Canonical Decompositions in Monadically Stable and Bounded Shrubdepth Graph Classes

Pierre Ohlmann
Institute of Informatics, University of Warsaw, Poland

Michał Pilipczuk
Institute of Informatics, University of Warsaw, Poland

Wojciech Przybyszewski
Institute of Informatics, University of Warsaw, Poland

Szymon Toruńczyk
Institute of Informatics, University of Warsaw, Poland

Abstract

We use model-theoretic tools originating from stability theory to derive a result we call the Finitary Substitute Lemma, which intuitively says the following. Suppose we work in a stable graph class $\mathscr{C}$, and using a first-order formula φ with parameters we are able to define, in every graph $G \in \mathscr{C}$, a relation R that satisfies some hereditary first-order assertion ψ. Then we are able to find a first-order formula φ' that has the same property, but additionally is *finitary*: there is finite bound $k \in \mathbb{N}$ such that in every graph $G \in \mathscr{C}$, different choices of parameters give only at most k different relations R that can be defined using φ'.

We use the Finitary Substitute Lemma to derive two corollaries about the existence of certain canonical decompositions in classes of well-structured graphs.

- We prove that in the Splitter game, which characterizes nowhere dense graph classes, and in the Flipper game, which characterizes monadically stable graph classes, there is a winning strategy for Splitter, respectively Flipper, that can be defined in first-order logic from the game history. Thus, the strategy is canonical.
- We show that for any fixed graph class $\mathscr{C}$ of bounded shrubdepth, there is an $\mathcal{O}(n^2)$-time algorithm that given an n-vertex graph $G \in \mathscr{C}$, computes in an isomorphism-invariant way a structure H of bounded treedepth in which G can be interpreted. A corollary of this result is an $\mathcal{O}(n^2)$-time isomorphism test and canonization algorithm for any fixed class of bounded shrubdepth.

2012 ACM Subject Classification Theory of computation → Finite Model Theory

Keywords and phrases Model Theory, Stability Theory, Shrubdepth, Nowhere Dense, Monadically Stable

Digital Object Identifier 10.4230/LIPIcs.ICALP.2023.135

Category Track B: Automata, Logic, Semantics, and Theory of Programming

Related Version *Full Version*: https://arxiv.org/abs/2303.01473

Funding This paper is part of a project that have received funding from the European Research Council (ERC) (grant agreement No 948057 – BOBR).

Acknowledgements We are indebted to Pierre Simon for enlightening discussions about classic results in stability theory that greatly helped in deriving the results presented in this work. We also thank Jakub Gajarský, Rose McCarty, and Marek Sokołowski for their contribution in discussions around the topic of this work.

1 Introduction

Stability theory is a well-established branch of model theory devoted to the study of stable theories, or equivalently classes of structures that are models of such theories. Here, we say that a formula[1] $\varphi(\bar{x};\bar{y})$ is *stable* on a class of relational structures $\mathscr{C}$ if there is an integer $k \in \mathbb{N}$ such that for every $\mathbf{M} \in \mathscr{C}$, one cannot find tuples $\bar{u}_1, \ldots, \bar{u}_k \in \mathbf{M}^{\bar{x}}$ and $\bar{v}_1, \ldots, \bar{v}_k \in \mathbf{M}^{\bar{y}}$ such that for all $i, j \in \{1, \ldots, k\}$,

$$\mathbf{M} \models \varphi(\bar{u}_i, \bar{v}_j) \qquad \text{if and only if} \qquad i < j.$$

Then $\mathscr{C}$ is stable if every formula is stable on $\mathscr{C}$. Intuitively, this means that using a fixed formula, one cannot interpret arbitrarily long total orders in structures from $\mathscr{C}$. We refer to the textbooks of Pillay [24] or of Tent and Ziegler [27] for an introduction to stability.

The goal of this paper is to use certain classic results of stability theory, particularly the understanding of *forking* in stable theories, to derive statements about the existence of canonical decompositions in certain classes of well-structured graphs. Here, we model graphs as relational structures with one binary adjacency relation that is symmetric.

Finitary Substitute Lemma. Our main model-theoretic tool is the Finitary Substitute Lemma, which we state below in a simplified form; see Lemma 10 for a full statement.

To state the lemma, we need some definitions. A formula $\varphi(\bar{x};\bar{y})$ is *finitary* on a class of structures $\mathscr{C}$ if there exists $k \in \mathbb{N}$ such that for every $\mathbf{M} \in \mathscr{C}$, we have

$$\left|\left\{\varphi(\mathbf{M}^{\bar{x}}, \bar{v}) \colon \bar{v} \in \mathbf{M}^{\bar{y}}\right\}\right| \leqslant k,$$

where $\varphi(\mathbf{M}^{\bar{x}}, \bar{v}) = \{\bar{u} \in \mathbf{M}^{\bar{x}} \mid \mathbf{M} \models \varphi(\bar{u}, \bar{v})\}$. In other words, $\varphi(\bar{x};\bar{y})$ is finitary on $\mathscr{C}$ if by substituting different parameters for $\bar{y}$ in any model $\mathbf{M} \in \mathscr{C}$, one can define only at most k different relations on $\bar{x}$-tuples. Next, a sentence ψ is *hereditary* if for every model $\mathbf{M}$ and its induced substructure $\mathbf{M}'$, $\mathbf{M} \models \psi$ implies $\mathbf{M}' \models \psi$. Finally, for a relation $R(\bar{x})$ present in the signature, formula $\varphi(\bar{x})$ (possibly with parameters), and sentence ψ, by $\psi[R(\bar{x})/\varphi(\bar{x})]$ we mean the sentence derived from ψ by substituting every occurrence of R with formula φ.

▶ **Lemma 1** (Finitary Substitute Lemma, simplified version). *Let $\mathscr{C}$ be a stable class of structures. Suppose $\varphi(\bar{x};\bar{y})$ is a formula and ψ a hereditary sentence such that for every $G \in \mathscr{C}$,*

$$\textit{there exists } \bar{s} \in G^{\bar{y}} \textit{ such that } G \models \psi[R(\bar{x})/\varphi(\bar{x};\bar{s})]. \tag{1}$$

Then there exists a formula $\varphi'(\bar{x}, \bar{z})$ that also satisfies (1)*, but is additionally finitary on $\mathscr{C}$.*

Thus, intuitively, the Finitary Substitute Lemma says that in stable classes, every relation that is definable with parameters can be replaced by a finitary one, as long as we care that the relation satisfies some hereditary first-order assertion. The main observation of this paper is that this can be used in the context of various graph decompositions. Intuitively, if every step in decomposing the graph can be defined by a first-order formula with parameters, and the validity of the step can be verified using a hereditary first-order sentence, then we can use the Finitary Substitute Lemma to derive an equivalent definition of a step that is finitary. This yields only a bounded number of different steps that can be taken, making it possible to construct a decomposition that, in a certain sense, is canonical.

---

[1] All formulas considered in this paper are first-order, unless explicitly stated.

Classes of bounded shrubdepth. Our first application concerns classes of *bounded shrubdepth*. The concept of shrubdepth was introduced by Ganian et al. [17] to capture dense graphs that are well-structured in a shallow way. On one hand, classes of bounded shrubdepth are exactly those that can be interpreted, using first-order formulas with two free variables, in classes of forests of bounded depth. On the other hand, graphs from any fixed class of bounded shrubdepth admit certain decompositions, called *connection models*, which are essentially clique expressions of bounded depth. (See Section 5.1 for a definition of a connection model.) Thus, in particular every graph class of bounded shrubdepth has bounded cliquewidth, but classes of bounded shrubdepth are in addition stable [17].

Shrubdepth is a dense counterpart of *treedepth*, defined as follows: the treedepth of a graph G is the smallest integer d such that G is a subgraph of the ancestor/descendant closure of a rooted forest of depth at most d. In particular, every class of graphs of bounded treedepth has bounded shrubdepth; boundedness of treedepth and of shrubdepth is in fact equivalent assuming that the class excludes some biclique $K_{t,t}$ as a subgraph [17]. In essence, treedepth is a bounded-depth counterpart of treewidth in the same way as shrubdepth is a bounded-depth counterpart of cliquewidth.

In spite of the above, the combinatorial properties of shrubdepth are still much less understood than those of treedepth. For instance, a good understanding of subgraph obstacles allows one to construct suitable canonical decompositions for graphs of bounded treedepth. This allowed Bouland et al. [4] to design a graph isomorphism test that works in fixed-parameter time parameterized by the treedepth, or more precisely, in time $f(d) \cdot n^3 \log n$, where f is a computable function. While it is known that every class of bounded shrubdepth can be characterized by a finite number of forbidden induced subgraphs [17], it is unclear how to use just this result to design any kind of canonical decompositions for classes of bounded shrubdepth. Consequently, so far it was unknown whether the graph isomorphism problem can be solved in fixed-parameter time on classes of bounded shrubdepth[2]. The most efficient isomorphism test in this context is the one designed by Grohe and Schweitzer [20] for the cliquewidth parameterization: it works in XP time, that is, in time $n^{f(k)}$ where k is the cliquewidth and f is a computable function. See also the later work of Grohe and Neuen [19], which improves the XP running time and applies to the more general canonization problem.

We show that the Finitary Substitute Lemma can be used to bridge this gap by proving the following result.

▶ **Theorem 2.** *Let $\mathscr{C}$ be a class of graphs of bounded shrubdepth. Then there is a class $\mathscr{D}$ of binary structures of bounded treedepth and a mapping $\mathcal{A}\colon \mathscr{C} \to \mathscr{D}$ such that:*

- *For each $G \in \mathscr{C}$, the vertex set of G is contained in the domain of $\mathcal{A}(G)$ and the mapping $G \mapsto \mathcal{A}(G)$ is isomorphism-invariant.*
- *Given an n-vertex graph $G \in \mathscr{C}$, the structure $\mathcal{A}(G)$ has $\mathcal{O}(n)$ elements and can be computed in time $\mathcal{O}(n^2)$.*
- *There is a simple first-order interpretation I such that $G = \mathsf{I}(\mathcal{A}(G))$, for every $G \in \mathscr{C}$.*

Here, by isomorphism-invariance we mean that every isomorphism between $G, G' \in \mathscr{C}$ extends to an isomorphism between $\mathcal{A}(G)$ and $\mathcal{A}(G')$. Further, by a *simple* interpretation we mean a first-order interpretation that is 1-dimensional: vertices of G are interpreted in

[2] This statement is somewhat imprecise, as shrubdepth is defined as a parameter of a graph class, rather than of a single graph. By this we mean that there is a universal constant c such that for every graph class $\mathscr{C}$ of bounded shrubdepth, the isomorphism of graphs from $\mathscr{C}$ can be tested in $\mathcal{O}(n^c)$ time, with the constant hidden in the $\mathcal{O}(\cdot)$ notation possibly depending on $\mathscr{C}$.

single elements of $\mathcal{A}(G)$ (actually, every vertex is interpreted in itself). Thus, $\mathcal{A}(G)$ can be regarded as a canonical – obtained in an isomorphism-invariant way – sparse decomposition of G that encodes G faithfully and takes the form of a structure of bounded treedepth. We remark that certain logic-based sparsification procedures for classes of bounded shrubdepth were proposed in [8, 14], but these are insufficient for our applications, which we explain next.

The third point above together with the fact that $\mathcal{A}$ is isomorphism-invariant imply the following: for all $G, G' \in \mathscr{C}$, G and G' are isomorphic if and only if $\mathcal{A}(G)$ and $\mathcal{A}(G')$ are. We can now combine Theorem 2 with the approach of Bouland et al. [4] to give a fixed-parameter isomorphism test on classes of bounded shrubdepth.

▶ **Theorem 3.** *For every graph class $\mathscr{C}$ of bounded shrubdepth there is an $\mathcal{O}(n^2)$-time algorithm that given n-vertex graphs $G, G' \in \mathscr{C}$, decides whether G and G' are isomorphic.*

In fact, our algorithm solves also the general canonization problem, see Section 5.4.

We remark that the algorithm of Theorem 3 is non-uniform, in the sense that we obtain a different algorithm for every class $\mathscr{C}$. Despite the existence of parameters such as rankdepth [22] or SC-depth [17] that are suited for the treatment of single graphs and are equivalent in terms of boundedness on classes to shrubdepth, we do not know how to make our algorithm uniform even for the rankdepth or SC-depth parameterizations.

Finally, we believe that the construction behind our proof of Theorem 2 can be used to obtain an alternative proof of a result of Hliněný and Gajarský [13], later reproved by Chen and Flum [8]: the expressive power of first-order and monadic second-order logic coincide on classes of bounded shrubdepth. This direction will be explored in future research.

Nowhere dense and monadically stable classes. Second, we use the Finitary Substitute Lemma to provide canonical strategies in game characterizations of two important concepts in structural graph theory: nowhere dense classes and monadically stable classes. In both cases, a strategy in the game can be regarded as decompositions of the graph in question.

We start with some definitions. A *unary lift* of a class of graphs $\mathscr{C}$ is any class of structures $\mathscr{C}^+$ such that every member of $\mathscr{C}^+$ is obtained from a graph belonging to $\mathscr{C}$ by adding any number of unary predicates on vertices. A class of graphs $\mathscr{C}$ is *monadically stable* if every unary lift of $\mathscr{C}$ is stable. On the other hand, a class of graphs $\mathscr{C}$ is *nowhere dense* if for every $d \in \mathbb{N}$ there exists t such that no graph in $\mathscr{C}$ contains the d-subdivision of K_t as a subgraph.

Nowhere denseness is the most fundamental concept of uniform sparsity in graphs considered in the theory of Sparsity; see the monograph of Nešetřil and Ossona de Mendez [23] for an introduction to this area. A pinnacle result of this theory was derived by Grohe et al. [18]: the model-checking problem for first-order logic is fixed-parameter tractable on any nowhere dense graph class. As observed by Adler and Adler [2] using earlier results of Podewski and Ziegler [25], monadically stable classes are dense counterparts of nowhere dense classes in the following sense: every nowhere dense class is monadically stable, and nowhere denseness and monadic stability coincide when we assume the class to be sparse, for instance to exclude some biclique $K_{t,t}$ as a subgraph. This motivated the following conjecture [1], which is an object of intensive studies for the last few years: The model-checking problem for first-order logic is fixed-parameter tractable on every monadically stable class of graphs $\mathscr{C}$.

To approach this conjecture, it is imperative to obtain a better structural understanding of graphs from monadically stable classes. This is the topic of several very recent works [5, 6, 7, 11, 15]. In this work we are particularly interested in the results of Gajarský et al. [15], who characterized monadically stable classes of graphs through a game model called the *Flipper game*, which reflects the characterization of nowhere dense classes through the *Splitter game*, due to Grohe et al. [18].

The radius-r Splitter game is played on a graph G between two players: *Splitter* and *Connector*. In every round, Connector first chooses any vertex u and the current *arena* – graph on which the game is played – gets restricted to a ball of radius r around u. Then Splitter removes any vertex of the graph. The game finishes, with Splitter's win, when the arena becomes empty. Splitter's goal is to win the game as quickly as possible, while Connector's goal is to avoid losing for as long as possible. The *Flipper game* is defined similarly, except that the moves of *Flipper* – who replaces Splitter – are as follows. Instead of removing a vertex, Flipper selects any subset of vertices F and performs a *flip*: replaces all edges with both endpoints in F with non-edges, and vice versa. Also, the game finishes when the arena consists of one vertex.

Grohe et al. [18] proved that a class of graphs $\mathscr{C}$ is nowhere dense if and only if for every radius $r \in \mathbb{N}$ there exists $\ell \in \mathbb{N}$ such that on every graph from $\mathscr{C}$, Splitter can win the radius-r Splitter game within at most ℓ rounds. This characterization is the backbone of their model-checking result for nowhere dense classes, as a strategy in the Splitter game provides a shallow decomposition of the graph in question, useful for understanding its first-order properties. Very recently, Gajarský et al. [15] proved an analogous characterization of monadically stable classes in terms of the Flipper game, and subsequently Dreier et al. [10] used this characterization to prove fixed-parameter tractability of the model-checking first order logic on monadically stable classes of graphs which possess so-called sparse neighborhood covers. Given this state-of-the-art, it is clear that a better understanding of strategies for Splitter and Flipper in the respective games may lead to a deeper insight into decompositional properties of nowhere dense and monadically stable graph classes.

In the Splitter game, we prove using just basic compactness, that in any arena there is only a bounded number of possible Splitter's moves that are *progressing*: lead to an arena where the Splitter can win in one less round. (See Theorem 6 for a formal statement.) So this gives a transparent canonical strategy for Splitter: just play all progressive moves one by one, in any order. Obtaining a similar canonicity result for strategies in the Flipper game requires the full power of our Finitary Substitute Lemma, discussed above.

In the interest of space, we have omitted from this version our results about canonical strategies in the Flipper game, as well as most proofs. For a complete exposition, we refer to the full version of the paper.

2 Preliminaries

Models. We work with first-order logic over a fixed signature Σ that consists of (possibly infinitely many) constant symbols and relation symbols. A *model* is a Σ-structure, and is typically denoted $\mathbf{M}, \mathbf{N}$, etc. We usually do not distinguish between a model and its domain, when writing, for instance, $m \in \mathbf{M}$ or $X \subseteq \mathbf{M}$. A graph G is viewed as a model over the signature consisting of one binary relation denoted E, indicating adjacency between vertices.

If $\bar{x}$ is a finite set of variables, then we write $\varphi(\bar{x})$ to denote a first-order formula φ with free variables contained in $\bar{x}$. We may also write $\varphi(\bar{x}_1, \ldots, \bar{x}_k)$ to denote a formula whose free variables are contained in $\bar{x}_1 \cup \ldots \cup \bar{x}_k$. We will write x instead of $\{x\}$ in case of a singleton set of variables, e.g. $\varphi(x, y)$ will always refer to a formula with two free variables x and y. We sometimes write $\varphi(\bar{x}; \bar{y})$ to distinguish a partition of the set of free variables of φ into two parts, $\bar{x}$ and $\bar{y}$; this partition plays an implicit role in some definitions. A Σ-formula $\varphi(\bar{x})$ *with parameters* from a set $A \subseteq \mathbf{M}$ is a formula $\varphi(\bar{x})$ over the signature $\Sigma \uplus A$, where the elements of A are treated as constant symbols (which are interpreted by themselves).

If U is a set and $\bar{x}$ is a set of variables, then $U^{\bar{x}}$ denotes the set of all $\bar{x}$*-tuples* $\bar{a}\colon \bar{x} \to U$ of $\bar{x}$ in U. For a formula $\varphi(\bar{x})$ (with or without parameters) and an $\bar{x}$-tuple $\bar{m} \in \mathbf{M}^{\bar{x}}$, we write $\mathbf{M} \models \varphi(\bar{m})$ if the valuation $\bar{m}$ satisfies the formula $\varphi(\bar{x})$ in $\mathbf{M}$. For a formula $\varphi(\bar{x};\bar{y})$ and a tuple $\bar{b} \in \mathbf{M}^{\bar{y}}$ we denote by $\varphi(\mathbf{M}^{\bar{x}};\bar{b})$ the set of all $\bar{a} \in \mathbf{M}^{\bar{x}}$ such that $\mathbf{M} \models \varphi(\bar{a};\bar{b})$.

Theories and compactness. A *theory* T (over Σ) is a set of Σ-sentences. The theory of a class of structures $\mathscr{C}$ is the set of sentences that hold in every structure $\mathbf{M} \in \mathscr{C}$. For instance, the theory of a class of graphs $\mathscr{C}$ contains sentences expressing that the relation E is symmetric and irreflexive. A model of a theory T is a structure $\mathbf{M}$ such that $\mathbf{M} \models \varphi$ for all $\varphi \in T$. When a theory has a model, it is said to be *consistent*.

▶ **Theorem 4** (Compactness). *A theory T is consistent if and only if every finite subset of T is consistent.*

Elementary extensions. Let $\mathbf{M}$ and $\mathbf{N}$ be two structures with $\mathbf{M} \subseteq \mathbf{N}$, that is, the domain of $\mathbf{M}$ is contained in the domain of $\mathbf{N}$. Then $\mathbf{N}$ is an *elementary extension* of $\mathbf{M}$, written $\mathbf{M} \prec \mathbf{N}$, if for every formula $\varphi(\bar{x})$ (without parameters) and tuple $\bar{m} \in \mathbf{M}^{\bar{x}}$, the following equivalence holds:

$$\mathbf{M} \models \varphi(\bar{m}) \qquad \text{if and only if} \qquad \mathbf{N} \models \varphi(\bar{m}).$$

We also say that $\mathbf{M}$ is an *elementary substructure* of $\mathbf{N}$. In other words, $\mathbf{M}$ is an elementary substructure of $\mathbf{N}$ if $\mathbf{M}$ is an induced substructure of $\mathbf{N}$, where we imagine that $\mathbf{M}$ and $\mathbf{N}$ are each equipped with every relation R_φ of arity k (for $k \in \mathbb{N}$) that is defined by any fixed first-order formula $\varphi(x_1,\ldots,x_k)$. In this intuition, formulas of arity 0 correspond to Boolean flags, with the same valuation for both $\mathbf{M}$ and $\mathbf{N}$.

Interpretations and transductions. A *simple interpretation* I between signatures Σ and Γ is specified by a domain formula $\delta(x)$ and a formula $\alpha_R(x_1,\ldots,x_k)$ for each relation symbol $R \in \Gamma$ of arity k, with δ and the α_R's being in the signature Σ. For a given Σ-structure $\mathbf{M}$, the interpretation outputs the Γ-structure $\mathsf{I}(\mathbf{M})$ whose domain is $\delta(\mathbf{M})$ and in which the interpretation of each relation R of arity k consists of the tuples $\bar{m}$ such that $\mathbf{M} \models \alpha_R(\bar{m})$. In this paper, we only consider simple interpretations, and therefore we will call them interpretations for conciseness.

For an integer $k \in \mathbb{N}$ and a structure $\mathbf{M}$, we define $k \times \mathbf{M}$ to be the structure consisting of k disjoint copies of $\mathbf{M}$, together with a new symmetric binary relation S containing all pairs (m,m') such that m and m' originate from the same element of $\mathbf{M}$. A *transduction* from Σ to Γ consists of an integer k, unary symbols $U_1,\ldots,U_\ell$ and an interpretation I from $\Sigma \cup \{S,U_1,\ldots,U_\ell\}$ to Γ.

For a transduction T and an input Σ-structure $\mathbf{M}$, the output $\mathsf{T}(\mathbf{M})$ consists of all Γ-structures $\mathbf{N}$ such that there exists a coloring $\widehat{\mathbf{M}}$ of $k \times \mathbf{M}$ with fresh unary predicates $U_1,\ldots,U_\ell$ such that $\mathbf{B} = \mathsf{I}(\widehat{\mathbf{M}})$. We say that a class of Σ-structures $\mathscr{C}$ *transduces* a class of Γ-structure $\mathcal{D}$ if there exists a transduction T such that for every structure $\mathbf{N} \in \mathscr{D}$ there is $\mathbf{M} \in \mathscr{C}$ satisfying $\mathbf{N} \in \mathsf{T}(\mathbf{M})$.

Graphs. We use standard graph theory notation. For a graph parameter π, we say that a graph class $\mathscr{C}$ has *bounded* π if there exists $k \in \mathbb{N}$ such that $\pi(G) \leqslant k$ for all $G \in \mathscr{C}$. Similarly, a class of structures $\mathscr{C}$ has bounded π if the class of Gaifman graphs of structures in $\mathscr{C}$ has bounded π.

3 Canonical Splitter-strategies in nowhere dense graphs

In this section, we show how compactness can be used to derive canonical decompositions for nowhere dense classes. More precisely, we will show that in the Splitter game, which characterizes nowhere dense classes [18], there is a constant k (depending only on the graph class $\mathscr{C}$) such that for any graph in $\mathscr{C}$ there are at most k optimal Splitter moves. This will allow us to illustrate the general methodology used in the paper.

Splitter game. First, we recall the rules of the Splitter game. The radius-r Splitter game is played on a graph G by two players, Splitter and Connector, in rounds $i = 1, 2, \ldots$ as follows. Initially the arena G_1 is the whole graph G. In the i-th round,

- Connector chooses a vertex $c_i \in G_i$;
- Splitter chooses a vertex $s_i \in G_i$ and we let $G_{i+1} = G_i[B^r_{G_i}(c_i)] - s_i$;
- Splitter wins if G_{i+1} is the empty graph, otherwise the game continues.

Here, $B^r_H(u) = \{v \in V(H) \mid \operatorname{dist}_H(u, v) \leqslant r\}$ denotes the ball of radius r around u in H.

The following result is instrumental in the celebrated proof of model-checking on nowhere dense classes [18].

▶ **Theorem 5** (Theorems 4.2 and 4.5 in [18]). *A class of graphs $\mathscr{C}$ is nowhere dense if and only if for every r, there exists ℓ such that on every graph $G \in \mathscr{C}$, Splitter can win the radius-r game in at most ℓ rounds.*

The *r-Splitter number* of a graph G is the minimal ℓ such that Splitter wins the radius-r game in ℓ rounds. Fix a nowhere dense class $\mathscr{C}$ and a radius r, and let ℓ be as in the theorem (hence ℓ is an upper bound to all r-Splitter numbers of graphs in $\mathscr{C}$). Observe that for a given $\ell' \leqslant \ell$ there is a first-order sentence expressing that Splitter wins the radius-r game in $\leqslant \ell'$ rounds, and therefore, there is a first-order sentence expressing that G has Splitter number ℓ'. Given a Connector move $c \in V(G)$, we say that a Splitter move $s \in V(G)$ is *r-progressing against c* if the r-Splitter number of $G[B_r(c)] - s$ is strictly smaller than the r-Splitter number of $G[B_r(c)]$. In other words, playing s is strictly better for Splitter than not playing any vertex. Again, since an upper bound to Splitter numbers depends only on $\mathscr{C}$, this can be expressed by a formula $\varphi_r(s; c)$. This leads to the following result.

▶ **Theorem 6.** *Let $\mathscr{C}$ be a nowhere dense class of graphs, and $r \in \mathbb{N}$. There is a constant k such that for every graph $G \in \mathscr{C}$, and every Connector move c, there are at most k progressing moves against c in G.*

In particular, this gives an isomorphism-invariant strategy for Splitter: simply play all progressing moves (either one by one, in any order, or all at once in an extended variant of the game considered in [18], where Splitter can remove a bounded number of vertices in each turn, instead of just one.) The idea of the proof is to extend, by compactness, progressive moves towards outside the model (in an elementary extension), and conclude by observing that "being a progressive move" is a definable and hereditary property.

Proof. Let T be the theory of $\mathscr{C}$. Note that T contains the sentence "Splitter wins the radius-r game in $\leqslant \ell$ rounds". Our aim is to prove that for some k, it contains the sentence "for all connector moves c, there are at most k progressing Splitter moves against c". We show that for any model of T and any connector move c, there are finitely many progressing Splitter moves against c; the result then follows from an easy application of compactness.

Assume for contradiction that there is a model $\mathbf{M}$ of T and a connector move $c \in \mathbf{M}$ such that Splitter has infinitely many progressing moves against c. We now let T' be the theory over the signature extended by a constant corresponding to each element $m \in \mathbf{M}$ and an additional constant s, such that T' consists of:

- all sentences in T,
- all sentences (with parameters in $\mathbf{M}$) which hold in $\mathbf{M}$,
- a sentence expressing that s is a progressing move against c, and
- for each $m \in \mathbf{M}$, the sentence $s \neq m$.

Since every finite subset T'' of T' mentions finitely many $m \in \mathbf{M}$, one can construct a model of T'' by starting from $\mathbf{M}$ and setting s to be one of those progressing moves that are not mentioned. We conclude from compactness (Theorem 4) that T' is consistent.

Let $\mathbf{N}$ be a model of T'. By construction $\mathbf{N}$ is an elementary extension of $\mathbf{M}$ – in particular, $\mathbf{N}[B_r(c)]$ has the same Splitter number ℓ' as $\mathbf{M}[B_r(c)]$ – and contains a progressing move $s \in \mathbf{N} - \mathbf{M}$ against c. This means that $\mathbf{N}[B_r(c)] - s$ has Splitter number $< \ell'$. But $\mathbf{M}[B_r(c)]$ is a subgraph of $\mathbf{N}[B_r(c)] - s$ with Splitter number ℓ': this is absurd. ◀

The next section presents more elaborate tools from stability theory that will allow us to extend the above idea to different settings.

4 Stability, forking, and Finitary Substitution

This section collects notions and a few basic results from stability theory. The purpose is to give a self-contained exposition culminating in our Finitary Substitution Lemma; for more context and explanations we refer to [27].

4.1 Stability and definability of types

We say that a formula $\varphi(\bar{x};\bar{y})$ defines a *ladder* of order k in a model $\mathbf{M}$ if there are sequences $\bar{a}_1, \dots \bar{a}_k \in \mathbf{M}^{\bar{x}}$ and $\bar{b}_1, \dots, \bar{b}_k \in \mathbf{M}^{\bar{y}}$ satisfying

$$\mathbf{M} \models \varphi(\bar{a}_i;\bar{b}_j) \qquad \text{if and only if} \qquad i < j, \qquad \text{for } 1 \leqslant i,j \leqslant k.$$

For a formula $\varphi(\bar{x};\bar{y})$ we call the largest k such that φ defines a ladder of order k the *ladder index* of φ in $\mathbf{M}$. If no such k exists, we say that the ladder index of φ is ∞.

We say that φ is *stable* in $\mathbf{M}$ if its ladder index is finite. We say that φ is stable in a theory T if it is stable in all models of T. Moreover, we say that a model (or a theory) is stable if every formula is stable.

We now state a fundamental result about stable formulas; it states that sets definable by stable formulas with parameters in some elementary extension can actually be defined from the model itself.

▶ **Theorem 7** (Definability of types). *Let $\mathbf{M} \prec \mathbf{N}$ be two models and $\varphi(\bar{x};\bar{y})$ be a stable formula of ladder index d in $\mathbf{M}$. For every $\bar{n} \in \mathbf{N}^{\bar{y}}$ there is a formula $\psi(\bar{x})$, which is a positive boolean combination of formulas of the form $\psi(\bar{x};\bar{m})$ using a tuple $\bar{m}$ of $2d+1$ parameters from $\mathbf{M}$, such that for every $\bar{a} \in \mathbf{M}^{\bar{x}}$,*

$$\mathbf{N} \models \varphi(\bar{a};\bar{n}) \quad \textit{if and only if} \quad \mathbf{M} \models \psi(\bar{a}).$$

4.2 Forking in stable theories

We move on to the definition of forking, which was first defined by Shelah in order to study stable theories [26], and later grew to become the central notion of stability theory. In stable theories, forking coincides with the simpler notion of dividing, so by a slight abuse we will only work with dividing (and call it forking). We first need to formally introduce types, then we give a definition of forking in stable theories and a few useful properties.

Types. Fix a model $\mathbf{M}$ over a signature Σ. A set π of formulas in variables $\bar{x}$ with parameters from $A \subseteq \mathbf{M}$ is called a *partial type over* A if it is *consistent*: for every finite subset $\pi' \subseteq \pi$ there is $\bar{m} \in \mathbf{M}^{\bar{x}}$ which satisfies all the formulas from π' (i.e. for every formula $\varphi(\bar{x}) \in \pi'$ we have $\mathbf{M} \models \varphi(\bar{m})$). We sometimes write $\pi(\bar{x})$ to explicitly mention free variables. Partial types p which are maximal are called *types*; this amounts to stating that for every formula $\varphi(\bar{x})$ with parameters from A, either $\varphi(\bar{x}) \in p$ or $\neg\varphi(\bar{x}) \in p$. Observe that for sets $A \subseteq B \subseteq \mathbf{M}$ every type p over A can be seen as a partial type over B. We denote the set of types over A in variables $\bar{x}$ by $S_{\bar{x}}(A)$.

For a tuple $\bar{a} \in \mathbf{M}^{\bar{x}}$ and a set $A \subseteq \mathbf{M}$ of parameters, the type of $\bar{a}$ over A, denoted $\mathrm{tp}(\bar{a}/A) \in S_{\bar{x}}(A)$, is the set of all formulas $\varphi(\bar{x})$ with parameters from A such that $\mathbf{M} \models \varphi(\bar{a})$. It follows from compactness that for every $p \in S_{\bar{x}}(\mathbf{M})$ there is some $\mathbf{N} \succ \mathbf{M}$ and an $\bar{x}$-tuple $\bar{n} \in \mathbf{N}^{\bar{x}}$ such that $\mathrm{tp}(\bar{n}/\mathbf{M}) = p$.

Forking. Fix a stable model $\mathbf{M}$ over a signature Σ and a set $A \subseteq \mathbf{M}$. Let $\varphi(\bar{x};\bar{y})$ be a formula without parameters and let $\bar{b} \in \mathbf{M}^{\bar{y}}$. We say that $\varphi(\bar{x};\bar{b})$ *forks over* A if there is an elementary extension $\mathbf{N} \succ \mathbf{M}$, a sequence $\bar{b}_1, \bar{b}_2, \ldots \in \mathbf{N}^{\bar{y}}$ satisfying $\mathrm{tp}(\bar{b}_i/A) = \mathrm{tp}(\bar{b}/A)$ for every i and an integer k such that $S = \{\varphi(\bar{x};\bar{b}_i) : i \in \mathbb{N}\}$ is k-inconsistent: no k-element subset of S is consistent. For a type $p \in S_{\bar{x}}(B)$ over a set $B \subseteq \mathbf{M}$, we say that p forks over A if there is a formula $\varphi(\bar{x};\bar{b}) \in p$ which forks over A.

We will make use of the following important property of forking which is often called *(full) existence*.

▶ **Theorem 8** (See [27, Corollary 7.2.7]). *Let $\mathbf{M}$ be a stable model and let $A \subseteq B \subseteq \mathbf{M}$. For every $p \in S_{\bar{x}}(A)$ there is some $q \in S_{\bar{x}}(B)$ such that $p \subseteq q$ and q does not fork over A.*

Finitary formulas. We say that a formula $\varphi(\bar{x};\bar{y})$ is *finitary* in a theory T if for every model $\mathbf{M}$ of T, the set $\{\varphi(\mathbf{M}^{\bar{x}};\bar{m}) : m \in \mathbf{M}^{\bar{y}}\}$ is finite. By compactness, this is equivalent to the following assertion: there exists $k \in \mathbb{N}$ such that $|\{\varphi(\mathbf{M}^{\bar{x}};\bar{m}) : m \in \mathbf{M}^{\bar{y}}\}| \leqslant k$ for every model $\mathbf{M}$ of T. We now relate forking and finitary formulas.

▶ **Theorem 9** (Special case of [27, Theorem 8.5.1][3]). *Let $\mathbf{M}$ be a stable model, $\mathbf{N}$ an elementary extension of $\mathbf{M}$, $\varphi(\bar{x};\bar{y})$ a formula, $\bar{n} \in \mathbf{N}^{\bar{y}}$, and $A \subseteq \mathbf{M}$. If $\mathrm{tp}(\bar{n}/\mathbf{M})$ does not fork over A, then there is a finitary formula $\varphi'(\bar{x};\bar{z})$ and a tuple $\bar{r} \in \mathbf{M}^{\bar{z}}$ such that $\varphi(\mathbf{N}^{\bar{x}};\bar{n}) \cap \mathbf{M}^{\bar{x}} = \varphi'(\mathbf{M}^{\bar{x}};\bar{r})$.*

Combining Theorems 8 and 9 yields the following statement.

▶ **Lemma 10.** *Let $\mathbf{M}$ be a stable model over the signature Σ, $\varphi(\bar{x};\bar{y})$ a Σ-formula, and ψ a sentence over signature $\Sigma \cup \{R\}$, where $R \notin \Sigma$ has arity $|\bar{x}|$. Let $\bar{s} \in \mathbf{M}^{\bar{y}}$ be such that $\mathbf{M} \models \psi[R(\bar{x})/\varphi(\bar{x};\bar{s})]$. Then there is an elementary extension $\mathbf{N}$ of $\mathbf{M}$, a tuple $\bar{s}' \in \mathbf{N}^{\bar{y}}$, a finitary formula $\varphi'(\bar{x};\bar{z})$ and a tuple $\bar{r} \in \mathbf{M}^{\bar{z}}$, such that $\mathbf{N} \models \psi[R(\bar{x})/\varphi(\bar{x};\bar{s}')]$ and $\varphi(\mathbf{N}^{\bar{x}};\bar{s}') \cap \mathbf{M}^{\bar{x}} = \varphi'(\mathbf{M}^{\bar{x}};\bar{r})$.*

Proof. Consider $p = \mathrm{tp}(\bar{s}/\emptyset)$. By Theorem 8, p extends to a type $q \in S_{\bar{y}}(\mathbf{M})$ which does not fork over $\emptyset$. By compactness there is an elementary extension $\mathbf{N} \succ \mathbf{M}$ and a tuple $\bar{s}' \in \mathbf{N}^{\bar{y}}$ such that $\mathrm{tp}(\bar{s}'/\mathbf{M}) = q$. In particular $\mathrm{tp}(\bar{s}'/\emptyset) = p = \mathrm{tp}(\bar{s}/\emptyset)$, and therefore $\mathbf{N} \models \psi[R(\bar{x})/\varphi(\bar{x};\bar{s}')]$ as required. Applying Theorem 9 we get a finitary formula $\varphi'(\bar{x};\bar{z})$ and a tuple $\bar{r} \in \mathbf{M}^{\bar{z}}$ with the wanted properties. ◀

[3] Formally, [27, Theorem 8.5.1] speaks about definability with *imaginaries*, which is known to be equivalent to the existence of finitary formulas (see for instance [9, Lemma 1.3.2 (5), Lemma 1.3.7]).

4.3 Finitary Substitute Lemma

Recall from Section 3 that applying our method requires a mechanism for moving the wanted property ψ back towards the structure $\mathbf{M}$ we started from. This is formalized by the following definition. In a theory T, and given two sentences ψ and ψ' over the signature $\Sigma \cup \{R\}$, we say that a sentence ψ *induces* ψ' *on semi-elementary substructures* if for every model $\mathbf{M}$ of T, for every elementary extension $\mathbf{N}$ and for every $\mathcal{R} \subseteq \mathbf{N}^k$, where k is the arity of R,

$$\mathbf{N}[R/\mathcal{R}] \models \psi \qquad \text{implies} \qquad \mathbf{M}[R/\mathcal{R}|_{\mathbf{M}}] \models \psi'.$$

As an important special case, if ψ is hereditary then ψ induces ψ on semi-elementary substructures. We are now ready to state our main model-theoretic tool.

▶ **Lemma 11** (Finitary Substitute Lemma). *Let T be a theory with signature Σ, $\varphi(\bar{x};\bar{y})$ a stable formula, and ψ, ψ' be sentences over the signature $\Sigma \cup \{R\}$, where $R \notin \Sigma$ has arity $|\bar{x}|$, such that ψ induces ψ' on semi-elementary substructures. Assume that $T \models \exists \bar{s}.\psi[R(\bar{x})/\varphi(\bar{x};\bar{s})]$. Then there is a finitary formula $\varphi'(\bar{x};\bar{z})$ such that $T \models \exists \bar{s}.\psi'[R(\bar{x})/\varphi'(\bar{x};\bar{s})]$.*

The proof follows from Lemma 10 by applying compactness; we refer to the full version for details.

5 Canonization of graphs of bounded shrubdepth

In this section, we prove Theorems 2 and 3 which we now recall for convenience.

▶ **Theorem 2.** *Let $\mathscr{C}$ be a class of graphs of bounded shrubdepth. Then there is a class $\mathscr{D}$ of binary structures of bounded treedepth and a mapping $\mathcal{A}\colon \mathscr{C} \to \mathscr{D}$ such that:*

- *For each $G \in \mathscr{C}$, the vertex set of G is contained in the domain of $\mathcal{A}(G)$ and the mapping $G \mapsto \mathcal{A}(G)$ is isomorphism-invariant.*
- *Given an n-vertex graph $G \in \mathscr{C}$, the structure $\mathcal{A}(G)$ has $\mathcal{O}(n)$ elements and can be computed in time $\mathcal{O}(n^2)$.*
- *There is a simple first-order interpretation I such that $G = \mathsf{I}(\mathcal{A}(G))$, for every $G \in \mathscr{C}$.*

▶ **Theorem 3.** *For every graph class $\mathscr{C}$ of bounded shrubdepth there is an $\mathcal{O}(n^2)$-time algorithm that given n-vertex graphs $G, G' \in \mathscr{C}$, decides whether G and G' are isomorphic.*

The proof is broken into three parts.

- The first part combines insights about classes of bounded shrubdepth with our Finitary Substitute Lemma developed in the previous section, to conclude that the first level in a shrubdepth decomposition (which we will call a dicing, defined below) can be defined using finitary formulas. This result is stated as Theorem 12 below.
- The second part builds on Theorem 12 to propose a canonical transformation from classes of bounded shrubdepth to classes of bounded treedepth. This proves Theorem 2.
- In the third part, we show how Theorem 3 is derived from Theorem 2, and also establish a stronger result about the canonization problem.

We start by recalling a few preliminaries about shrubdepth in Section 5.1, and proceed with the three parts outlined above in Sections 5.2, 5.3 and 5.4.

5.1 Preliminaries on shrubdepth

Shrubdepth. The decomposition notion underlying shrubdepth is that of *connection models*, defined as follows. Let G be a graph. A *connection model* for G consists of:

- a finite set of labels Labels;
- a labelling $\mathsf{label}\colon V(G) \to \mathsf{Labels}$;
- a rooted tree T whose leaf set coincides with the vertex set of G; and
- for every non-leaf node x of T, a symmetric relation $\mathsf{Adj}(x) \subseteq \mathsf{Labels} \times \mathsf{Labels}$, called the *adjacency table* at x.

The rule is as follows: for every distinct vertices u, v of G, u and v are adjacent in G if and only if $(\mathsf{label}(u), \mathsf{label}(v)) \in \mathsf{Adj}(x)$, where x is the lowest common ancestor of u and v in T.

The *depth* of a connection model is the depth of T. The *shrubdepth* of a graph class $\mathscr{C}$ is the least integer d with the following property: there exists a finite set of labels Labels such that every graph $G \in \mathscr{C}$ has a connection model of depth at most d that uses label set Labels.

Dicings. Our inductive proof requires manipulating the first level (just below the root) of a connection model; we will call this a dicing. Formally, for a graph G, a pair $(\mathcal{P}, \mathcal{L})$ of partitions of the vertex set of G is called a *dicing* of G if for every pair of vertices u, v belonging to different parts of $\mathcal{P}$, whether u and v are adjacent in G depends only on the pair of parts of $\mathcal{L}$ that u and v belong to. In other words, there is a symmetric relation $Z \subseteq \mathcal{L} \times \mathcal{L}$ such that for all u, v belonging to different parts of $\mathcal{P}$,

$$u \text{ and } v \text{ are adjacent in } G \qquad \text{if and only if} \qquad (\mathcal{L}(u), \mathcal{L}(v)) \in Z,$$

where $\mathcal{L}(w)$ denotes the part of $\mathcal{L}$ to which w belongs. In the context of a dicing $(\mathcal{P}, \mathcal{L})$, partition $\mathcal{P}$ will be called the *component partition*, and partition $\mathcal{L}$ will be called the *label partition*. The *order* of a dicing $(\mathcal{P}, \mathcal{L})$ is $|\mathcal{L}|$, the number of parts in the label partition.

Dicings appear naturally in connection models for shrubdepth: given a connection model for a graph G, using "having a common ancestor below the root" as component partition $\mathcal{P}$ and label-classes as label partition $\mathcal{L}$ defines a dicing of G.

5.2 Definability of canonical dicings

We say that a formula $\varphi(\bar{x}; \bar{y})$ with $|\bar{x}| = 2$ *defines a partition* if for every graph G and $\bar{b} \in G^{\bar{y}}$, $\varphi(G^{\bar{x}}; \bar{b})$ is an equivalence relation on the vertex set of G. (Note that for different choices of $\bar{b}$, φ can yield different equivalence relations.) Abusing the notation, by $\varphi(G^{\bar{x}}; \bar{b})$ we will also denote the partition of the vertex set into the equivalence classes of $\varphi(G^{\bar{x}}; \bar{b})$. Recall that a formula $\varphi(\bar{x}; \bar{y})$ is said to be finitary in (the theory of) a graph class $\mathscr{C}$ if there exists k such that for all graph $G \in \mathscr{C}$,

$$|\{\varphi(G^{\bar{x}}; \bar{b}) : \bar{b} \in G^{\bar{y}}\}| \leqslant k.$$

This section is focused on establishing the following result.

▶ **Theorem 12.** *Let $\mathscr{C}$ be a class of shrubdepth at most d, where $d > 1$. Then there exists a hereditary class $\mathscr{C}'$ of shrubdepth at most $d - 1$, finitary first-order formulas $\varphi(\bar{x}; \bar{y})$ and $\lambda(\bar{x}; \bar{y})$, each defining a partition, and $\ell \in \mathbb{N}$, such that the following holds: for every graph $G \in \mathscr{C}$ there exists $\bar{s} \in G^{\bar{y}}$ such that*

- *$(\varphi(G^{\bar{x}}; \bar{s}), \lambda(G^{\bar{x}}; \bar{s}))$ is a dicing of G of order at most ℓ; and*
- *for every part A of $\varphi(G^{\bar{x}}; \bar{s})$, we have $G[A] \in \mathscr{C}'$.*

On a high level, this proves that connection models can be defined using first-order formulas $\varphi(\bar{x};\bar{y}), \lambda(\bar{x};\bar{y})$ and parameters $\bar{s} \in G^{\bar{y}}$. While a good start towards sparsification, this alone would be insufficient for our needs, as different choices of $\bar{s}$ may lead to many different connection models, and choosing an arbitrary $\bar{s}$ would not give an isomorphism-invariant construction. This difficulty is overcome by the finitariness of φ and λ: our construction will take into account all of the (boundedly many) possible dicings (see Section 5.3).

The proof of Theorem 12 is broken into three parts as follows.

- The first part consists of proving that the label partition $\mathcal{L}$ can be chosen to be definable as a partition $\lambda(G^{\bar{x}};\bar{s})$ into $\bar{s}$-types. This is achieved thanks to a more general result of Bonnet et al. [3] pertaining to classes of bounded VC-dimension.
- We then show that the component partition $\mathcal{P}$ can be chosen to be definable by a formula $\varphi(G^{\bar{x}};\bar{s})$ using the same parameters $\bar{s}$. This relies on known properties of classes of bounded shrubdepth [16].
- We then apply our Finitary Substitute Lemma (Lemma 11) and prove that φ and λ can be taken to be finitary.

Definability of the label partition. For a subset of vertices S of a graph G we let $\mathcal{L}_S$ denote the partition of the vertex set of G into neighborhood classes with respect to S: u and v belong to the same part of $\mathcal{L}_S$ if and only if

$$\{w \in S \mid u \text{ and } w \text{ are adjacent}\} = \{w \in S \mid v \text{ and } w \text{ are adjacent}\}.$$

Note that we have $|\mathcal{L}_S| \leqslant 2^{|S|}$. It turns out that label partitions can be taken of this form.

▶ **Lemma 13** (follows from Theorem 3.5 of [3]). *Let $\mathscr{C}$ be graph class of bounded shrubdepth. Then for every graph $G \in \mathscr{C}$ and dicing $(\mathcal{P},\mathcal{L})$ of G of order at most t, there exists $S \subseteq V(G)$ with $|S| \leqslant \mathcal{O}(t^2)$ such that $(\mathcal{P},\mathcal{L}_S)$ is also a dicing of G.*

Definability of the component partition. We now show that the component partition $\mathcal{P}$ can also be defined using a first-order formula.

▶ **Lemma 14.** *Let $\mathscr{C}$ be a graph class of bounded shrubdepth and $t \in \mathbb{N}$ be an integer. There exist formulas $\varphi(\bar{x};\bar{y})$ and $\lambda(\bar{x};\bar{y})$, both defining a partition, such that the following holds: for every graph $G \in \mathscr{C}$ and dicing $(\mathcal{P},\mathcal{L})$ of G of order at most t, there exists $\bar{s} \in G^{\bar{y}}$ such that*

$$(\mathcal{P}',\mathcal{L}') = \big(\varphi(G^{\bar{x}};\bar{s}), \lambda(G^{\bar{x}};\bar{s})\big)$$

is a dicing of G of order at most $2^{\mathcal{O}(t^2)}$. Further, every part of $\mathcal{P}'$ is entirely contained in some part of $\mathcal{P}$.

Proof sketch. By Lemma 13, there exists a vertex subset S with $|S| \leqslant \mathcal{O}(t^2)$ such that $(\mathcal{P},\mathcal{L}_S)$ is also a dicing of G, with relation $Z \subseteq \mathcal{L}_S \times \mathcal{L}_S$. We let H denote the graph obtained by "flipping according to the dicing $(\mathcal{P},\mathcal{L}_S)$", meaning that we exchange edges and non-edges between pairs of parts in $\mathcal{L}_S$ that belong to Z. Since $(\mathcal{P},\mathcal{L}_S)$ is a dicing, connected components of H are contained in single parts of $\mathcal{P}$; let $\mathcal{P}'$ denote the partition of $V(G) = V(H)$ into connected components in H. Since H can be transduced from G, it has bounded shrubdepth, and thus we get that each part of $\mathcal{P}'$ have diameter bounded by a constant; this is because every class of bounded shrubdepth does not admit arbitrarily long induced paths [16]. We deduce that there is a formula expressing that two vertices belong to the same $\mathcal{P}'$-component, and the result follows. ◀

Finitariness of the definition. We are now ready to derive the theorem.

Proof sketch for Theorem 12. Let Labels be a large enough set of labels so that graphs in $\mathscr{C}$ admit connection models with labels in Labels, and let $\mathscr{C}'$ be the class of all graphs that admit a connection model of depth at most $d-1$ using the label set Labels. By Lemma 14, there exist formulas $\varphi(\bar{x};\bar{y})$ and $\lambda(\bar{x};\bar{y})$, depending only on $\mathscr{C}$, such that there is $\bar{s} \in G^{\bar{y}}$ for which $(\varphi(G^{\bar{x}};\bar{s}), \lambda(G^{\bar{x}};\bar{s}))$ is a dicing of G of order at most ℓ, where $\ell \in 2^{\mathcal{O}(|\mathsf{Labels}|^2)}$ is a constant depending only on $\mathscr{C}$. Moreover, every part of $\varphi(G^{\bar{x}};\bar{s})$ is entirely contained in a single part of $\mathcal{P}$, which implies that for every part A' of $\varphi(G^{\bar{x}},\bar{s})$ we have $G[A'] \in \mathscr{C}'$. It remains to transform φ and λ into finitary formulas. Let T be the theory of $\mathscr{C}$.

Let R be a relation symbol of arity 4 and consider the following assertion:

> "R is the product of two partitions $\mathcal{P}$ and $\mathcal{L}$ such that $(\mathcal{P},\mathcal{L})$ is a dicing of G of order at most ℓ. Moreover, for every part A of $\mathcal{P}$ it holds that $G[A] \in \mathscr{C}'$".

It follows from [16, Corollary 3.9] that the assertion above can be expressed by a first order sentence ψ over the signature $\{E, R\}$. Moreover, ψ is hereditary, so we may apply the Finitary Substitute Lemma to the formula $\eta(\bar{x}_1,\bar{x}_2;\bar{y}) = \varphi(\bar{x}_1;\bar{y}) \wedge \lambda(\bar{x}_2,\bar{y})$; we get a finitary $\eta'(\bar{x}_1,\bar{x}_2;\bar{y})$ such that

$$T \quad \text{implies} \quad \exists \bar{s}.\psi[R(\bar{x}_1,\bar{x}_2)/\eta'(\bar{x}_1,\bar{x}_2;\bar{s})].$$

Then the formulas

$$\varphi'(\bar{x};\bar{y}) = \exists z.\eta'(\bar{x},z,z;\bar{y}) \quad \text{and} \quad \lambda'(\bar{x};\bar{y}) = \exists z.\eta'(z,z,\bar{x};\bar{y})$$

yield the wanted result. ◀

5.3 Canonical reduction to bounded treedepth

With Theorem 12 in hand, we proceed to the proof of Theorem 2. Fix a class $\mathscr{C}$ of shrubdepth at most d.

Properties of the construction. We describe a construction that, given a graph $G \in \mathscr{C}$, constructs a structure $\mathcal{A}(G)$ of the following shape.

- $\mathcal{A}(G)$ is a structure over a signature consisting of several unary relations and one binary relation. Thus, we see $\mathcal{A}(G)$ as a vertex-colored directed graph, and we will apply the usual directed graphs terminology to $\mathcal{A}(G)$.
- The vertex set of G is contained in the vertex set of $\mathcal{A}(G)$. The elements of $V(G)$ will be called *leaves* of $\mathcal{A}(G)$. In $\mathcal{A}(G)$ there is a unary predicate marking all the leaves.
- In $\mathcal{A}(G)$ there is a specified vertex, called the *root*, such that for every vertex u of $\mathcal{A}(G)$ there is an arc from u to the root. The root is identified using a unary predicate.

The construction will satisfy the following properties.

- The mapping $G \mapsto \mathcal{A}(G)$ is isomorphism-invariant within the class $\mathscr{C}$.
- For every vertex u of $\mathcal{A}(G)$, there are at most c arcs with tail at u, for some constant c depending only on $\mathscr{C}$.
- There is a transduction T depending on $\mathscr{C}$ such that $\mathcal{A}(G) \in \mathsf{T}(G)$.
- The class $\{\mathcal{A}(G) \mid G \in \mathscr{C}\}$ has bounded treedepth.
- There is an interpretation I depending on $\mathscr{C}$ such that $G = \mathsf{I}(\mathcal{A}(G))$.
- Given G, $\mathcal{A}(G)$ can be computed in time $\mathcal{O}(n^2)$, where n is the vertex count of G.

We proceed by induction on d, the shrubdepth of $\mathscr{C}$; the base case $d = 1$ is obvious.

Preparation for the inductive construction. Suppose $d > 1$. Let $\varphi(\bar{x};\bar{y})$, $\lambda(\bar{x};\bar{y})$, $\ell \in \mathbb{N}$, and $\mathscr{C}'$ be the finitary formulas, the bound, and the class provided by Theorem 12. Since the shrubdepth of $\mathscr{C}'$ is at most $d-1$, by induction assumption we get a suitable mapping $\mathcal{A}'(\cdot)$, constant c', transduction T', and interpretation I' that satisfy the properties stated above for $\mathscr{C}'$.

Call a tuple $\bar{s} \in G^{\bar{y}}$ *good* if $(\varphi(G^{\bar{x}};\bar{s}), \lambda(G^{\bar{x}};\bar{s}))$ is a dicing of G of order at most ℓ satisfying that for every part A of $\varphi(G^{\bar{x}}, \bar{s})$ it holds that $G[A] \in \mathscr{C}'$. Define

$$\mathcal{F} = \{\, (\varphi(G^{\bar{x}};\bar{s}), \lambda(G^{\bar{x}};\bar{s})) \;:\; \bar{s} \in G^{\bar{y}} \text{ is a good tuple} \,\}.$$

By Theorem 12, we have

$$1 \leqslant |\mathcal{F}| \leqslant k$$

for some constant $k \in \mathbb{N}$ depending only on $\mathscr{C}$.

Let $\widehat{\mathcal{L}}$ be the coarsest partition that refines all label partitions of the dicings belonging to $\mathcal{F}$; that is, u, v are in the same part of $\widehat{\mathcal{L}}$ if and only if u, v are in the same part of $\mathcal{L}$ for each $(\mathcal{P}, \mathcal{L}) \in \mathcal{F}$. Similarly, let $\widehat{\mathcal{P}}$ be the coarsest partition that refines all component partitions of the dicings belonging to $\mathcal{F}$. Since $|\mathcal{F}| \leqslant k$ and $|\mathcal{L}| \leqslant \ell$ for each label partition featured in $\mathcal{F}$, we have

$$|\widehat{\mathcal{L}}| \leqslant \ell^k.$$

Moreover, every part of $\widehat{\mathcal{P}}$ is contained in a single part of any component partition featured in $\mathcal{F}$, hence $G[B] \in \mathscr{C}'$ for every part B of $\widehat{\mathcal{P}}$.

Let $\widehat{\mathcal{F}} = \{(\mathcal{P}, \widehat{\mathcal{L}}) \colon (\mathcal{P}, \mathcal{L}) \in \mathcal{F}\}$. Since $\widehat{\mathcal{L}}$ refines each label partition featured in $\mathcal{F}$, it follows that every element of $\widehat{\mathcal{F}}$ is a dicing of G. Then, for a component partition $\mathcal{P}$ featured in $\mathcal{F}$, let $Z_{\mathcal{P}} \subseteq \widehat{\mathcal{L}} \times \widehat{\mathcal{L}}$ be the symmetric relation witnessing that $(\mathcal{P}, \widehat{\mathcal{L}})$ is a dicing.

Definition of the construction. We now describe the structure $\mathcal{A}(G)$; see Figure 1. Construct:

- a root vertex r;
- for every part $L \in \widehat{\mathcal{L}}$, a vertex x_L;
- for every component partition $\mathcal{P}$ featured in $\mathcal{F}$, a vertex $y_{\mathcal{P}}$;
- for every component partition $\mathcal{P}$ featured in $\mathcal{F}$, and every part $A \in \mathcal{P}$, a vertex $z_{\mathcal{P},A}$; and
- for every component partition $\mathcal{P}$ featured in $\mathcal{F}$, and every (unordered) pair $LL' \in Z_{\mathcal{P}}$, a vertex $q_{\mathcal{P},LL'}$. (Note here that $Z_{\mathcal{P}}$ is symmetric, so we may treat its elements as unordered pairs of elements of $\widehat{\mathcal{L}}$.)

Further, for every part B of $\widehat{\mathcal{P}}$ we have $G[B] \in \mathscr{C}'$, hence we may apply the construction $\mathcal{A}'$ to $G[B]$, yielding a structure $H_B = \mathcal{A}'(G[B])$. Let r_B be the root of H_B. We add all structures H_B obtained in this way to $\mathcal{A}(G)$. We then connect these with the following arcs:

1. for every vertex u of $\mathcal{A}(G)$ there is an arc (u, r);
2. for every vertex of the form $z_{\mathcal{P},A}$ there is an arc $(z_{\mathcal{P},A}, y_{\mathcal{P}})$;
3. for every vertex of the form $q_{\mathcal{P},LL'}$, there are arcs $(q_{\mathcal{P},LL'}, x_L)$, $(q_{\mathcal{P},LL'}, x_{L'})$, and $(q_{\mathcal{P},LL'}, y_{\mathcal{P}})$;
4. for every part B of $\widehat{\mathcal{P}}$ and every component partition $\mathcal{P}$ featured in $\mathcal{F}$, there is an arc $(r_B, z_{\mathcal{P},A})$, where A is the unique part of $\mathcal{P}$ that contains B;
5. for every vertex u of G there is an arc (u, x_L), where L is the unique part of $\widehat{\mathcal{L}}$ that contains u. (Recall that the vertex set of G is the union of the leaf sets of H_B for $B \in \widehat{\mathcal{P}}$.)

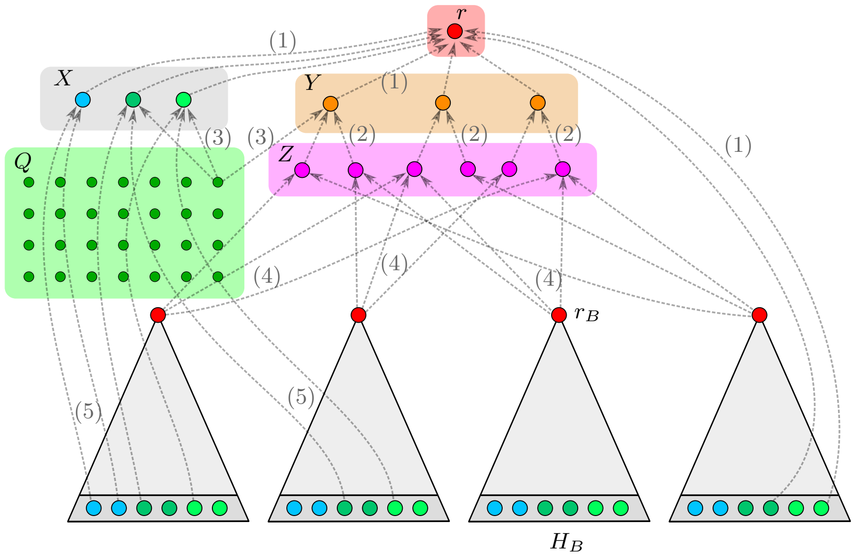

Figure 1 Inductive construction of $\mathcal{A}(G)$. Vertices of the form r, $y_{\mathcal{P}}$, $z_{\mathcal{P},A}$, $q_{\mathcal{P},LL'}$ are depicted in red, orange, violet, and green, respectively. Vertices of the form x_L are depicted in the top-left corner of the figure in different soft colors (which do not correspond to unary predicates), matching the colors of vertices of G that point to them; thus the soft color partition is $\widehat{\mathcal{L}}$. We depict a few representatives for each type of arcs.

Finally, we add five fresh unary predicates, called R, X, Y, Z, Q, respectively selecting the root r, the vertices of the form x_L, the vertices of the form $y_{\mathcal{P}}$, the vertices of the form $z_{\mathcal{P},A}$, and the vertices of the form $q_{\mathcal{P},LL'}$. Note here that H contains more unary predicates: those that come with structures H_B constructed by induction. These include a unary relation selecting the leaves.

This concludes the construction of $\mathcal{A}(G)$. We do not include detailed proofs of the properties listed above, and refer instead to the full version. That the transformation is isomorphism-invariant and that every element is the tail of a bounded number of arcs follows directly from the construction. Also, it is quite straightforward to transduce $\mathcal{A}(G)$ from G, by guessing good tuples $\bar{s}_1, \ldots, \bar{s}_{k'}$; then it follows from the results of Ganian et al. [16] that the class of $\mathscr{D} = \{\mathcal{A}(G) : G \in \mathscr{C}\}$ has bounded shrubdepth. This, together with the sparsity of $\mathcal{A}(G)$ following from the bound on outdegrees, implies that $\mathscr{D}$ in fact has bounded treedepth. Further, there is no difficulty in interpreting G in $\mathcal{A}(G)$. To compute $\mathcal{A}(G)$ in quadratic time, we rely on algorithmic meta-theorems over graphs of bounded cliquewidth obtained from combining [12, 21]. Theorem 2 follows.

5.4 Canonization and isomorphism test

We now use Theorem 2 to prove Theorem 3, that is, give a quadratic-time isomorphism test for any class of graphs of bounded shrubdepth. As mentioned in the introduction, in fact we solve the more general canonization problem, defined as follows.

For a class of structures $\mathscr{C}$, a *canonization map* for $\mathscr{C}$ is a mapping $\mathfrak{c}$ with the following property: for every $\mathbf{M} \in \mathscr{C}$, $\mathfrak{c}(\mathbf{M})$ is a total order on elements of $\mathbf{M}$ so that if $\mathbf{M}, \mathbf{M}' \in \mathscr{C}$ are isomorphic, then associating elements with same index in $\mathfrak{c}(\mathbf{M})$ and in $\mathfrak{c}(\mathbf{M}')$ yields an isomorphism between $\mathbf{M}$ and $\mathbf{M}'$. Note that if there is a canonization map $\mathfrak{c}$ for $\mathscr{C}$ that is efficiently computable, then this immediately gives an isomorphism test within the same time complexity.

For classes of bounded treedepth, Bouland et al. [4] gave a relatively easy fixed-parameter isomorphism test. Their techniques can be easily extended to the canonization problem for binary structures of bounded treedepth.

▶ **Theorem 15** (Adapted from [4])**.** *Let $\mathcal{D}$ be a class of binary structures of bounded treedepth. There exists a canonization map $\mathfrak{c}$ on $\mathcal{D}$ that is computable in time $\mathcal{O}(n \log^2 n)$, where n is the size of the universe of the input structure.*

We can now prove the main result of this section.

▶ **Theorem 16.** *Let $\mathscr{C}$ be a class of graphs of bounded shrubdepth. There exists a canonization map $\mathfrak{c}$ on $\mathscr{C}$ that is computable in time $\mathcal{O}(n^2)$, where n is the vertex count of the input graph.*

Proof. Let $\mathscr{D}$ be the class of bounded treedepth and $\mathcal{A}\colon \mathscr{C} \to \mathscr{D}$ be the mapping provided by Theorem 2 for the class $\mathscr{C}$. Then to get a suitable canonization map for $\mathscr{C}$, it suffices to compose $\mathcal{A}$ with the canonization map for $\mathscr{D}$, provided by Theorem 15, and restrict the output order to the vertex set of the original graph. ◀

As discussed, Theorem 3 follows immediately from Theorem 16.

References

1. Algorithms, Logic and Structure Workshop in Warwick – Open Problem Session. URL: `https://warwick.ac.uk/fac/sci/maths/people/staff/daniel_kral/alglogstr/openproblems.pdf`, 2016.
2. Hans Adler and Isolde Adler. Interpreting nowhere dense graph classes as a classical notion of model theory. *Eur. J. Comb.*, 36:322–330, 2014. doi:10.1016/j.ejc.2013.06.048.
3. Édouard Bonnet, Jan Dreier, Jakub Gajarský, Stephan Kreutzer, Nikolas Mählmann, Pierre Simon, and Szymon Torunczyk. Model Checking on Interpretations of Classes of Bounded Local Cliquewidth. In *LICS '22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 54:1–54:13. ACM, 2022. doi:10.1145/3531130.3533367.
4. Adam Bouland, Anuj Dawar, and Eryk Kopczyński. On Tractable Parameterizations of Graph Isomorphism. In *7th International Symposium on Parameterized and Exact Computation, IPEC 2012*, volume 7535 of *Lecture Notes in Computer Science*, pages 218–230. Springer, 2012. doi:10.1007/978-3-642-33293-7_21.
5. Samuel Braunfeld and Michael C. Laskowski. Characterizations of monadic NIP. *Transactions of the American Mathematical Society, Series B*, 8(30):948–970, 2021. doi:10.1090/btran/94.
6. Samuel Braunfeld and Michael C. Laskowski. Existential characterizations of monadic NIP. *CoRR*, abs/2209.05120, 2022. doi:10.48550/arXiv.2209.05120.
7. Samuel Braunfeld, Jaroslav Nešetřil, Patrice Ossona de Mendez, and Sebastian Siebertz. Decomposition horizons: from graph sparsity to model-theoretic dividing lines. *CoRR*, abs/2209.11229, 2022. doi:10.48550/arXiv.2209.11229.
8. Yijia Chen and Jörg Flum. FO-Definability of Shrub-Depth. In *28th EACSL Annual Conference on Computer Science Logic, CSL 2020*, volume 152 of *LIPIcs*, pages 15:1–15:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.CSL.2020.15.
9. Artem Chernikov. Lecture notes on stability theory. AMS Open Math Notes (OMN: 201905.110792), 2019.
10. Jan Dreier, Nikolas Mählmann, and Sebastian Siebertz. First-order model checking on structurally sparse graph classes. *CoRR*, abs/2302.03527, 2023. doi:10.48550/arXiv.2302.03527.
11. Jan Dreier, Nikolas Mählmann, Sebastian Siebertz, and Szymon Torunczyk. Indiscernibles and flatness in monadically stable and monadically NIP classes. *CoRR*, abs/2206.13765, 2022. doi:10.48550/arXiv.2206.13765.

12 Fedor V. Fomin and Tuukka Korhonen. Fast FPT-approximation of branchwidth. In *54th Annual ACM SIGACT Symposium on Theory of Computing, STOC '22*, pages 886–899. ACM, 2022. doi:10.1145/3519935.3519996.

13 Jakub Gajarský and Petr Hliněný. Faster Deciding MSO Properties of Trees of Fixed Height, and Some Consequences. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2012*, volume 18 of *LIPIcs*, pages 112–123. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2012. doi:10.4230/LIPIcs.FSTTCS.2012.112.

14 Jakub Gajarský, Stephan Kreutzer, Jaroslav Nešetřil, Patrice Ossona de Mendez, Michał Pilipczuk, Sebastian Siebertz, and Szymon Toruńczyk. First-Order Interpretations of Bounded Expansion Classes. *ACM Trans. Comput. Log.*, 21(4):29:1–29:41, 2020. doi:10.1145/3382093.

15 Jakub Gajarský, Nikolas Mählmann, Rose McCarty, Pierre Ohlmann, Michał Pilipczuk, Wojciech Przybyszewski, Sebastian Siebertz, Marek Sokolowski, and Szymon Torunczyk. Flipper games for monadically stable graph classes. *CoRR*, abs/2301.13735, 2023. doi:10.48550/arXiv.2301.13735.

16 Robert Ganian, Petr Hliněný, Jaroslav Nešetřil, Jan Obdržálek, and Patrice Ossona de Mendez. Shrub-depth: Capturing Height of Dense Graphs. *Log. Methods Comput. Sci.*, 15(1), 2019. doi:10.23638/LMCS-15(1:7)2019.

17 Robert Ganian, Petr Hliněný, Jaroslav Nešetřil, Jan Obdržálek, Patrice Ossona de Mendez, and Reshma Ramadurai. When Trees Grow Low: Shrubs and Fast MSO_1. In *37th International Symposium on Mathematical Foundations of Computer Science 2012, MFCS 2012*, volume 7464 of *Lecture Notes in Computer Science*, pages 419–430. Springer, 2012. doi:10.1007/978-3-642-32589-2_38.

18 Martin Grohe, Stephan Kreutzer, and Sebastian Siebertz. Deciding First-Order Properties of Nowhere Dense Graphs. *J. ACM*, 64(3):17:1–17:32, 2017. doi:10.1145/3051095.

19 Martin Grohe and Daniel Neuen. Isomorphism, canonization, and definability for graphs of bounded rank width. *Commun. ACM*, 64(5):98–105, 2021. doi:10.1145/3453943.

20 Martin Grohe and Pascal Schweitzer. Isomorphism Testing for Graphs of Bounded Rank Width. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015*, pages 1010–1029. IEEE Computer Society, 2015. doi:10.1109/FOCS.2015.66.

21 Wojciech Kazana and Luc Segoufin. Enumeration of monadic second-order queries on trees. *ACM Trans. Comput. Log.*, 14(4):25:1–25:12, 2013. doi:10.1145/2528928.

22 O-joung Kwon, Rose McCarty, Sang-il Oum, and Paul Wollan. Obstructions for bounded shrub-depth and rank-depth. *J. Comb. Theory, Ser. B*, 149:76–91, 2021. doi:10.1016/j.jctb.2021.01.005.

23 Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity: Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and Combinatorics*. Springer, 2012. doi:10.1007/978-3-642-27875-4.

24 A. Pillay. *Geometric Stability Theory*. Oxford logic guides. Clarendon Press, 1996. URL: https://books.google.pl/books?id=k6FK_Gqal2EC.

25 Klaus-Peter Podewski and Martin Ziegler. Stable graphs. *Fundamenta Mathematicae*, 100(2):101–107, 1978. doi:10.4064/fm-100-2-101-107.

26 Saharon Shelah. *Classification theory – and the number of non-isomorphic models*, volume 92 of *Studies in logic and the foundations of mathematics*. North-Holland Publishing Co., 1978.

27 Katrin Tent and Martin Ziegler. *A Course in Model Theory*. Lecture Notes in Logic. Cambridge University Press, 2012. doi:10.1017/CBO9781139015417.