, applied to variables,
- binary distance atoms with radius $r \leqslant \rho^-(k,q)$,
- unary distance atoms with radius $r \leqslant \rho^+(k,q)$,
- formulas
$$\exists y \left(\mathrm{dist}(x_i,y) \leqslant \rho^+(k+1,q-1) \wedge \varphi(\bar{x}y) \right),$$
where $q \geqslant 1$, x_i is a variable in $\bar{x}$, and $\varphi(\bar{x}y)$ has distance rank $(k+1,q-1)$.

A formula with distance rank (k,q) has quantifier rank at most q. Observe by (12) that the distance atoms have progressively shorter radii as one traverses the quantifiers inwards.

Formulas of distance rank (k,q) are called *local* formulas of distance rank (k,q) in [DT25b; DT26], as opposed to the more general formulas which also allow the unrestricted quantification $\exists y\ \psi(\bar{x}y)$. Since we only consider local formulas in this paper, we omit this adjective.

10.2 Type trees

In the context of usual first-order logic, the *type* of quantifier rank q of a tuple $\bar{a}$ in a structure A is defined as the set of all formulas of quantifier rank q that are satisfied by $\bar{a}$ in A (see Section 3). It is well known [Sco79] that for structures over a finite relational signature, such types can be identified with *trees*, also called *Hintikka types*, defined recursively as follows:

$$\mathrm{tree}_q^A(\bar{a}) = \begin{cases} \mathrm{atp}^A(\bar{a}) & \text{if } q = 0, \\ \{\mathrm{tree}_{q-1}^A(\bar{a}b) \,:\, b \in V(A)\} & \text{if } q > 0. \end{cases}$$

In the setting of dist-FO formulas of distance rank (k,q), the (local) *type of distance rank* (k,q) of a tuple $\bar{a}$ is defined [DT25b] as the set of (local) dist-FO formulas of distance rank (k,q) satisfied by $\bar{a}$. In this section, we adapt the construction of trees to this setting. This is useful in the algorithmic context, as trees are much smaller objects than types.

Type trees. In the two definitions below, let G be a partial graph, $k, q \in \mathbb{N}$, and $\bar{a} \in V(G)^{\leqslant k}$, and put $\ell := |\bar{a}|$. The *atomic* (k,q)*-type*, denoted $\mathrm{atp}_{k,q}(\bar{a})$, is the set of all atomic formulas φ of distance rank (k,q) with free variables among $x_1, \ldots, x_\ell$, such that $G \models \varphi(\bar{a})$. Note that this set is finite, as an atom mentioning a color that does not occur in G is not satisfiable in G.

Define the (k,q)*-type tree* of $\bar{a}$ as

$$\mathrm{tree}_{k,q}(\bar{a}) := \big(\mathrm{atp}_{k,q}(\bar{a}), \mathcal{E}_{k,q}(\bar{a})\big)$$

where $\mathcal{E}_{k,q}(\bar{a}) = \emptyset$ if $q = 0$, and otherwise

$$\mathcal{E}_{k,q}(\bar{a}) := \{\big(i, \mathrm{tree}_{k+1,q-1}(\bar{a}b)\big) \,:\, i \in [\ell],\ b \in V(G),\ \mathrm{dist}(a_i, b) \leqslant \rho^+(k+1, q-1)\}.$$

Note that type trees are "local" in the sense that new vertices are always quantified close to previous ones (see Lemma 10.4). The members of $\mathcal{E}_{k,q}(\bar{a})$ are the *extensions* of $\mathrm{tree}_{k,q}(\bar{a})$. The position i of an extension (i, μ) is called the *guard*, and we say that i *admits* b if $\mathrm{dist}(a_i, b) \leqslant \rho^+(k+1, q-1)$. We view $\mathrm{tree}_{k,q}(\bar{a})$ as a labeled, rooted tree of depth q, with the root labeled $\mathrm{atp}_{k,q}(\bar{a})$ and the root of each type tree μ of an extension $(i, \mu) \in \mathcal{E}_{k,q}(\bar{a})$ attached as a child of the root via an edge labeled i. We call ℓ the *arity* of the type tree, and (k,q) its *rank*.

Representation. Type trees are represented in the standard RAM model in a canonical (say, sorted, duplicate-free) form, so that equality of two type trees τ_1, τ_2 can be tested in time $O(\|\tau_1\| + \|\tau_2\|)$. By $\|\tau\|$ we denote the size of this representation (which includes k, q, and ℓ).

Observation 10.1 (canonical form)**.** *A type tree given by a representation whose sets need not be sorted or duplicate-free can be brought into canonical form in polynomial time.*

The next lemma states that the type trees indeed represent the types. It is the reason for computing type trees in the first place.

Lemma 10.2 (evaluation)**.** *Let $\varphi(\bar{x})$ be a* dist-FO *formula of distance rank (k,q), G be a partial graph and $\bar{a} \in V(G)^\ell$ with $\ell = |\bar{x}| \leqslant k$. Then whether $G \models \varphi(\bar{a})$ holds is decidable in polynomial time, given $\mathrm{tree}^G_{k,q}(\bar{a})$ and φ.*

Notation. For a (k,q)-type tree $\tau = \mathrm{tree}^G_{k,q}(\bar{a})$ and a dist-FO formula $\varphi(\bar{x})$ of distance rank (k,q), we write $\tau \models \varphi$, if $G \models \varphi(\bar{a})$. The lemma implies that this does not depend on G and $\bar{a}$, only on the type tree τ.

Proof of Lemma 10.2. By induction on q. Write $\tau := \mathrm{tree}_{k,q}(\bar{a})$. By definition of distance rank, φ is a Boolean combination of atoms and of existential components, and we evaluate these components separately. An atom holds of $\bar{a}$ if and only if it occurs in the root label $\mathrm{atp}_{k,q}(\bar{a})$ of τ. Now let $q \geqslant 1$ and put $H := \rho^+(k+1, q-1)$. Every existential component χ of φ can be written as

$$\exists y\ (\mathrm{dist}(x_i, y) \leqslant H \wedge \psi(\bar{x}y))$$

with $i \in [\ell]$ and ψ of distance rank $(k+1, q-1)$. By the inductive assumption, whether $G \models \psi(\bar{c})$ holds is determined by $\mathrm{tree}_{k+1,q-1}(\bar{c})$ and ψ, and by the notation above, this holds if $\mathrm{tree}_{k+1,q-1}(\bar{c}) \models \psi$. Unfolding the semantics, then applying the inductive assumption, and finally the definition of $\mathcal{E}_{k,q}(\bar{a})$, we obtain

$$\begin{aligned} G \models \chi(\bar{a}) &\iff \exists b \in V(G)\colon\ \mathrm{dist}(a_i, b) \leqslant H \text{ and } G \models \psi(\bar{a}b) \\ &\iff \exists b \in V(G)\colon\ \mathrm{dist}(a_i, b) \leqslant H \text{ and } \mathrm{tree}_{k+1,q-1}(\bar{a}b) \models \psi \\ &\iff \exists \mu\colon\ (i, \mu) \in \mathcal{E}_{k,q}(\bar{a}) \text{ and } \mu \models \psi. \end{aligned}$$

So χ is decided by evaluating ψ, via the inductive assumption, on the extensions of τ with guard i.

This proof yields a polynomial-time algorithm. ∎

Corollary 10.3. *Fix $\ell \leqslant k$ and q, a partial graph G, and $\bar{a}, \bar{b} \in V(G)^\ell$. Then $\mathrm{tree}^G_{k,q}(\bar{a}) = \mathrm{tree}^G_{k,q}(\bar{b})$ if and only if $G \models \varphi(\bar{a}) \Leftrightarrow \varphi(\bar{b})$ for every* dist-FO *formula $\varphi(\bar{x})$ of distance rank (k,q) with $|\bar{x}| = \ell$.*

Proof. The rightwards implication follows from Lemma 10.2. The converse is proved by an easy induction on q. ∎

Lemma 10.4 (locality)**.** *Let $k \geqslant 1$, $q \in \mathbb{N}$, let G be a partial graph, and let $a \in V(G)$. Every node at depth j of $\mathrm{tree}_{k,q}(a)$ is $\mathrm{tree}_{k+j,q-j}(\bar{v})$ for some tuple $\bar{v}$ from $\mathrm{Ball}_{\rho^+(k,q)}(a)$, and every color mentioned in the root label intersects this ball.*

Proof. By the definition of extensions and induction on j, a node at depth j is $\mathrm{tree}_{k+j,q-j}(\bar{v})$ for a tuple $\bar{v} = (v_0, \ldots, v_j)$ with $v_0 = a$ in which every v_i with $i \geqslant 1$ is within distance $\rho^+(k+i, q-i)$ of an earlier entry. By induction on i and the triangle inequality, $\mathrm{dist}(a, v_i) \leqslant \sum_{i'=1}^{i} \rho^+(k+i', q-i')$. The root label mentions colors only in atoms $Y(x_i)$ and $\mathrm{dist}(x_i, Y) < r$ with $r \leqslant \rho^+(k+j, q-j)$, and such an atom holds of $\bar{v}$ only if a vertex of Y is within distance r of v_i. Hence every entry of $\bar{v}$, and some vertex of every color in an atom in the root label, is within distance

$$\sum_{i=1}^{j} \rho^+(k+i, q-i) + \rho^+(k+j, q-j) \leqslant \rho^+(k,q)$$

of a, where the inequality uses $\rho^+(k+i, q-i) \leqslant 2^{-i}\rho^+(k,q)$, which follows by applying (12) i times. ∎

10.3 Operations on type trees

Lemma 10.5 (ordered coordinate restriction)**.** *There is an operation $\restriction$ on type trees such that for all $k, q \in \mathbb{N}$, partial graphs G, tuples $\bar{w} \in V(G)^\ell$ for $\ell \leqslant k$ and ordered lists J of distinct positions in $[\ell]$,*

$$\mathrm{tree}_{k,q}(\bar{w}) \restriction J = \mathrm{tree}_{k,q}(\bar{w} \restriction J).$$

Given τ and J one can compute $\tau \restriction J$ in time polynomial in $\|\tau\|$.

Proof. By induction on q. Let $J = (j_1, \ldots, j_p)$, and $H := \rho^+(k+1, q-1)$. Renaming variables gives

$$\mathrm{atp}_{k,q}(\bar{w} \restriction J) = \{\psi(x_1, \ldots, x_p) \,:\, \psi(x_{j_1}, \ldots, x_{j_p}) \in \mathrm{atp}_{k,q}(\bar{w})\},$$

which settles the root label. Let $q \geqslant 1$, and put $J^+ := (j_1, \ldots, j_p, \ell+1)$, so that $(\bar{w} \restriction J)\, b = (\bar{w}b) \restriction J^+$ for every vertex b. Unfolding the definition of $\mathcal{E}_{k,q}(\bar{w} \restriction J)$, then applying the inductive assumption, and finally the definition of $\mathcal{E}_{k,q}(\bar{w})$, we obtain

$$\begin{aligned}
\mathcal{E}_{k,q}(\bar{w} \restriction J) &= \{\big(i,\ \mathrm{tree}_{k+1,q-1}((\bar{w}b) \restriction J^+)\big) \,:\, i \in [p],\ b \in V(G),\ \mathrm{dist}(w_{j_i}, b) \leqslant H\} \\
&= \{\big(i,\ \mathrm{tree}_{k+1,q-1}(\bar{w}b) \restriction J^+\big) \,:\, i \in [p],\ b \in V(G),\ \mathrm{dist}(w_{j_i}, b) \leqslant H\} \\
&= \{(i,\ \mu \restriction J^+) \,:\, i \in [p],\ (j_i, \mu) \in \mathcal{E}_{k,q}(\bar{w})\}.
\end{aligned}$$

So the guard j_i is renumbered to i, as the variable x_{j_i} is renamed to x_i in the root. The right-hand side depends only on $\mathrm{tree}_{k,q}(\bar{w})$ and J and we define $\tau \restriction J$ accordingly. It is computed by one pass over τ, followed by restoring canonical form. As no subtree grows, this takes polynomial time by Observation 10.1. ∎

Lemma 10.6 (rank shift)**.** *There is a polynomial-time algorithm which maps* $\mathrm{tree}_{k,q}(\bar{a})$ *to* $\mathrm{tree}_{k+1,q-1}(\bar{a})$ *for all* $k \in \mathbb{N}$, $q \geqslant 1$, *partial graphs* G *and nonempty* $\bar{a} \in V(G)^{\leqslant k}$.

Proof. Since $\mathrm{dist}(a_1, a_1) = 0$, the guard 1 admits a_1, so $\mathrm{tree}_{k,q}(\bar{a})$ has an extension $(1, \mu)$. Every such μ is $\mathrm{tree}_{k+1,q-1}(\bar{a}b)$ for some vertex b, and applying the algorithm of Lemma 10.5 yields $\mu \restriction (1, \ldots, |\bar{a}|) = \mathrm{tree}_{k+1,q-1}(\bar{a})$. ∎

The following is an adaptation of [DT26, Lemma 3.1].

Lemma 10.7 (far composition)**.** *There is an operation* $*$ *on type trees such that for all* $k, q \in \mathbb{N}$, *partial graphs* G *and nonempty tuples* $\bar{a}_1, \bar{a}_2$ *with* $|\bar{a}_1\bar{a}_2| \leqslant k$ *and* $\mathrm{dist}(\bar{a}_1, \bar{a}_2) > \rho^-(k, q)$,

$$\mathrm{tree}_{k,q}(\bar{a}_1) * \mathrm{tree}_{k,q}(\bar{a}_2) = \mathrm{tree}_{k,q}(\bar{a}_1\bar{a}_2).$$

Given τ_1, τ_2 *one can compute* $\tau_1 * \tau_2$ *in time* $\mathrm{poly}_{k,q}(\|\tau_1\| + \|\tau_2\|)$.

Proof. By induction on q. Write $\ell_i := |\bar{a}_i|$ and $\tau_i := \mathrm{tree}_{k,q}(\bar{a}_i)$. Note that ℓ_i can be deduced from τ_i.

Consider an atom of distance rank (k, q) that mentions variables of both $\bar{a}_1$ and $\bar{a}_2$. If the atom was true, we would have $\mathrm{dist}(\bar{a}_1, \bar{a}_2) \leqslant 1$ for an equality, edge, or non-edge atom, and $\mathrm{dist}(\bar{a}_1, \bar{a}_2) \leqslant \rho^-(k, q)$ for a binary distance atom, each contradicting $\mathrm{dist}(\bar{a}_1, \bar{a}_2) > \rho^-(k, q)$. Every other atom mentions only variables of $\bar{a}_1$ or only variables of $\bar{a}_2$, and whether it holds is recorded in $\mathrm{atp}_{k,q}(\bar{a}_1)$ or $\mathrm{atp}_{k,q}(\bar{a}_2)$. Hence $\mathrm{atp}_{k,q}(\bar{a}_1\bar{a}_2)$ is determined by $\mathrm{atp}_{k,q}(\bar{a}_1)$ and $\mathrm{atp}_{k,q}(\bar{a}_2)$. This settles the root label.

Let now $q \geqslant 1$. Put $H := \rho^+(k+1, q-1)$, and let $*$ also denote the operation given by the inductive assumption for rank $(k+1, q-1)$. Write $\bar{a}_1\bar{a}_2 = (c_1, \ldots, c_{\ell_1+\ell_2})$, and split $\mathcal{E}_{k,q}(\bar{a}_1\bar{a}_2)$ by guards into $\mathcal{E}_1 \cup \mathcal{E}_2$, where

$$\mathcal{E}_1 := \{(i,\ \mathrm{tree}_{k+1,q-1}(\bar{a}_1\bar{a}_2b))\ :\ i \in [\ell_1],\ b \in V(G),\ \mathrm{dist}(c_i, b) \leqslant H\},$$
$$\mathcal{E}_2 := \{(i,\ \mathrm{tree}_{k+1,q-1}(\bar{a}_1\bar{a}_2b))\ :\ \ell_1 < i \leqslant \ell_1 + \ell_2,\ b \in V(G),\ \mathrm{dist}(c_i, b) \leqslant H\}.$$

Consider $\mathcal{E}_1$, and let i, b satisfy its conditions. Then $\mathrm{dist}(\bar{a}_1, b) \leqslant H$, and the triangle inequality, the hypothesis, (12) and (11) give

$$\mathrm{dist}(\bar{a}_1b, \bar{a}_2) \ \geqslant\ \mathrm{dist}(\bar{a}_1, \bar{a}_2) - \mathrm{dist}(\bar{a}_1, b) \ >\ \rho^-(k, q) - H \ \geqslant\ H \ \geqslant\ \rho^-(k+1, q-1).$$

Hence, we may apply the inductive assumption for rank $(k+1, q-1)$ to the tuples $\bar{a}_1b$ and $\bar{a}_2$. Moreover, let J be the list of positions with $\bar{a}_1\bar{a}_2b = (\bar{a}_1b\bar{a}_2) \restriction J$. Applying in this order the property of J, Lemma 10.5 and the inductive assumption, and finally the definition of $\mathcal{E}_{k,q}(\bar{a}_1)$, we obtain

$$\begin{aligned}
\mathcal{E}_1 &= \{(i,\ \mathrm{tree}_{k+1,q-1}(\bar{a}_1b\bar{a}_2) \restriction J)\ :\ i \in [\ell_1],\ b \in V(G),\ \mathrm{dist}(c_i, b) \leqslant H\} \\
&= \{(i,\ (\mathrm{tree}_{k+1,q-1}(\bar{a}_1b) * \mathrm{tree}_{k+1,q-1}(\bar{a}_2)) \restriction J)\ :\ i \in [\ell_1],\ b \in V(G),\ \mathrm{dist}(c_i, b) \leqslant H\} \\
&= \{(i,\ (\mu * \mathrm{tree}_{k+1,q-1}(\bar{a}_2)) \restriction J)\ :\ (i, \mu) \in \mathcal{E}_{k,q}(\bar{a}_1)\},
\end{aligned}$$

where $\mathrm{tree}_{k+1,q-1}(\bar{a}_2)$ is obtained from τ_2 by Lemma 10.6. Symmetrically, with the guard i of $\bar{a}_2$ renumbered to $\ell_1 + i$ as the variables of its root are shifted by ℓ_1, and without reordering since $\bar{a}_1\bar{a}_2b$ already ends with b,

$$\mathcal{E}_2 = \{(\ell_1 + i,\ \mathrm{tree}_{k+1,q-1}(\bar{a}_1) * \mu)\ :\ (i, \mu) \in \mathcal{E}_{k,q}(\bar{a}_2)\}.$$

The right-hand sides depend only on τ_1 and τ_2 and we define $\tau_1 * \tau_2$ accordingly. By Lemmas 10.5 and 10.6 and the inductive assumption, the work per input extension is $\mathrm{poly}_{k,q}(\|\tau_1\| + \|\tau_2\|)$, and the recursion has fixed depth q. Restoring canonical form by Observation 10.1 stays within this bound, which proves the stated running time. ∎

11 Scatter types and far trees

The type tree of a tuple $\bar{a}$ records extensions only by vertices near $\bar{a}$. A resolve, however, can bring vertices far from $\bar{a}$ within guard range, so lifting type trees through a resolve (Section 12) also requires the type trees of such far away vertices. For a vertex b far from $\bar{a}$, the type tree of $\bar{a}b$ is obtained from those of $\bar{a}$ and b by far composition (Lemma 10.7). As for quantifiers over far witnesses in the proof of Gaifman's locality theorem, it therefore suffices to know which type trees are realized far from $\bar{a}$. We collect them in the *far tree* of $\bar{a}$.

The far tree is not determined by the type tree of $\bar{a}$ alone. It also depends on global information about G. As in Gaifman's proof, this global information is captured by an analogue of the basic local sentences, which we call the *scatter type* of G. In this section we show that the far tree of $\bar{a}$ is computable from the type tree of $\bar{a}$ and the scatter type of G.

Section 11.1 defines scatter types, and Section 11.2 uses them to compute the far tree. The model checking algorithm maintains the type trees of all vertices, and Lemma 11.3 shows that the scatter type can be computed from these.

11.1 Scatter types

Let G be a partial graph. A set S of vertices of G is *r-scattered* if all vertices in S have pairwise distance larger than r in G. Fix, once and for all, a linear order on $V(G)$. The algorithms below use it without further mention. For a set $X \subseteq V(G)$ and $r \in \mathbb{N}$, let $S_r(X)$ denote the inclusion-wise maximal r-scattered subset of X constructed by the greedy process that repeatedly adds the least available vertex. If ν is a type tree of distance rank (k, q) and arity one, write

$$[\nu]^G := \{b \in V(G) \,:\, \mathrm{tree}_{k,q}(b) = \nu\}$$

for the set of its realizations.

Definition 11.1 (scatter type). Let G be a partial graph and $k, q \in \mathbb{N}$. The *(k,q)-scatter type* $\mathrm{stp}^G_{k,q}$ is the function with domain

$$\big\{(r, \mathrm{tree}_{k+i,q-i}(b)) \;\big|\; i \in [q],\ b \in V(G),\ r \in [9^{k+i}\rho^-(k+i, q-i)]\big\}$$

defined by

$$\mathrm{stp}^G_{k,q}(r, \nu) := \min\big(|S_r([\nu]^G)|,\ k+q\big).$$

For $q = 0$ the domain is empty and $\mathrm{stp}_{k,0}$ is the empty function. We write $\mathrm{stp}_{k,q} = \mathrm{stp}^G_{k,q}$.

We represent a scatter type canonically as a sorted list of its triples $(r, \nu, \mathrm{stp}_{k,q}(r, \nu))$, using the canonical representation of ν. Its representation size, denoted by $\|\mathrm{stp}_{k,q}\|$, is the total size of this list, including the encodings of the radii and values.

Lemma 11.2. *For every partial graph G and $q \geqslant 1$, the function $\mathrm{stp}_{k+1,q-1}$ is the restriction of $\mathrm{stp}_{k,q}$ to $\mathrm{domain}(\mathrm{stp}_{k+1,q-1})$.*

Proof. Level i of $\mathrm{domain}(\mathrm{stp}_{k+1,q-1})$ equals level $i+1$ of $\mathrm{domain}(\mathrm{stp}_{k,q})$: the ranks $\big((k+1)+i, (q-1)-i\big) = \big(k+(i+1), q-(i+1)\big)$ and the radius bounds coincide, the value caps $(k+1)+(q-1) = k+q$ coincide, and both functions use the same canonical set $S_r([\nu]^G)$. ∎

In the model checking application, the type trees of all vertices are maintained by a dynamic-programming algorithm, and the scatter type is computed from them. Recall that $|G| = |V(G)|$ and $\|G\| = |V(G)| + |R(G)|$.

Lemma 11.3. *There is an algorithm which, given a nonempty partial graph G, numbers $k \geqslant 1$ and $q \in \mathbb{N}$, its distance matrix D_G, and the family $(\mathrm{tree}_{k,q}(a) : a \in V(G))$, computes $\mathrm{stp}^G_{k,q}$ in time*

$$|G|^2 \cdot \mathrm{poly}_{k,q}\left(\max_{a \in V(G)} \|\mathrm{tree}_{k,q}(a)\|\right).$$

Proof. Put $M := \max_{a \in V(G)} \|\mathrm{tree}_{k,q}(a)\|$. For every $a \in V(G)$ and $i \in [q]$, applying Lemma 10.6 i times to $\mathrm{tree}_{k,q}(a)$ yields $\mathrm{tree}_{k+i,q-i}(a)$, of size $\mathrm{poly}_{k,q}(M)$. This takes $|G| \cdot \mathrm{poly}_{k,q}(M)$ time and gives $\mathrm{domain}(\mathrm{stp}_{k,q})$, as each radius is $O_{k,q}(1)$; in particular, $|\mathrm{domain}(\mathrm{stp}_{k,q})| = O_{k,q}(|G|)$.

For each $(r, \nu) \in \mathrm{domain}(\mathrm{stp}_{k,q})$, run the greedy process defining $S_r([\nu]^G)$, stopping after $k+q$ vertices were discovered. It scans $V(G)$ in order, testing membership in $[\nu]^G$ by comparing the trees computed above with ν. After adding a vertex v, scan the row $D_G(v, \cdot)$ and mark all vertices u with $D_G(v,u) \leqslant r$ as unavailable. At most $k+q$ rows are scanned, so one run takes $|G| \cdot \mathrm{poly}_{k,q}(M)$ time, and the $|\mathrm{domain}(\mathrm{stp}_{k,q})| = O_{k,q}(|G|)$ runs together take the stated time. ∎

11.2 The far tree

Definition 11.4. For a partial graph G, $k \in \mathbb{N}$, $q \geqslant 1$, and a nonempty tuple $\bar{a} \in V(G)^{\leqslant k}$, define

$$\mathrm{ftree}_{k,q}(\bar{a}) := \{\mathrm{tree}_{k+1,q-1}(b) \,:\, b \in V(G),\ \mathrm{dist}(\bar{a}, b) > \rho^-(k+1, q-1)\}.$$

The goal of this subsection is to compute $\mathrm{ftree}_{k,q}(\bar{a})$ from $\mathrm{tree}_{k,q}(\bar{a})$ and $\mathrm{stp}_{k,q}$ (Lemma 11.6). Our main tool is the following clustering lemma [DT26, Corollary 3.4 and Lemma 3.5].

Lemma 11.5 ([DT26])**.** *Let $r \in \mathbb{N}$, let $\ell \geqslant 1$, let G be a graph, and let $\bar{a} = (a_1, \ldots, a_\ell) \in V(G)^\ell$. There exist a radius R with $r \leqslant R \leqslant 9^{\ell-1} r$ and a partition $\mathcal{B}$ of $\{a_1, \ldots, a_\ell\}$ with the following property.*

For every $X \subseteq V(G)$, inclusion-wise maximal $4R$-scattered subset $S \subseteq X$, and $H \geqslant 4R + r$,

$$\begin{aligned} \exists x \in X \colon \mathrm{dist}^G(x, \bar{a}) > r \quad \Longleftrightarrow \quad & \exists x \in X \colon r < \mathrm{dist}^G(x, \bar{a}) \leqslant H, \\ & \textit{or} \quad \big|\{B \in \mathcal{B} \,:\, \exists x \in X \colon \mathrm{dist}^G(x, B) \leqslant r\}\big| < |S|. \end{aligned}$$

Moreover, the radius R and the partition $\mathcal{B}$ can be determined by tests of the form $\mathrm{dist}^G(a_i, a_j) \leqslant s$, where $i, j \in [\ell]$ and $s \leqslant 8 \cdot 9^{\ell-1} r$.

Lemma 11.6. *Fix $k \in \mathbb{N}$ and $q \geqslant 1$, let G be a partial graph, and let $\bar{a} \in V(G)^{\leqslant k}$ be nonempty. Then $\mathrm{ftree}_{k,q}(\bar{a})$ is computable from $\mathrm{tree}_{k,q}(\bar{a})$ and $\mathrm{stp}_{k,q}$ in time*

$$\|\mathrm{stp}_{k,q}\| \cdot \mathrm{poly}_{k,q}(\|\mathrm{tree}_{k,q}(\bar{a})\|).$$

Proof. Let $\tau := \mathrm{tree}_{k,q}(\bar{a})$, $\sigma := \mathrm{stp}_{k,q}$, $r := \rho^-(k+1, q-1)$, $H := \rho^+(k+1, q-1)$, and $\ell := |\bar{a}|$. The entries with $i = 1$ in the domain of σ let us compute the set

$$\mathcal{T} := \{\mathrm{tree}_{k+1,q-1}(b) \,:\, b \in V(G)\}.$$

Then

$$\mathrm{ftree}_{k,q}(\bar{a}) = \{\nu \in \mathcal{T} \,:\, \exists b \in [\nu]^G \colon \mathrm{dist}(\bar{a}, b) > r\}.$$

Rewriting the target. The tree τ only records vertices within distance H of $\bar{a}$, while a vertex $b \in [\nu]^G$ with $\mathrm{dist}(\bar{a}, b) > r$ may be farther away than H. We will now use Lemma 11.5 to replace the existence of such a vertex by two conditions that we can check. Apply the lemma with radius r to the Gaifman graph of G (with vertex set $V(G)$ and edge set $R(G)$) and to $\bar{a}$, and obtain a radius R and a partition $\mathcal{B}$ of $\{a_1, \dots, a_\ell\}$. Since $\ell \leqslant k$,

$$4R \leqslant 4 \cdot 9^{\ell-1} r \leqslant 9^{k+1} r, \qquad 4R + r \leqslant (4 \cdot 9^{\ell-1} + 1) r \leqslant 9^k r \leqslant H.$$

The first chain shows that $(4R, \nu)$ lies in the domain of σ for every $\nu \in \mathcal{T}$. The second is the hypothesis $H \geqslant 4R + r$ of Lemma 11.5, so the lemma applies to every $X \subseteq V(G)$ together with $S_{4R}(X)$. Define

$$\mathcal{N} := \{\mathrm{tree}_{k+1,q-1}(b) \,:\, b \in V(G),\ r < \mathrm{dist}(\bar{a}, b) \leqslant H\},$$
$$\mathcal{N}_B := \{\mathrm{tree}_{k+1,q-1}(b) \,:\, b \in V(G),\ \mathrm{dist}(b, B) \leqslant r\} \qquad \text{for } B \in \mathcal{B}.$$

For $\nu \in \mathcal{T}$, Lemma 11.5 applied to $X := [\nu]^G$ and $S := S_{4R}([\nu]^G)$ states that

$$\exists b \in [\nu]^G \colon \mathrm{dist}(\bar{a}, b) > r \quad \Longleftrightarrow \quad \nu \in \mathcal{N} \ \text{ or } \ \big|\{B \in \mathcal{B} \,:\, \nu \in \mathcal{N}_B\}\big| < |S_{4R}([\nu]^G)|.$$

As $\sigma(4R, \nu) = \min(|S_{4R}([\nu]^G)|, k+q)$ by Definition 11.1, and $\big|\{B \in \mathcal{B} \,:\, \nu \in \mathcal{N}_B\}\big| \leqslant |\mathcal{B}| \leqslant \ell \leqslant k < k+q$, we can equivalently replace the bound $|S_{4R}([\nu]^G)|$ above by $\sigma(4R, \nu)$. Therefore,

$$\mathrm{ftree}_{k,q}(\bar{a}) = \{\nu \in \mathcal{T} \,:\, \nu \in \mathcal{N} \ \text{ or } \ \big|\{B \in \mathcal{B} \,:\, \nu \in \mathcal{N}_B\}\big| < \sigma(4R, \nu)\}.$$

In this expression, $\mathcal{T}$ and the values $\sigma(4R, \nu)$ can be read from σ. It remains to compute R, $\mathcal{B}$, $\mathcal{N}$, and the sets $\mathcal{N}_B$ from τ.

Reading R and $\mathcal{B}$ off τ. By Lemma 11.5, R and $\mathcal{B}$ are determined by tests $\mathrm{dist}(a_i, a_j) \leqslant s$ with $i, j \in [\ell]$ and

$$s \leqslant 8 \cdot 9^{\ell-1} r \leqslant 9^k r \leqslant H \leqslant \rho^-(k, q).$$

These are binary distance atoms of radius at most $\rho^-(k, q)$ and hence determined by the label $\mathrm{atp}_{k,q}(\bar{a})$ of the root of τ. We represent every set $B \in \mathcal{B}$ by its set of positions $I_B := \{i \in [\ell] \,:\, a_i \in B\}$, which is determined by the equality atoms in the root atomic type of τ.

Reading $\mathcal{N}$ and $\mathcal{N}_B$ off τ. The trees attached to the root of τ are exactly the trees $\mathrm{tree}_{k+1,q-1}(\bar{a}b)$ with $b \in V(G)$ and $\mathrm{dist}(\bar{a}, b) \leqslant H$. Such a tree μ determines two things about b: by Lemma 10.5, $\mu \restriction (\ell+1) = \mathrm{tree}_{k+1,q-1}(b)$, and since the root atomic type $\mathrm{atp}_{k+1,q-1}(\bar{a}b)$ of μ contains the binary distance atoms of radius at most $\rho^-(k+1, q-1) = r$, we have $\mu \models \mathrm{dist}(x_j, x_{\ell+1}) \leqslant r$ if and only if $\mathrm{dist}(a_j, b) \leqslant r$, for every $j \in [\ell]$. Therefore,

$$\mathcal{N} = \{\mu \restriction (\ell+1) \,:\, (i, \mu) \in \mathcal{E}_{k,q}(\bar{a}),\ \mu \models \textstyle\bigwedge_{j \in [\ell]} \mathrm{dist}(x_j, x_{\ell+1}) > r\}.$$

Likewise, since $r \leqslant H$, the condition $\mathrm{dist}(b, B) \leqslant r$ implies $\mathrm{dist}(\bar{a}, b) \leqslant H$, and

$$\mathcal{N}_B = \{\mu \restriction (\ell+1) \,:\, (i, \mu) \in \mathcal{E}_{k,q}(\bar{a}),\ \mu \models \textstyle\bigvee_{j \in I_B} \mathrm{dist}(x_j, x_{\ell+1}) \leqslant r\}.$$

The right-hand sides depend only on τ. Substituting them into the expression for $\mathrm{ftree}_{k,q}(\bar{a})$ above gives a description in terms of τ and σ only.

Running time. The radius R and the sets I_B are read off the root label of τ. For all extensions (i, μ) of τ, the trees $\mu \restriction (\ell+1)$ are obtained by Lemma 10.5 and the atoms $\mathrm{dist}(x_j, x_{\ell+1}) \leqslant r$ are read off the root of μ, in time $\mathrm{poly}_{k,q}(\|\tau\|)$ in total. Collecting them yields $\mathcal{N}$ and the at most k sets $\mathcal{N}_B$ within the same bound. The expression for $\mathrm{ftree}_{k,q}(\bar{a})$ is then evaluated by one scan of σ: for every entry $(4R, \nu, \sigma(4R, \nu))$ with ν of rank $(k+1, q-1)$, membership of ν in $\mathcal{N}$ and in each $\mathcal{N}_B$ is tested by comparing ν with every element of these at most $k+1$ sets, in time $O(k \cdot \|\tau\| \cdot (\|\nu\| + \|\tau\|))$. Summed over all entries of σ this is $O(\|\sigma\| \cdot \mathrm{poly}_{k,q}(\|\tau\|))$, which gives the stated running time. ∎

12 Lifting type trees through merges and resolves

The model checking algorithm will process the construction sequence step by step and maintain the type trees of all vertices. We first show how to lift type trees through a merge, and then through a resolve.

Lemma 12.1 (merge lift)**.** *Fix $\ell \leqslant k$ and $q \in \mathbb{N}$. Let G be a partial graph and $\bar{a} \in V(G)^\ell$. Given $\mathrm{tree}^G_{k,q}(\bar{a})$ and two distinct colors P, Q, one can compute $\mathrm{tree}^H_{k,q}(\bar{a})$, where $H = \mathrm{merge}_{P,Q}(G)$, in time*

$$\mathrm{poly}_{k,q}(\|\mathrm{tree}^G_{k,q}(\bar{a})\|).$$

Proof. We proceed by induction on q.

Root labels. The partial graphs G and H have the same resolved pairs, so all distances agree, and only the colors P and Q change. Hence, for every vertex u and $s \in \mathbb{N}$,

$$\begin{aligned} H \models P(u) &\iff G \models P(u) \vee Q(u), \\ H \models Q(u) &\iff \mathit{false}, \\ H \models \mathrm{dist}(u,P) < s &\iff G \models (\mathrm{dist}(u,P) < s) \vee (\mathrm{dist}(u,Q) < s), \\ H \models \mathrm{dist}(u,Q) < s &\iff \mathit{false}, \end{aligned}$$

and every other atom has the same interpretation in G and H. This allows us to rewrite every atom φ into a quantifier-free formula φ' of the same distance rank, so that $H \models \varphi(\bar{a})$ if and only if $G \models \varphi'(\bar{a})$. Since whether or not $G \models \varphi'(\bar{a})$ is determined by $\mathrm{atp}^G_{k,q}(\bar{a})$, it follows that the new root label $\mathrm{atp}^H_{k,q}(\bar{a})$ can be computed from $\mathrm{atp}^G_{k,q}(\bar{a})$. Only atoms mentioning P or Q change, so the new root label is obtained by one pass over the old root label and is no larger. This settles the root label.

Extensions. Let $q \geqslant 1$. Since $\mathrm{dist}^H = \mathrm{dist}^G$, every guard admits the same vertices in G and H, so the inductive assumption gives $\mathcal{E}^H_{k,q}(\bar{a}) = \{(i, F(\mu)) : (i, \mu) \in \mathcal{E}^G_{k,q}(\bar{a})\}$, where $F(\mu)$ is obtained inductively. Together with the root label, this gives $\mathrm{tree}^H_{k,q}(\bar{a})$.

Running time. The transformation replaces atoms in the label of every node and keeps the guards, so it visits each node of the input tree once, and no tree grows. By Observation 10.1, restoring canonical form takes time polynomial in the size of the tree. ∎

A resolve can bring vertices far from a tuple within guard range. Their type trees are not recorded in the old type tree of the tuple, and we obtain them from the far tree instead.

Lemma 12.2 (resolve lift). *Let $k \geqslant 1$ and $q \in \mathbb{N}$. Given a partial graph G, its distance matrix D_G, colors P, Q, a sign $\pm$, and $(\mathrm{tree}^G_{k,q}(a) : a \in V(G))$, one can compute $H = \mathrm{resolve}^{\pm}_{P,Q}(G)$, its distance matrix D_H, and $(\mathrm{tree}^H_{k,q}(a) : a \in V(G))$ in time*

$$|G|^2 \cdot \mathrm{poly}_{k,q}(M+T),$$

where

$$M := \max_{a \in V(G)} \|\mathrm{tree}^G_{k,q}(a)\|,$$
$$T := \max_{a \in V(G)} \left|\left\{\mathrm{tree}^G_{k+i,q-i}(\bar{v}) \colon 0 \leqslant i \leqslant q,\ \ell \in [k+i],\ \bar{v} \in \left(\mathrm{Ball}^H_{\rho^+(k,q)}(a)\right)^\ell\right\}\right|.$$

Each output tree has size $\mathrm{poly}_{k,q}(T)$.

Proof. Assume the resolve is positive, as the negative case is symmetric. By induction on q, we construct an auxiliary map F on trees of arbitrary nonempty tuples, so that $\mathrm{tree}^H_{k,q}(\bar{a}) = F(\mathrm{tree}^G_{k,q}(\bar{a}))$ for all $\ell \leqslant k$ and $\bar{a} \in V(G)^\ell$. Fix $\bar{a} \in V(G)^\ell$ and put $\tau := \mathrm{tree}^G_{k,q}(\bar{a})$ and $\sigma := \mathrm{stp}^G_{k,q}$.

Root labels. We may assume that P^G and Q^G are nonempty, as otherwise $H = G$. We rewrite every atom φ in H to an equivalent formula φ' in G, so that φ' has the same distance rank as φ, and $H \models \varphi(\bar{a}) \iff G \models \varphi'(\bar{a})$. Below, u, v are vertices, Y is a color, and $s \in \mathbb{N}$.

First-order atoms. Equality, color, and non-edge atoms $N(x,y)$ have the same interpretation in G and H. By definition of the positive resolve,

$$H \models E(u,v) \quad\iff\quad G \models E(u,v) \vee \Big(\big((P(u) \wedge Q(v)) \vee (Q(u) \wedge P(v))\big) \wedge u \neq v \wedge \neg N(u,v)\Big).$$

Distances. We have $R(H) = R(G) \cup P^G Q^G$. A shortest path from u to v in H is a path in G, or runs in G from u to P^G or Q^G, then crosses to the other part or returns to the same part through the other one, and then runs in G to v: otherwise, the segment between its first and last vertex in $P^G \cup Q^G$ could be replaced by one pair from $P^G Q^G$, or by two through the other part if both ends lie in the same part. Hence

$$\mathrm{dist}^H(u,v) = \min \begin{Bmatrix} \mathrm{dist}^G(u,v), & \\ \mathrm{dist}^G(u,P) + 1 + \mathrm{dist}^G(Q,v), & \mathrm{dist}^G(u,Q) + 1 + \mathrm{dist}^G(P,v), \\ \mathrm{dist}^G(u,P) + 2 + \mathrm{dist}^G(P,v), & \mathrm{dist}^G(u,Q) + 2 + \mathrm{dist}^G(Q,v) \end{Bmatrix}. \tag{13}$$

Unary distance atoms. Taking the minimum of (13) over $v \in Y^G$ gives

$$\begin{aligned} H \models \mathrm{dist}(u,Y) < s \quad\iff\quad & G \models \mathrm{dist}(u,Y) < s \\ & \vee \mathrm{dist}(u,P) < (s - 1 - \mathrm{dist}^G(Q,Y)) \vee \mathrm{dist}(u,Q) < (s - 1 - \mathrm{dist}^G(P,Y)) \\ & \vee \mathrm{dist}(u,P) < (s - 2 - \mathrm{dist}^G(P,Y)) \vee \mathrm{dist}(u,Q) < (s - 2 - \mathrm{dist}^G(Q,Y)). \end{aligned}$$

The values $\mathrm{dist}^G(P,Y)$ and $\mathrm{dist}^G(Q,Y)$ do not depend on u, and are computed once. A unary-distance comparison whose upper bound is nonpositive, or from which an infinite distance is subtracted, is interpreted as *false*.

Binary distance atoms. Define

$$\begin{aligned}&\delta^s_{P,Q}(x,y) :=\\ &\bigvee_{s_1+1+s_2=s} \big(\mathrm{dist}(x,P) < s_1{+}1 \wedge \mathrm{dist}(y,Q) < s_2{+}1\big) \vee \big(\mathrm{dist}(x,Q) < s_1{+}1 \wedge \mathrm{dist}(y,P) < s_2{+}1\big)\\ \vee &\bigvee_{s_1+2+s_2=s} \big(\mathrm{dist}(x,P) < s_1{+}1 \wedge \mathrm{dist}(y,P) < s_2{+}1\big) \vee \big(\mathrm{dist}(x,Q) < s_1{+}1 \wedge \mathrm{dist}(y,Q) < s_2{+}1\big).\end{aligned}$$

By (13),

$$H \models \mathrm{dist}(u,v) \leqslant s \quad\iff\quad G \models (\mathrm{dist}(u,v) \leqslant s) \vee \delta^s_{P,Q}(u,v). \tag{14}$$

Root label. For an atom φ, let φ' be the corresponding formula as obtained above. Then $H \models \varphi(\bar{a})$ if and only if $G \models \varphi'(\bar{a})$ and φ' has the same distance rank as φ. Since whether or not $G \models \varphi'(\bar{a})$ is determined by $\mathrm{atp}^G_{k,q}(\bar{a})$, it follows that the new root label $\mathrm{atp}^H_{k,q}(\bar{a})$ can be computed from $\mathrm{atp}^G_{k,q}(\bar{a})$. Algorithmically, there are $O_{k,q}(1)$ atoms without colors, and each nonempty color contributes $O_{k,q}(1)$ further atoms.

Extensions. Let $q \geqslant 1$, and put $r := \rho^-(k+1, q-1)$ and $R := \rho^+(k+1, q-1)$. On trees of rank $(k+1, q-1)$, F denotes the map given by the inductive assumption, using the restriction of σ to $\mathrm{domain}(\mathrm{stp}_{k+1,q-1})$ from Lemma 11.2 and the same precomputed distances. By the inductive assumption,

$$\mathcal{E}^H_{k,q}(\bar{a}) = \{\big(i, F(\mathrm{tree}^G_{k+1,q-1}(\bar{a}b))\big) \,:\, i \in [\ell], \mathrm{dist}^H(a_i, b) \leqslant R\}. \tag{15}$$

For $i \in [\ell]$, put

$$\gamma_i(x_1, \ldots, x_{\ell+1}) := \delta^R_{P,Q}(x_i, x_{\ell+1}).$$

By (14),

$$H \models \mathrm{dist}(a_i, b) \leqslant R \quad\iff\quad G \models \mathrm{dist}(a_i, b) \leqslant R \;\vee\; \gamma_i(\bar{a}b). \tag{16}$$

The right-hand side of (16) is equivalent to the disjunction of three conditions:

1. $\mathrm{dist}^G(a_i, b) \leqslant R$,
2. $\mathrm{dist}^G(\bar{a}, b) \leqslant R$ and $G \models \gamma_i(\bar{a}b)$,
3. $\mathrm{dist}^G(\bar{a}, b) > r$ and $G \models \gamma_i(\bar{a}b)$.

This holds since $r \leqslant R$ by (11). Consequently,

$$\mathcal{E}^H_{k,q}(\bar{a}) = \mathcal{E}_{\mathrm{old}} \cup \mathcal{E}_{\mathrm{near}} \cup \mathcal{E}_{\mathrm{far}},$$

where

$$\begin{aligned}\mathcal{E}_{\mathrm{old}} &:= \{\big(i,\ F(\mathrm{tree}^G_{k+1,q-1}(\bar{a}b))\big) \,:\, i \in [\ell],\ b \in V(G),\ \mathrm{dist}^G(a_i, b) \leqslant R\},\\ \mathcal{E}_{\mathrm{near}} &:= \{\big(i,\ F(\mathrm{tree}^G_{k+1,q-1}(\bar{a}b))\big) \,:\, i \in [\ell],\ b \in V(G),\ \mathrm{dist}^G(\bar{a}, b) \leqslant R,\ G \models \gamma_i(\bar{a}b)\},\\ \mathcal{E}_{\mathrm{far}} &:= \{\big(i,\ F(\mathrm{tree}^G_{k+1,q-1}(\bar{a}b))\big) \,:\, i \in [\ell],\ b \in V(G),\ \mathrm{dist}^G(\bar{a}, b) > r,\ G \models \gamma_i(\bar{a}b)\}.\end{aligned}$$

The first set is obtained via

$$\mathcal{E}_{\text{old}} = \{(i, F(\mu)) \,:\, (i,\mu) \in \mathcal{E}^G_{k,q}(\bar{a})\}.$$

Every witness b for $\mathcal{E}_{\text{near}}$ is admitted by some guard j in G, so $\text{tree}^G_{k+1,q-1}(\bar{a}b)$ occurs as an extension (j,μ) of τ with $\mu \models \gamma_i$. Conversely, every such μ equals $\text{tree}^G_{k+1,q-1}(\bar{a}b)$ for some b with $\text{dist}^G(a_j, b) \leqslant R$ and $G \models \gamma_i(\bar{a}b)$. Hence, using the notation $\mu \models \gamma$ introduced after Lemma 10.2,

$$\mathcal{E}_{\text{near}} = \{(i, F(\mu)) \,:\, i \in [\ell],\ (j,\mu) \in \mathcal{E}^G_{k,q}(\bar{a}),\ \mu \models \gamma_i\}.$$

For $\mathcal{E}_{\text{far}}$, consider b with $\text{dist}^G(\bar{a}, b) > r$. Since $r = \rho^-(k+1, q-1)$ is exactly the separation required by Lemma 10.7 at rank $(k+1, q-1)$, that lemma gives

$$\text{tree}^G_{k+1,q-1}(\bar{a}b) = \text{tree}^G_{k+1,q-1}(\bar{a}) * \text{tree}^G_{k+1,q-1}(b),$$

where $\text{tree}^G_{k+1,q-1}(\bar{a})$ is obtained from $\tau = \text{tree}^G_{k,q}(\bar{a})$ by Lemma 10.6. As b ranges over all such vertices, $\text{tree}^G_{k+1,q-1}(b)$ ranges exactly over $\text{ftree}^G_{k,q}(\bar{a})$, computed from τ and σ by Lemma 11.6. Therefore

$$\mathcal{E}_{\text{far}} = \{\big(i,\ F(\text{tree}^G_{k+1,q-1}(\bar{a}) * \nu)\big) \,:\, i \in [\ell],\ \nu \in \text{ftree}^G_{k,q}(\bar{a}),\ \text{tree}^G_{k+1,q-1}(\bar{a}) * \nu \models \gamma_i\}.$$

So the new root label, $\mathcal{E}_{\text{old}}$, $\mathcal{E}_{\text{near}}$, and $\mathcal{E}_{\text{far}}$ are computed from τ together with σ and the values $\text{dist}^G(P, Y)$ and $\text{dist}^G(Q, Y)$, which do not depend on $\bar{a}$. This defines $F(\tau) = \text{tree}^H_{k,q}(\bar{a})$.

Running time. Computing H takes time $O(\|H\|) = O(|G|^2)$. Scanning D_G gives $\text{dist}^G(u, P)$ and $\text{dist}^G(u, Q)$ for all vertices u, and thus $\text{dist}^G(P, Y)$ and $\text{dist}^G(Q, Y)$ for all nonempty colors Y, in time $O(|G|^2)$. Then (13) gives D_H in time $O(|G|^2)$, and Lemma 11.3 computes $\sigma := \text{stp}^G_{k,q}$ from D_G in time $|G|^2 \cdot \text{poly}_{k,q}(M)$. Every tree $\text{tree}^G_{k+i,q-i}(b)$ is obtained from $\text{tree}^G_{k,q}(b)$ by at most q applications of Lemma 10.6, so it has size $\text{poly}_{k,q}(M)$, and $\|\sigma\| \leqslant |G| \cdot \text{poly}_{k,q}(M)$, since the domain of σ has $O_{k,q}(1)$ entries per vertex.

Number of calls. Fix $a \in V(G)$ and consider the recursion computing $F(\text{tree}^G_{k,q}(a))$. We collect the distinct old extension trees admitted by at least one guard in H before recursing, retaining the list of guards for each tree. The calls at depth j are on trees $\text{tree}^G_{k+j,q-j}(\bar{v})$ for tuples $\bar{v}$ reached by the guards in H. By Lemma 10.4 applied to H, every such $\bar{v}$ is a tuple from $\text{Ball}^H_{\rho^+(k,q)}(a)$. So by definition of T, a call makes at most T recursive calls, and as the depth is q, there are $\text{poly}_{k,q}(T)$ calls for a.

Cost per call and output size. A recursive input is a child of the previous input, or is obtained by rank shift and far composition with a tree of arity one recorded in σ. These trees have size $\text{poly}_{k,q}(M)$, so by Lemmas 10.6 and 10.7 and the depth bound q, every recursive input and every candidate extension tree has size $\text{poly}_{k,q}(M)$. For a call with positive remaining depth, restricting σ and applying Lemma 11.6 takes time $O(\|\sigma\| \cdot \text{poly}_{k,q}(M)) = O(|G| \cdot \text{poly}_{k,q}(M))$. The far tree has at most $|G|$ members, so composing them and testing the formulas γ_i (which are quantifier-free formulas of distance rank $(k+1, q-1)$, so are determined by the root label of a tree of rank $(k+1, q-1)$) takes the same time. We scan the candidates, test their guards, and compare each admitted candidate with the list of at most T distinct admitted trees collected so far. Including the candidates from the old extensions, this takes $O(|G| \cdot \text{poly}_{k,q}(M+T))$ time.

The colors recorded in each new root label meet $\mathrm{Ball}^{H}_{\rho^+(k,q)}(a)$ by Lemma 10.4 applied to H, and as vertices of distinct colors have distinct old trees, there are at most T of them. So each new root label has size $O_{k,q}(T)$ and is computed in time $O(|G| \cdot \|\tau\|)$. Together with the branching and depth bounds, this gives output size $\mathrm{poly}_{k,q}(T)$; restoring canonical form has the same polynomial bound by Observation 10.1.

Altogether the recursion for a takes time $O(|G| \cdot \mathrm{poly}_{k,q}(M+T))$, and for all vertices $O(|G|^2 \cdot \mathrm{poly}_{k,q}(M+T))$. Adding the construction of H and D_H, the distance preprocessing, and the scatter type gives the claimed bound. ∎

13 The model checking algorithm

We now give the model checking algorithm, proving Theorem 4.3. Recall the width notion used in the statement.

Definition 4.1. For $r, q \in \mathbb{N}$, the *radius-r depth-q width* of a construction sequence $G_1, \dots, G_\lambda$ is defined as:

$$\max_{t \in [\lambda-1]} \max_{v \in V(G)} |\operatorname{Types}^{G_t}_{q,q}(\mathrm{Ball}^{G_{t+1}}_r(v))|.$$

The width bounds numbers of ordinary types, and the following lemma converts such bounds into bounds on type trees, as required by Lemma 12.2.

Lemma 13.1 (type tree sizes)**.** *Let $k \geqslant 1$ and $q \in \mathbb{N}$, and put $r := \rho^+(k,q)$ and $d := r + k + q$. The following holds for every partial graph G.*

(i) $\big|\{\mathrm{tree}^G_{k+i,q-i}(\bar v) : \bar v \in X^\ell\}\big| \leqslant |\operatorname{Types}^G_{d,d}(X)|$ *for all $i \in \{0,\dots,q\}$, $\ell \in [k+i]$, and $X \subseteq V(G)$.*

(ii) $\|\mathrm{tree}^G_{k,q}(a)\| \leqslant \mathrm{poly}_{k,q}(|\operatorname{Types}^G_{d,d}(\mathrm{Ball}^G_r(a))|)$ *for every $a \in V(G)$.*

Proof. We first show that ordinary types determine type trees: for all $k', q' \in \mathbb{N}$, every partial graph G, and all nonempty tuples $\bar a, \bar b \in V(G)^{\leqslant k'}$ of the same length,

$$\mathrm{tp}^G_{\rho^+(k',q')+q'}(\bar a) = \mathrm{tp}^G_{\rho^+(k',q')+q'}(\bar b) \quad\Longrightarrow\quad \mathrm{tree}^G_{k',q'}(\bar a) = \mathrm{tree}^G_{k',q'}(\bar b). \tag{17}$$

By expressing distance atoms by paths along $E \vee N$, every dist-FO formula of distance rank (k',q') is equivalent to a first-order formula of quantifier rank at most $\rho^+(k',q') + q'$. Thus, the left-hand side of (17) implies that $\bar a$ and $\bar b$ satisfy the same dist-FO formulas of distance rank (k',q'). By Corollary 10.3, this yields the right-hand side.

For (i), extend every tuple $\bar v \in X^\ell$ to a tuple $\bar v' \in X^d$ by repeating its last coordinate. Then $\mathrm{tp}_d(\bar v')$ determines $\mathrm{tp}_d(\bar v)$, which determines $\mathrm{tree}_{k+i,q-i}(\bar v)$ by (17), as $\rho^+(k+i,q-i) + q - i \leqslant \rho^+(k,q) + q \leqslant d$ by (12). So distinct type trees yield distinct elements of $\operatorname{Types}^G_{d,d}(X)$.

For (ii), put $X := \mathrm{Ball}_r(a)$ and $K := |\operatorname{Types}^G_{d,d}(X)|$. As every vertex has exactly one color, vertices of distinct colors have distinct type trees, so at most K colors meet X by (i). By Lemma 10.4, every node of $\mathrm{tree}_{k,q}(a)$ is $\mathrm{tree}_{k+j,q-j}(\bar v)$ for a tuple $\bar v$ from X, and every color in a true atom of its root meets X. So every root has size $O_{k,q}(K)$, and by (i), every node has at most K children per guard, and thus at most $(k+q)K$ children. As the depth is q, this gives $\|\mathrm{tree}_{k,q}(a)\| \leqslant \mathrm{poly}_{k,q}(K)$. ∎

Theorem 4.3. *There is an algorithm which, given a sentence φ and an n-vertex graph G together with its construction sequence of length λ, determines whether $G \models \varphi$ in time $\lambda \cdot n^2 \cdot \mathrm{poly}_\varphi(K)$, where K is the radius-r depth-d width of the construction sequence, for some r and d which are computable from φ.*

Proof. Let q be the quantifier rank of φ, and let r, d be as in Lemma 13.1 for $k = 1$ and this q. We prove the statement with radius r and depth d. If $n \leqslant 1$, evaluate φ directly, so assume $n \geqslant 2$.

Let $G_1, \dots, G_\lambda$ be the given construction sequence of G, and let K be its radius-r depth-d width. We iteratively compute $(\mathrm{tree}_{1,q}^{G_t}(a) : a \in V(G))$ for $t = 1, \dots, \lambda$. To save time we also maintain the distance matrix D_{G_t}. Initially, for $t = 1$, there are no resolved pairs and each nonempty part is a singleton. Hence, every guard admits only vertices of the guarded tuple itself, so each $\mathrm{tree}_{1,q}^{G_1}(a)$ is determined by the part containing a, and can therefore be computed in time $O_q(1)$. Initialize also D_{G_1} in $O(n^2)$ time by putting 0 on its diagonal and ∞ elsewhere.

Now let $t < \lambda$, and assume the trees in G_t have been computed. All these trees have size $\mathrm{poly}_q(K)$: for every $t' \in [\lambda - 1]$ and $a \in V(G)$, Lemma 13.1(ii) bounds $\|\mathrm{tree}_{1,q}^{G_{t'}}(a)\|$ by a polynomial in $|\operatorname{Types}_{d,d}^{G_{t'}}(\mathrm{Ball}_r^{G_{t'}}(a))|$. Using $R(G_{t'}) \subseteq R(G_{t'+1})$ and therefore $\mathrm{Ball}_r^{G_{t'}}(a) \subseteq \mathrm{Ball}_r^{G_{t'+1}}(a)$, this is at most $|\operatorname{Types}_{d,d}^{G_{t'}}(\mathrm{Ball}_r^{G_{t'+1}}(a))| \leqslant K$ by the definition of the width.

For a merge, construct G_{t+1} in time $O(n)$, keep $D_{G_{t+1}} = D_{G_t}$ unchanged, and apply Lemma 12.1 to each tree of arity one, taking $O(n \cdot \mathrm{poly}_q(K))$ time in total. For a resolve, apply Lemma 12.2 to G_t, the supplied matrix D_{G_t}, and all its trees of arity one. The lemma returns both $D_{G_{t+1}}$ and the new trees. Its parameter M is $\mathrm{poly}_q(K)$ by the bound above. Its parameter T is

$$T := \max_{a \in V(G)} \left| \left\{ \mathrm{tree}_{1+i,q-i}(\bar{v}) \,:\, \bar{v} \in \left(\mathrm{Ball}_{\rho^+(1,q)}^{G_{t+1}}(a)\right)^\ell,\ 0 \leqslant i \leqslant q,\ \ell \in [i+1] \right\} \right|.$$

The tuples $\bar{v}$ above lie in $X := \mathrm{Ball}_r^{G_{t+1}}(a)$ since $\rho^+(1,q) \leqslant r$. By Lemma 13.1(i), there are at most $|\operatorname{Types}_{d,d}^{G_t}(X)| \leqslant K$ trees for each i and arity ℓ. There are at most $(q+1)^2$ choices of i and ℓ, so $T \leqslant (q+1)^2 K$.

As each step takes $O(n^2 \cdot \mathrm{poly}_q(K))$ time, the total time is $O(\lambda n^2 \cdot \mathrm{poly}_\varphi(K))$. After $\lambda - 1$ steps we have all the type trees in G_λ. These final trees also have size $\mathrm{poly}_q(K)$: if the last step is a merge, no tree grows; if it is a resolve, use the output-size bound $\mathrm{poly}_q(T)$ in Lemma 12.2 and $T \leqslant (q+1)^2 K$.

As all pairs of G_λ are resolved, we can turn each quantifier of φ into a local one without changing its truth value. Construct $\varphi^{\mathrm{loc}}(x)$ from φ by introducing a fresh variable x, and replacing every subformula $\exists y\, \psi$ of φ that lies in the scope of quantifiers $\exists z_1, \dots, \exists z_j$ by

$$\exists y \left(\mathrm{dist}(x,y) \leqslant \rho^+(j+2, q-j-1) \wedge \psi \right),$$

where ψ is transformed in the same way. This yields a dist-FO formula $\varphi^{\mathrm{loc}}(x)$ such that, for every $a \in V(G)$,

$$G_\lambda \models \varphi \quad \Longleftrightarrow \quad G_\lambda \models \varphi^{\mathrm{loc}}(a).$$

Note that φ^{loc} has distance rank $(1, q)$.

Now choose any $a \in V(G)$, and use Lemma 10.2 for $k = 1$ to compute whether $G_\lambda \models \varphi^{\mathrm{loc}}(a)$, and hence whether $G_\lambda \models \varphi$, from $\mathrm{tree}_{1,q}^{G_\lambda}(a)$ in time $\mathrm{poly}\big(|\varphi^{\mathrm{loc}}| + \|\mathrm{tree}_{1,q}^{G_\lambda}(a)\|\big) \leqslant \mathrm{poly}_\varphi(K)$. As $E(G_\lambda) = E$ and φ is a sentence over graphs, $G_\lambda \models \varphi$ if and only if $G \models \varphi$. This fits the preceding time bound and completes the proof. ∎

AI methodology

We developed the approach to the main theorem in summer 2025, without AI assistance. By summer 2026, we had a draft of the present paper, and the companion paper [Dre+26] was nearly complete. The draft already had essentially the current proof structure. It contained a complete proof of the result of Part I, and definitions and lemma statements for the other parts, though some of their proofs were still sketches. Believing that these sketches could be completed, we used GPT-6 to help check the argument.

In Part II, we had proposed a proof of a stronger version of Lemma 8.1, in which the number of unary predicates is bounded by a constant independent of the length of the input merge sequence. This would have let us apply the results of [Dre+26] without controlling the sequence length. The model identified a gap in this proof. We expected the gap to be easy to repair with the methods already in the paper, but the model did not find a repair.

We therefore asked a narrower question: are there construction sequences of width $n^{o(1)}$ that are locally definable using a constant number of unary predicates? The model observed that a slightly weaker width bound of n^{ε}, for any fixed $\varepsilon > 0$, can be achieved by tuning the parameter q in Theorem 4.3 in [Dre+26], so that it produces chains of length $O(r/\varepsilon)$. Since the number of unary predicates depends only on the length of the input sequence, this gives a bound independent of the size of the graph. We immediately recognized that this was enough to close the gap within our original proof strategy.

The text of Parts I and II was written by the authors. We used LLMs to help write some proofs in Part III that are mechanical adaptations of proofs from [DT26]. GPT-6 also proposed a shorter proof of Lemma 5.6 in Part I, which we verified and adopted. Finally, we used LLMs to streamline the presentation. The models used were GPT-6, Claude Opus 5.5 and Claude Fable 5.1.

The authors verified the correctness and originality of all content, including references.

References

[AA14] Hans Adler and Isolde Adler. "Interpreting nowhere dense graph classes as a classical notion of model theory". In: *Eur. J. Comb.* 36 (2014), pp. 322–330. DOI: 10.1016/j.ejc.2013.06.048.

[ALS16] *Algorithms, Logic and Structure Workshop in Warwick, Open Problem Session*. [Online; accessed 23-Jan-2023]. 2016. URL: https://warwick.ac.uk/fac/sci/maths/people/staff/daniel_kral/alglogstr/openproblems.pdf.

[Asc+16] Matthias Aschenbrenner, Alf Dolich, Deirdre Haskell, Dugald Macpherson, and Sergei Starchenko. "Vapnik–Chervonenkis density in some theories without the independence property, I". In: *Transactions of the American Mathematical Society* 368.8 (2016), pp. 5889–5949. DOI: 10.1090/tran/6659.

[Ber+26] Steffen van Bergerem, Martin Grohe, Charlotte Lenz, and Nicole Schweikardt. "A Rank-Preserving Gaifman Normal Form for First-Order Logic on Weighted Structures". In: *arXiv preprint arXiv:2609.37337* (2026).

[Bon+20] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. "Twin-width I: tractable FO model checking". In: *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*. IEEE, 2020, pp. 601–612. DOI: 10.1109/FOCS46700.2020.00062.

[Bon+22a] Édouard Bonnet, Dibyayan Chakraborty, Eun Jung Kim, Noleen Köhler, Raul Lopes, and Stéphan Thomassé. "Twin-Width VIII: Delineation and Win-Wins". In: *17th International Symposium on Parameterized and Exact Computation, IPEC 2022, September 7-9, 2022, Potsdam, Germany*. Vol. 249. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 9:1–9:18. DOI: 10.4230/LIPIcs.IPEC.2022.9.

[Bon+22b] Édouard Bonnet, Jan Dreier, Jakub Gajarský, Stephan Kreutzer, Nikolas Mählmann, Pierre Simon, and Szymon Toruńczyk. "Model Checking on Interpretations of Classes of Bounded Local Cliquewidth". In: *LICS '22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science, Haifa, Israel, August 2 - 5, 2022*. ACM, 2022, 54:1–54:13. DOI: 10.1145/3531130.3533367.

[Bon+22c] Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Toruńczyk. "Twin-Width IV: Ordered Graphs and Matrices". In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2022*. ACM, 2022, pp. 924–937. DOI: 10.1145/3519935.3520037.

[Bon+25] Édouard Bonnet, Samuel Braunfeld, Ioannis Eleftheriadis, Colin Geniet, Nikolas Mählmann, Michal Pilipczuk, Wojciech Przybyszewski, and Szymon Torunczyk. "Separability Properties of Monadically Dependent Graph Classes". In: *52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, Aarhus, Denmark, July 8-11, 2025*. Vol. 334. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025, 147:1–147:19. DOI: 10.4230/LIPICS.ICALP.2025.147.

[Bra+22] Samuel Braunfeld, Jaroslav Nešetřil, Patrice Ossona de Mendez, and Sebastian Siebertz. *Decomposition horizons: from graph sparsity to model-theoretic dividing lines*. 2022. eprint: 2209.11229.

[BS85] J.T. Baldwin and S. Shelah. "Second-order quantifiers and the complexity of theories". In: *Notre Dame Journal of Formal Logic* 26.3 (1985), pp. 229–303.

[CMR00] Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. "Linear time solvable optimization problems on graphs of bounded clique-width". In: *Theory of Computing Systems* 33.2 (2000), pp. 125–150.

[Cou90] Bruno Courcelle. "The monadic second-order logic of graphs. I. Recognizable sets of finite graphs". In: *Information and Computation* 85.1 (1990), pp. 12–75.

[DGK07] Anuj Dawar, Martin Grohe, and Stephan Kreutzer. "Locally excluding a minor". In: *22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007)*. IEEE. 2007, pp. 270–279.

[DGL11] Alfred Dolich, John Goodrick, and David Lippel. "Dp-Minimality: Basic Facts and Examples". In: *Notre Dame Journal of Formal Logic* 52.3 (2011), pp. 267–288. DOI: 10.1215/00294527-1435456.

[DGP26] Jan Dreier, Jakub Gajarskỳ, and Michał Pilipczuk. "Efficient reversal of transductions of sparse graph classes". In: *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*. 2026, pp. 1180–1191.

[Din+96] Guoli Ding, Bogdan Oporowski, James G. Oxley, and Dirk Vertigan. "Unavoidable Minors of Large 3-Connected Binary Matroids". In: *J. Comb. Theory, Ser. B* 66.2 (1996), pp. 334–360. DOI: 10.1006/jctb.1996.0026.

[DK09a] Anuj Dawar and Stephan Kreutzer. "Domination Problems in Nowhere-Dense Classes". In: *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2009)*. Vol. 4. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2009, pp. 157–168. DOI: 10.4230/LIPIcs.FSTTCS.2009.2315.

[DK09b] Anuj Dawar and Stephan Kreutzer. "Parameterized complexity of first-order logic". In: *Electronic Colloquium on Computational Complexity, TR09-131*. 2009, p. 39.

[DKT13] Zdenek Dvorák, Daniel Král, and Robin Thomas. "Testing first-order properties for subclasses of sparse graphs". In: *J. ACM* 60.5 (2013), 36:1–36:24. DOI: 10.1145/2499483.

[DMS23] Jan Dreier, Nikolas Mählmann, and Sebastian Siebertz. "First-Order Model Checking on Structurally Sparse Graph Classes". In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023*. ACM, 2023, pp. 567–580. DOI: 10.1145/3564246.3585186.

[DMT24] Jan Dreier, Nikolas Mählmann, and Szymon Toruńczyk. "Flip-Breakability: A Combinatorial Dichotomy for Monadically Dependent Graph Classes". In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024*. ACM, 2024, pp. 1550–1560. DOI: 10.1145/3618260.3649739. Full version: arXiv:2403.15201.

[Dra+26a] Karolina Drabik, Maël Dumas, Colin Geniet, Jakub Nowakowski, Michał Pilipczuk, and Szymon Toruńczyk. "Variants of Merge-Width and Applications". In: *arXiv preprint arXiv:2602.23867* (2026).

[Dra+26b] Karolina Drabik, Maël Dumas, Nikolas Mählmann, Wojciech Przybyszewski, and Szymon Toruńczyk. "Flips and Merge-Width in Sparse Graphs". In: *arXiv preprint arXiv:2602.12926* (2026).

[Dre+23] Jan Dreier, Nikolas Mählmann, Sebastian Siebertz, and Szymon Toruńczyk. "Indiscernibles and Flatness in Monadically Stable and Monadically NIP Classes". In: *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*. Vol. 261. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 125:1–125:18. DOI: 10.4230/LIPIcs.ICALP.2023.125.

[Dre+24] Jan Dreier, Ioannis Eleftheriadis, Nikolas Mählmann, Rose McCarty, Michał Pilipczuk, and Szymon Toruńczyk. "First-Order Model Checking on Monadically Stable Graph Classes". In: *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024*. IEEE, 2024. DOI: 10.1109/FOCS61266.2024.00012.

[Dre+26] Jan Dreier, Nikolas Mählmann, Michał Pilipczuk, and Szymon Toruńczyk. *Monadically dependent graph classes have almost bounded merge-width*. To appear in arXiv. 2026.

[Dre+27] Jan Dreier, Nikolas Mählmann, Rose McCarty, Michał Pilipczuk, and Szymon Toruńczyk. "Neighborhood Complexity and Radius-1 Merge-Width in Monadically Dependent Graph Classes". In: *Proceedings of the 2027 ACM-SIAM Symposium on Discrete Algorithms (SODA 2027)*. To appear. arXiv:2607.10941. 2027.

[DT25a] Jan Dreier and Szymon Toruńczyk. "Merge-Width and First-Order Model Checking". In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. STOC '25. Association for Computing Machinery, 2025, pp. 1944–1955. DOI: 10.1145/3717823.3718259. Full version available at arXiv:2502.18065V2.

[DT25b] Jan Dreier and Szymon Toruńczyk. "Merge-width and First-order Model-checking". In: *arXiv preprint* (2025). arXiv: 2502.18065V2.

[DT26] Jan Dreier and Szymon Toruńczyk. *A Rank-Preserving Locality Theorem*. 2026. arXiv: 2606.23180 [cs.LO].

[EK17] Kord Eickmeyer and Ken-ichi Kawarabayashi. "FO model checking on map graphs". In: *International Symposium on Fundamentals of Computation Theory*. Springer. 2017, pp. 204–216.

[FG01a] Jörg Flum and Martin Grohe. "Fixed-parameter tractability, definability, and model-checking". In: *SIAM Journal on Computing* 31.1 (2001), pp. 113–145.

[FG01b] Markus Frick and Martin Grohe. "Deciding first-order properties of locally tree-decomposable structures". In: *Journal of the ACM* 48.6 (2001), pp. 1184–1206.

[Gai82] Haim Gaifman. "On Local and Non-Local Properties". In: *Proceedings of the Herbrand Symposium*. Vol. 107. Studies in Logic and the Foundations of Mathematics. Elsevier, 1982, pp. 105–135. DOI: 10.1016/S0049-237X(08)71879-2.

[Gaj+20a] Jakub Gajarský, Petr Hliněný, Jan Obdržálek, Daniel Lokshtanov, and M. S. Ramanujan. "A New Perspective on FO Model Checking of Dense Graph Classes". In: *ACM Trans. Comput. Logic* 21.4 (2020). ISSN: 1529-3785. DOI: 10.1145/3383206.

[Gaj+20b] Jakub Gajarský, Stephan Kreutzer, Jaroslav Nešetřil, Patrice Ossona de Mendez, Michał Pilipczuk, Sebastian Siebertz, and Szymon Toruńczyk. "First-Order Interpretations of Bounded Expansion Classes". In: *ACM Trans. Comput. Log.* 21.4 (2020), 29:1–29:41. DOI: 10.1145/3382093.

[Gaj+22] Jakub Gajarský, Michał Pilipczuk, Wojciech Przybyszewski, and Szymon Toruńczyk. "Twin-Width and Types". In: *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022*. Vol. 229. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 123:1–123:21. DOI: 10.4230/LIPIcs.ICALP.2022.123.

[Gaj+23] Jakub Gajarský, Nikolas Mählmann, Rose McCarty, Pierre Ohlmann, Michał Pilipczuk, Wojciech Przybyszewski, Sebastian Siebertz, Marek Sokołowski, and Szymon Toruńczyk. "Flipper Games for Monadically Stable Graph Classes". In: *50th International Colloquium on Automata, Languages, and Programming, ICALP 2023*. Vol. 261. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 128:1–128:16. DOI: 10.4230/LIPIcs.ICALP.2023.128. Full version available at arXiv:2301.13735.

[Gan+15] Robert Ganian, Petr Hlineny, Daniel Král, Jan Obdrzálek, Jarett Schwartz, and Jakub Teska. "FO model checking of interval graphs". In: *Logical Methods in Computer Science* 11 (2015).

[GK11] Martin Grohe and Stephan Kreutzer. "Methods for Algorithmic Meta Theorems". In: *Model Theoretic Methods in Finite Combinatorics*. Vol. 558. Contemporary Mathematics. American Mathematical Society, 2011, pp. 181–205. DOI: 10.1090/conm/558/11051.

[GKS17] Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. "Deciding First-Order Properties of Nowhere Dense Graphs". In: *J. ACM* 64.3 (2017), 17:1–17:32. DOI: 10.1145/3051095.

[GM14] Sylvain Guillemot and Dániel Marx. "Finding small patterns in permutations in linear time". In: *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*. SIAM. 2014, pp. 82–101.

[Gra+04] Sylvain Gravier, Frédéric Maffray, Jérôme Renault, and Nicolas Trotignon. "Ramsey-type results on singletons, co-singletons and monotone sequences in large collections of sets". In: *European Journal of Combinatorics* 25.5 (2004), pp. 719–734. ISSN: 0195-6698. DOI: https://doi.org/10.1016/j.ejc.2003.10.004.

[Gro07] Martin Grohe. "Logic, Graphs, and Algorithms". In: *Electron. Colloquium Comput. Complex.* TR07-091 (2007).

[Mat04] Jiří Matoušek. "Bounded VC-Dimension Implies a Fractional Helly Theorem". In: *Discrete Comput. Geom.* 31.2 (2004), pp. 251–255.

[Neš+21] Jaroslav Nešetřil, Patrice Ossona de Mendez, Michał Pilipczuk, Roman Rabinovich, and Sebastian Siebertz. "Rankwidth Meets Stability". In: *Proceedings of the Thirty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '21. SIAM, 2021, pp. 2014–2033.

[NM12] Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity - Graphs, Structures, and Algorithms*. Vol. 28. Algorithms and combinatorics. Springer, 2012. DOI: 10.1007/978-3-642-27875-4.

[NO08] Jaroslav Nešetřil and Patrice Ossona de Mendez. "Grad and classes with bounded expansion I. Decompositions". In: *European Journal of Combinatorics* 29.3 (2008), pp. 760–776. ISSN: 0195-6698. DOI: 10.1016/j.ejc.2006.07.013.

[NO11] Jaroslav Nešetřil and Patrice Ossona de Mendez. "On nowhere dense graphs". In: *Eur. J. Comb.* 32.4 (2011), pp. 600–617. DOI: 10.1016/j.ejc.2011.01.006.

[OS06] Sang-il Oum and Paul Seymour. "Approximating clique-width and branch-width". In: *Journal of Combinatorial Theory, Series B* 96.4 (2006), pp. 514–528. ISSN: 0095-8956. DOI: 10.1016/j.jctb.2005.10.006.

[Prz26] Wojciech Przybyszewski. "Flipping, Forking, and Decompositions in Monadically Stable and Monadically Dependent Classes". PhD thesis. University of Warsaw, June 2026.

[PST18] Michał Pilipczuk, Sebastian Siebertz, and Szymon Toruńczyk. "On the number of types in sparse graphs". In: *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*. ACM, 2018, pp. 799–808. DOI: 10.1145/3209108.3209178. arXiv: 1705.09336 [cs.DM].

[PZ78] Klaus-Peter Podewski and Martin Ziegler. "Stable graphs". In: *Fundamenta Mathematicae* 100.2 (1978), pp. 101–107.

[RS95] Neil Robertson and P. D. Seymour. "Graph minors. XIII: the disjoint paths problem". In: *J. Comb. Theory Ser. B* 63.1 (Jan. 1995), pp. 65–110. ISSN: 0095-8956. DOI: 10.1006/jctb.1995.1006.

[Sau72] Norbert Sauer. "On the density of families of sets". In: *Journal of Combinatorial Theory, Series A* 13.1 (1972), pp. 145–147. ISSN: 0097-3165. DOI: 10.1016/0097-3165(72)90019-2.

[Sco79] Dana Scott. "A Note on Distributive Normal Forms". In: *Essays in Honour of Jaakko Hintikka: On the Occasion of His Fiftieth Birthday on January 12, 1979*. Springer Netherlands, 1979, pp. 75–90. DOI: 10.1007/978-94-009-9860-5_5.

[See96] Detlef Seese. "Linear time computable problems and first-order descriptions". In: *Mathematical Structures in Computer Science* 6.6 (1996), pp. 505–526. DOI: 10.1017/S0960129500070079.

[She69] Saharon Shelah. "Stable Theories". In: *Israel Journal of Mathematics* 7 (1969), pp. 187–202. DOI: 10.1007/BF02787611.

[She72] Saharon Shelah. "A combinatorial problem; stability and order for models and theories in infinitary languages." In: *Pacific Journal of Mathematics* 41.1 (1972), pp. 247–261. DOI: pjm/1102968432.

[Sim15] Pierre Simon. *A Guide to NIP Theories*. Lecture Notes in Logic. Cambridge University Press, 2015. DOI: 10.1017/CBO9781107415133.

[ST21] Pierre Simon and Szymon Toruńczyk. "Ordered graphs of bounded twin-width". In: *arXiV* abs/2102.06881 (2021). arXiv: 2102.06881.

[Tor23] Szymon Toruńczyk. "Flip-width: Cops and Robber on dense graphs". In: *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023*. IEEE, 2023, pp. 663–700. DOI: 10.1109/FOCS57990.2023.00045. Full version available at arXiv:2302.00352.

Appendix

A Forbidden patterns

In this appendix we give a self-contained proof of Lemma 5.11, which is a corollary of [ST21, Lemma C.1]. The proofs in this appendix apply to arbitrary classes of relational structures, not just classes of finite graphs. Recall the notion of (φ, ψ)-patterns.

Definition 5.10. Fix first-order formulas $\varphi(\bar{x}, z)$ and $\psi(\bar{y}, z)$. A (φ, ψ)-*pattern* of order $s \times t$ in a graph G consists of tuples $\bar{a}_1, \dots, \bar{a}_s \in V(G)^{|\bar{x}|}$, $\bar{b}_1, \dots, \bar{b}_t \in V(G)^{|\bar{y}|}$, and vertices $(c_{i,j} : i \in [s], j \in [t])$ such that the following holds for all $i, i' \in [s]$ and $j, j' \in [t]$:

1. $G \models \varphi(\bar{a}_i, c_{i,j})$ and $G \not\models \varphi(\bar{a}_{i'}, c_{i,j})$ if $i' > i$,
2. $G \models \psi(\bar{b}_j, c_{i,j})$ and $G \not\models \psi(\bar{b}_{j'}, c_{i,j})$ if $j' > j$.

We call a (φ, ψ)-pattern *clean* if the above conditions hold with $\neq$ in place of $>$ in both conditions.

We define dp-minimal classes [DGL11] as follows.

Definition A.1. A class $\mathscr{C}$ of structures is *dp-minimal* if for every first-order formulas $\varphi(\bar{x}, z), \psi(\bar{y}, z)$ in the signature of $\mathscr{C}$ there is a constant $c \in \mathbb{N}$ such that no structure $G \in \mathscr{C}$ contains a clean (φ, ψ)-pattern of order $c \times c$.

We will prove the following two statements.

Lemma A.2. *Every monadically dependent class $\mathscr{C}$ of structures is dp-minimal.*

Lemma A.3. *Let $\mathscr{C}$ be a dp-minimal class of structures, and let $\varphi(\bar{x}, z)$ and $\psi(\bar{y}, z)$ be first-order formulas in the signature of $\mathscr{C}$. Then there is a constant $t \in \mathbb{N}$ such that no structure $G \in \mathscr{C}$ contains a (φ, ψ)-pattern of order $t \times t$.*

Clearly, Lemma A.2 and Lemma A.3 imply Lemma 5.11, repeated below.

Lemma 5.11. *For every monadically dependent graph class $\mathscr{C}$ and formulas $\varphi(\bar{x}, z), \psi(\bar{y}, z)$, there is some $t \in \mathbb{N}$ such that no graph $G \in \mathscr{C}$ contains a (φ, ψ)-pattern of order $t \times t$.*

We first prove Lemma A.2. This follows easily from the following characterization of monadic dependence in terms of the independence property. A formula $\varphi(\bar{x}, \bar{y})$ has the *independence property* on a class $\mathscr{C}$ of structures if for every $n \in \mathbb{N}$ there is some $G \in \mathscr{C}$ and tuples $\bar{a}_1, \dots, \bar{a}_n \in V(G)^{|\bar{x}|}$ and $(\bar{b}_J)_{J \subseteq [n]}$ in $V(G)^{|\bar{y}|}$ such that for all $i \in [n]$ and $J \subseteq [n]$,

$$G \models \varphi(\bar{a}_i, \bar{b}_J) \quad \iff \quad i \in J.$$

We use the following result of Baldwin and Shelah [BS85, Theorem 8.1.8] (see also [Prz26, Theorem 2.58 and Proposition 2.65] for a discussion).

Theorem A.4. *A class $\mathscr{C}$ is monadically dependent if and only if for all $k \in \mathbb{N}$, no first-order formula $\varphi(\bar{x}, \bar{y})$ has the independence property on the class $\mathscr{C}^{(k)}$ of all k-colorings of graphs in $\mathscr{C}$.*

Proof of Lemma A.2. Suppose that $\mathscr{C}$ is monadically dependent, and let $\varphi(\bar{x}, z)$ and $\psi(\bar{y}, z)$ be first-order formulas in the signature of $\mathscr{C}$. Let $G \in \mathscr{C}$ and suppose that $\bar{a}_1, \dots, \bar{a}_m \in V(G)^{|\bar{x}|}$, $\bar{b}_1, \dots, \bar{b}_m \in V(G)^{|\bar{y}|}$, and $\bar{c}_{i,j} \in V(G)^{|z|}$ for $i, j \in [m]$ form a clean (φ, ψ)-pattern of order $m \times m$ in G.

Fix an arbitrary binary relation $B \subseteq [m] \times [m]$, and let G^+ be the expansion of G with a unary relation U interpreted as $\{c_{i,j} : (i,j) \in B\}$. Consider the following formula:

$$\gamma(\bar{x}, \bar{y}) := \exists z \, U(z) \wedge \varphi(\bar{x}, z) \wedge \psi(\bar{y}, z).$$

By the definition of a clean pattern, the above imply that

$$G^+ \models \gamma(\bar{a}_i, \bar{b}_j) \quad \text{if and only if} \quad (i,j) \in B. \tag{18}$$

Indeed, if $(i,j) \in B$ then the existential quantifier is witnessed by $c_{i,j}$. Conversely, a witness to the existential quantifier must be of the form $c_{i',j'}$ for some $(i',j') \in B$, and then the definition of a clean pattern implies that $i = i'$ and $j = j'$, thus $(i,j) \in B$.

Let k be the largest integer such that $2^k \leqslant m$. Fix an arbitrary injection $h \colon P([k]) \to [m]$ from the powerset of $[k]$ into $[m]$. As $B \subseteq [m] \times [m]$ may be chosen arbitrarily, we may set

$$B := \{(i, h(J)) : i \in [k], J \subseteq [k], i \in J\}.$$

By (18), $G^+ \models \gamma(\bar{a}_i, \bar{b}_{h(J)})$ holds if and only if $i \in J$, for all $i \in [k]$ and $J \subseteq [k]$. As $\mathscr{C}$ is monadically dependent, Theorem A.4 implies that $k \leqslant O_{\mathscr{C},\gamma}(1)$. Since the formula γ depends only on the original formulas φ and ψ, this implies $k \leqslant O_{\mathscr{C},\varphi,\psi}(1)$. As $m < 2^{k+1}$, we conclude that $m \leqslant O_{\mathscr{C},\varphi,\psi}(1)$, as required. ∎

To prove Lemma A.3, we use the following version of Ramsey's theorem (see [DMT24, full version, Lemma 4.4]). For a tuple $\bar{a} = (a_1, \dots, a_\ell) \in [n]^\ell$, the *order type* $\operatorname{otp}(\bar{a})$ of $\bar{a}$ is the atomic type of $\bar{a}$ in the ordered set $([n], \leqslant)$. More explicitly, it stores the information, for each $i, j \in [\ell]$, whether $a_i < a_j$, $a_i = a_j$, or $a_i > a_j$.

Lemma A.5. *For every $k, \ell \in \mathbb{N}$ there is an unbounded, non-decreasing function $U_{k,\ell} \colon \mathbb{N} \to \mathbb{N}$ with the following property. For every $t \in \mathbb{N}$ and coloring*

$$c \colon [t]^\ell \times [t]^\ell \to [k]$$

there are subsets $I, J \subseteq [t]$ of size $U_{k,\ell}(t)$ such that $c(\bar{a}, \bar{b})$ only depends on $\operatorname{otp}(\bar{a})$ and $\operatorname{otp}(\bar{b})$, for all $\bar{a} \in I^\ell$ and $\bar{b} \in J^\ell$.

Proof of Lemma A.3. Let $\mathscr{C}$ be a dp-minimal class of structures, and let $\varphi(\bar{x}, z)$ and $\psi(\bar{y}, z)$ be first-order formulas in the signature of $\mathscr{C}$. Suppose that $(\bar{a}_i : i \in [t])$, $(\bar{b}_j : j \in [t])$ and $(c_{i,j} : i, j \in [t])$ form a (φ, ψ)-pattern of order $t \times t$ in some $G \in \mathscr{C}$. We prove that t is bounded by some constant depending on $\mathscr{C}$, φ, and ψ.

The proof proceeds in two steps, restricting the considered profiles depicted in the right-hand side part of Figure 4. In the first step, we arrive at a (φ, ψ)-pattern where all of the ?'s in those profiles are equal to some fixed boolean value 0 or 1. In the second step, we arrive at an (α, β)-pattern, for some related formulas α, β, for which the ?'s are further restricted to 0.

Step 1. As a first step we show that we can find sets $I, J \subseteq [t]$ of large size m, to ensure the following property:

There are $e_1, e_2 \in \{\mathit{true}, \mathit{false}\}$ such that for all $i, i' \in I$ and $j, j' \in J$,

3. $G \models \varphi(\bar{a}_{i'}, c_{i,j}) \Leftrightarrow e_1$ holds if $i' < i$,
4. $G \models \psi(\bar{b}_{j'}, c_{i,j}) \Leftrightarrow e_2$ holds if $j' < j$,

Indeed, define a coloring

$$c \colon [t]^2 \times [t]^2 \to \{\mathit{true}, \mathit{false}\}^2$$

where $c((i, i'), (j, j'))$ stores the truth values of $G \models \varphi(\bar{a}_{i'}, c_{i,j})$ and $G \models \psi(\bar{b}_{j'}, c_{i,j})$. By Lemma A.5, there are subsets $I, J \subseteq [t]$ of size $U(t)$, for some unbounded function $U = U_{4,2}$, such that $c((i, i'), (j, j'))$ only depends on $\operatorname{otp}(i, i')$ and $\operatorname{otp}(j, j')$ for all $i, i' \in I$ and $j, j' \in J$.

Let e_1 be the first component of $c((i, i'), (j, j'))$ for $i' < i \in I$ and $j = j' \in J$, and let e_2 be the second component of $c((i, i'), (j, j'))$ for $i = i' \in I$ and $j' < j \in J$. This choice satisfies the conditions 3. and 4. above.

Step 2. By reindexing and removing the last element of each of the sequences if necessary, assume that $I = J = [2m]$ for some $m \in \mathbb{N}$ with $m \geqslant U(t)/2 - 1$ for some unbounded function U, and that all conditions 1.–4. above hold. As a next step, we show that we can modify the formulas and sequences for which conditions 1.–4. hold with $e_1 = e_2 = \mathit{false}$.

Fix $i, j \in [2m]$ and for $i' \in [2m]$ let $f_{i'} \in \{\mathit{true}, \mathit{false}\}$ denote the truth value of $G \models \varphi(\bar{a}_{i'}, c_{i,j})$. Then $f_1, \ldots, f_{2m}$ consists of a sequence of e_1's, followed by a single *true* at position i, which is followed by a sequence of *false*'s. Grouping the sequence into consecutive pairs, and assuming that i is odd, we get a sequence of the form:

$$(e_1, e_1), \ldots, (e_1, e_1), (\mathit{true}, \mathit{false}), (\mathit{false}, \mathit{false}), \ldots, (\mathit{false}, \mathit{false}).$$

There is exactly one pair (p, q) in this sequence, namely the $\frac{i+1}{2}$th one, with distinct truth values, that is, for which $p \operatorname{XOR} q$ is *true*. A similar reasoning applies to ψ and $(\bar{b}_{j'} : j' \in [2m])$. This motivates us to consider the following formulas and sequences. Define

$$\alpha(\bar{x}_1 \bar{x}_2, z) \coloneqq \varphi(\bar{x}_1, z) \operatorname{XOR} \varphi(\bar{x}_2, z),$$
$$\beta(\bar{y}_1 \bar{y}_2, z) \coloneqq \psi(\bar{y}_1, z) \operatorname{XOR} \psi(\bar{y}_2, z).$$

For $i \in [m]$, let $\bar{u}_i$ be the $2|\bar{x}|$-tuple $\bar{a}_{2i-1}\bar{a}_{2i}$, and for $j \in [m]$, let $\bar{w}_j$ be the $2|\bar{y}|$-tuple $\bar{b}_{2j-1}\bar{b}_{2j}$. Finally, let $d_{i,j} = c_{2i-1,2j-1}$. By the observation above, for all $i, i', j, j' \in [m]$ we have that

(*a*) $G \models \alpha(\bar{u}_{i'}, d_{i,j})$ if and only if $i' = i$,

(*b*) $G \models \beta(\bar{w}_{j'}, d_{i,j})$ if and only if $j' = j$.

Thus, $(\bar{u}_i : i \in [m])$, $(\bar{w}_j : j \in [m])$, and $(d_{i,j} : i, j \in [m])$ form a clean (α, β)-pattern of order $m \times m$. Dp-minimality therefore bounds m by a constant depending only on $\mathscr{C}, \alpha$, and β, hence only on $\mathscr{C}, \varphi$, and ψ. Since $m \geqslant U(t)/2 - 1$ and U is unbounded and non-decreasing, this also bounds t by a constant depending only on $\mathscr{C}, \varphi$, and ψ, as required. ∎

B Corollary of Gaifman's locality theorem

In this appendix, we prove Lemma 4.9 and Lemma 6.1.

We will use Gaifman's Locality Theorem [Gai82]. We recommend [Gro07, Section 4.1] for an introduction. First some definitions.

Fix a finite relational signature σ of arity α. Two σ-formulas $\varphi(\bar{x})$ and $\psi(\bar{x})$ are *equivalent* if for every σ-structure A and every tuple $\bar{a} \in V(A)^{|\bar{x}|}$

$$A \models \varphi(\bar{a}) \Leftrightarrow A \models \psi(\bar{a}).$$

Let $\mathrm{adj}_\sigma(x_1, x_2)$ be the formula that tests whether x_1 and x_2 are adjacent in the Gaifman graph. More precisely, adj_σ tests for every relation $R \in \sigma$ of arity β with $2 \leqslant \beta \leqslant \alpha$, whether $x_1 \neq x_2$ and there exist elements $x_3, \dots, x_\beta$ such that some permutation of $x_1, \dots, x_\beta$ satisfies R. Thus, the quantifier rank of adj_σ can be bounded by $\alpha - 2$. We can naturally derive a formula $\mathrm{dist}(x_1, x_2) \leqslant r$ checking whether x_1 and x_2 have distance at most r in the Gaifman graph. We can (generously) bound the quantifier rank of $(\mathrm{dist}(x_1, x_2) \leqslant r)$ by $r + \alpha$ (quantifying candidate vertices for the path and then checking the required adjacencies "in parallel").

A σ-formula $\varphi(\bar{x})$ is *r-local* if every quantifier appearing in $\varphi(\bar{x})$ is of the form $\forall y \in \mathrm{Ball}_s(\bar{x})$ or $\exists y \in \mathrm{Ball}_s(\bar{x})$ for some $s \leqslant r$, where the ball checks are implemented using the dist formulas (which we exempt from this condition). A sentence ψ is *basic r-local* if there is some $r' \leqslant r$ such that ψ is of the form

$$\exists x_1, \dots, x_m \bigwedge_{1 \leqslant i < j \leqslant m} \mathrm{dist}(x_i, x_j) > 2r' \wedge \bigwedge_{1 \leqslant i \leqslant m} \varphi(x_i)$$

for an r'-local formula φ.

Intuitively: whether an r-local formula holds on a tuple depends only on the r-ball of that tuple, as made precise in the following lemma.

Lemma B.1. *Let $\varphi(\bar{x})$ be an r-local σ-formula, let A be a σ-structure, let $\bar{a} \in V(A)^{|\bar{x}|}$, and let $\mathrm{Ball}_r^A(\bar{a}) \subseteq B \subseteq V(A)$. Then $A \models \varphi(\bar{a}) \Leftrightarrow A[B] \models \varphi(\bar{a})$.*

Proof. When φ is evaluated in $A[B]$, its bounded quantifiers range over balls in the Gaifman graph of $A[B]$, which for arity above two can have fewer edges than the Gaifman graph of A restricted to B. Still, $\mathrm{Ball}_s^{A[B]}(\bar{a}) = \mathrm{Ball}_s^A(\bar{a})$ for all $s \leqslant r$, since a shortest path from $\bar{a}$ of length at most r survives in $A[B]$: every tuple witnessing its j-th edge has all entries at distance at most $j \leqslant r$ from $\bar{a}$, hence in B. So the bounded quantifiers of φ range over the same subsets of B in A and in $A[B]$, and since $A[B]$ is an induced substructure of A, the claim follows. ∎

The following theorem is one of the key tools for limiting the expressiveness of first-order logic.

Theorem B.2 (Gaifman's Locality Theorem [Gai82]). *For all $\alpha, q, k \in \mathbb{N}$ there is a $p = p(\alpha, q, k)$ computable from α, q, k with the following property. Let σ be a finite relational signature of arity at most α. Every σ-formula $\varphi(\bar{x})$ of quantifier rank at most q and with $|\bar{x}| \leqslant k$ free variables is equivalent to a σ-formula $\varphi^{\mathrm{loc}}(\bar{x})$, computable from φ, that is a Boolean combination of basic 7^q-local sentences and 7^q-local formulas,* **each of quantifier rank at most p.**

The above theorem was proven in [Gai82], but the paper does not explicitly state the highlighted bound p on the quantifier rank. Analyzing the proof in [Gai82] shows that, indeed, there exists a bound p that only depends on α, q, and k, but not on the specific choice of the signature σ. This observation is made explicit in the recent rank-preserving locality theorems presented in [DT26] and [Ber+26]. In all three papers [Gai82; DT26; Ber+26], the dependence of p on the arity α is subtle. Each of them produces a formula in some extension of first-order logic (FO). While these extensions of FO are equivalent in expressive power to plain FO, they each contain predicates measuring distances in the Gaifman graph of a structure. As discussed above, converting such distance formulas to plain FO introduces an additional nesting of quantifiers depending on the arity α.

Lemma 4.9 (Corollary of Gaifman locality). *For every $k, q, \alpha \in \mathbb{N}$, there are numbers $h = h(q)$ and $q' = q'(k, q, \alpha)$ computable from k, q, α, with the following property. Fix a finite relational signature σ of arity at most α and a σ-structure A. Then, for all $\bar{a} \in V(A)^k$ and $B \supseteq \mathrm{Ball}_h^A(\bar{a})$, $\mathrm{tp}_q^A(\bar{a})$ depends only on $\mathrm{tp}_{q'}^{A[B]}(\bar{a})$.*

Proof. Let $h := 7^q$, and let $q' := p$ be given by Theorem B.2 for α, q, k. Let $\bar{a}, \bar{b} \in V(A)^k$, $B_a \supseteq \mathrm{Ball}_h^A(\bar{a})$ and $B_b \supseteq \mathrm{Ball}_h^A(\bar{b})$ with $\mathrm{tp}_{q'}^{A[B_a]}(\bar{a}) = \mathrm{tp}_{q'}^{A[B_b]}(\bar{b})$. By Theorem B.2, every formula $\varphi(\bar{x})$ of quantifier rank at most q is equivalent to a Boolean combination of sentences and h-local formulas of quantifier rank at most q'. The sentences have a fixed truth value in A independent of $\bar{a}$ and $\bar{b}$. For an h-local formula $\psi(\bar{x})$ of quantifier rank at most q',

$$A \models \psi(\bar{a}) \iff A[B_a] \models \psi(\bar{a}) \iff A[B_b] \models \psi(\bar{b}) \iff A \models \psi(\bar{b}),$$

where the middle equivalence follows from the assumed equality of the q'-types. The first and third equivalences follow from Lemma B.1. Hence, $A \models \varphi(\bar{a}) \iff A \models \varphi(\bar{b})$ for all formulas of quantifier rank at most q, so $\mathrm{tp}_q^A(\bar{a}) = \mathrm{tp}_q^A(\bar{b})$. ∎

Lemma B.3. *For every $q, k, c \in \mathbb{N}$ there are numbers $r = r(q)$ and $T = T(q, k, c)$ with the following property. Let H be a graph expanded by at most c unary predicates, and let $U \subseteq V(H)$ be finite.*

Consider the set $F := \{a \in V(H) : \mathrm{dist}^H(a, U) > r\}$ of vertices that are far from U. We can bound

$$|\mathrm{Types}_{k,q}^H(F/U)| \leqslant T(q, k, c).$$

Proof. Let h and q' be given by Lemma 4.9 for k, q, and $\alpha = 2$. Since $h = h(q)$, the radius $r := h + 1$ depends only on q.

Let H^+ be the expansion of H by the unary predicates $P_u := \{u\}$ and $N_u := N^H(u)$ for all $u \in U$. Then H^+ has a finite relational signature of arity two, and its Gaifman graph is the underlying graph of H. Take a formula with parameters from U. Replace every atom $x = u$ or $u = x$ by $P_u(x)$, every atom $E(x, u)$ or $E(u, x)$ by $N_u(x)$, and every atom without variables by its truth value in H. This gives a parameter-free formula in the signature of H^+ that has the same quantifier rank and is satisfied by the same tuples. Hence $\mathrm{tp}_q^{H^+}(\bar{a})$ determines $\mathrm{tp}_q^H(\bar{a}/U)$.

Let $\bar{a} \in F^k$ and $B := \mathrm{Ball}_h^{H^+}(\bar{a})$. Every $w \in B$ satisfies $\mathrm{dist}^H(w, U) > r - h = 1$, so B avoids U and all neighbors of vertices in U. Thus every P_u and N_u is empty in $H^+[B]$. So $H^+[B]$ is the induced subgraph $H[B]$ expanded by empty predicates. Then $\mathrm{tp}_{q'}^{H[B]}(\bar{a})$ determines $\mathrm{tp}_{q'}^{H^+[B]}(\bar{a})$. By Lemma 4.9, $\mathrm{tp}_{q'}^{H^+[B]}(\bar{a})$ determines $\mathrm{tp}_q^{H^+}(\bar{a})$, and hence $\mathrm{tp}_q^H(\bar{a}/U)$. Here we crucially use that h and q' do not depend on the signature of H^+, which grows with $|U|$. So $|\mathrm{Types}_{k,q}^H(F/U)|$ is at most the number of parameter-free rank-q' types of k-tuples in graphs with c unary predicates. This number is finite and depends only on q', k, and c, hence only on q, k, and c. Call it $T(q,k,c)$. ■

The following is an adaptation of [PST18, Corollary 16].

Lemma 6.1. *For every $q, k, c \in \mathbb{N}$ there are numbers $r = r(q)$ and $T = T(q,k,c)$ with the following property. Let H be a graph expanded by at most c unary predicates, and let $U \subseteq V(H)$ be finite. Fix a tuple $\bar{u}$ of vertices of U of finite length. Let $F := \{a \in V(H) : \mathrm{dist}^H(a, U) > r\}$. Then*

$$\left|\{\mathrm{tp}_q^H(\bar{a}\bar{u}/U) : \bar{a} \in F^k\}\right| \leqslant T(q,k,c).$$

Proof. Let $r = r(q)$ and $T = T(q,k,c)$ be given by Lemma B.3, yielding the desired bound

$$\left|\{\mathrm{tp}_q^H(\bar{a}/U) : \bar{a} \in F^k\}\right| \leqslant T(q,k,c)$$

in case no tuple $\bar{u}$ is involved. To lift that result, it remains to show that for a fixed tuple $\bar{u}$ with elements from U, the type $\mathrm{tp}_q^H(\bar{a}\bar{u}/U)$ depends only on $\mathrm{tp}_q^H(\bar{a}/U)$. More precisely, it suffices to show that there exists a map $\varphi \mapsto \varphi_{\bar{u}}$ such that every formula $\varphi(\bar{x}\bar{y})$ with

$$\text{quantifier rank} \leqslant q, \qquad |\bar{x}| = k \text{ and } |\bar{y}| = |\bar{u}|, \qquad \text{and parameters from } U$$

gets mapped to a formula $\varphi_{\bar{u}}(\bar{x})$ of the same quantifier rank and again with parameters from U such that for every $\bar{a} \in F^k$

$$\varphi \in \mathrm{tp}_q^H(\bar{a}\bar{u}/U) \iff H \models \varphi(\bar{a}\bar{u}) \iff H \models \varphi_{\bar{u}}(\bar{a}) \iff \varphi_{\bar{u}} \in \mathrm{tp}_q^H(\bar{a}/U).$$

We can obtain a suitable $\varphi_{\bar{u}}(\bar{x})$ from $\varphi(\bar{x}\bar{y})$ by substituting in φ each free variable $y_i \in \bar{y}$ with the corresponding parameter $u_i \in \bar{u}$ from U. This translation depends only on $\bar{u}$ and not on $\bar{a}$. It does not change the quantifier rank, and still all parameters of $\varphi_{\bar{u}}$ are contained in U. ■

Remark B.4. In Lemma 6.1, the bound crucially relies on $\bar{u}$ being fixed, as the translation $\varphi \mapsto \varphi_{\bar{u}}$ differs for different $\bar{u}$. Without fixing $\bar{u}$ there is no bound independent of $|U|$: for $k = 1$, $a \in F$, and $\bar{u} = (u)$ ranging over U, the formulas $\psi_u(xy) := (y = u)$ (each using a different parameter u) show that the tuples au realize $|U|$ different types over U.

C Locally definable construction sequences in nowhere dense classes

We prove Theorem 4.8, that in nowhere dense classes, we can obtain definable construction sequences of length $O(n^{1+\varepsilon})$ for any $\varepsilon > 0$, with $O(n^\varepsilon)$ parameters and quantifier rank 0.

Theorem 4.8. *Let $\mathscr{C}$ be a nowhere dense graph class, $r \in \mathbb{N}$, and $\varepsilon > 0$. Every n-vertex graph $G \in \mathscr{C}$ has a radius-r definable construction sequence with one unary predicate, $O_{\mathscr{C},r,\varepsilon}(n^\varepsilon)$ parameters, quantifier rank 0, and length $O_{\mathscr{C},r,\varepsilon}(n^{1+\varepsilon})$. Such a sequence can be computed in time $O_{\mathscr{C},r,\varepsilon}(n^{1+\varepsilon})$.*

Proof sketch. Let $G \in \mathscr{C}$ have n vertices. We use the almost bounded *weak coloring numbers* of nowhere dense classes [NO11; GKS17], through the following consequence [Dra+26b]. For $v \notin U$, let $S_s(v, U)$ be the set of vertices in U reachable from v by a path of length at most s whose internal vertices avoid U. There is a total order on $V(G)$ such that $|S_{r+1}(v, U)| \leqslant O_{\mathscr{C},r,\varepsilon}(n^\varepsilon)$ for every prefix U of this order and every $v \notin U$. Such an order can be computed in time $O_{\mathscr{C},r,\varepsilon}(n^{1+\varepsilon})$ by the algorithm of Grohe, Kreutzer, and Siebertz [GKS17] for weak coloring orders: every vertex of $S_{r+1}(v, U)$ is weakly $(r+1)$-reachable from v, since it is the smallest vertex on the witnessing path.

For each prefix U, consider the partial graph H_U on $V(G)$ whose positive resolved pairs are the edges of $G[V(G) - U]$, whose negative relation is empty, and whose parts are the atomic types in G over U. Thus the vertices of U are singleton parts, and the remaining vertices are grouped by their neighborhoods in U. Starting with $U = V(G)$, remove the last vertex u of U one at a time. To pass from H_U to $H_{U'}$, where $U' = U - \{u\}$, first resolve positively between $\{u\}$ and each part outside U adjacent to u in G, and then merge parts with the same atomic type over U'. Finally, when U is empty, one negative resolve within the unique part completes the sequence. There are at most $|E(G)|$ positive resolves and $n-1$ merges, giving length $O_{\mathscr{C},r,\varepsilon}(n^{1+\varepsilon})$ by the sparsity of $\mathscr{C}$.

To implement the construction, first compute the atomic-type partitions over the prefixes in forward order by partition refinement: when adding u, split off $\{u\}$ and refine the remaining parts by adjacency to u. Using adjacency lists, this takes $O(n + |E(G)|)$ time in total. Record these splits and the resulting adjacent parts outside the prefix. Processing the prefixes in reverse order, resolve positively against the recorded adjacent parts and undo the recorded splits by merges. Thus the sequence can be computed in time $O_{\mathscr{C},r,\varepsilon}(n^{1+\varepsilon})$, including the computation of the order.

To see local definability, consider two consecutive partial graphs H, H' during the transition from H_U to $H_{U'}$, and put $B = \mathrm{Ball}_r^{H'}(v)$. Vertices of U' are still isolated, so their balls are singletons and are easy to define. Otherwise $B \subseteq V(G) - U'$. Every resolved pair of H' is an edge of G, so every vertex of U' adjacent in G to a vertex of B belongs to $S_{r+1}(v, U')$. All vertices of U' outside this separator are therefore nonadjacent to every vertex of B. Thus, on B, the atomic types in G over U are determined by adjacency and equality tests against the parameters

$$S_{r+1}(v, U') \cup \{u\}.$$

This set has size $O_{\mathscr{C},r,\varepsilon}(n^\varepsilon)$. Each atomic-type part is therefore quantifier-free definable, and merging parts corresponds to taking disjunctions of their definitions. The resolved edges are the original edges between vertices outside U, together with edges from u to the parts already resolved against it, so they too have quantifier-free definitions using these parameters. One unary predicate naming B restricts these definitions to the required domain. The final negative resolve is immediate, since just before it the positive relation is $E(G)$ and there is only one part. Thus the same bounds apply to $H[\mathrm{Ball}_r^{H'}(v)]$, including the intermediate merges and resolves. ∎