PRESERVATION AND DECOMPOSITION THEOREMS
FOR BOUNDED DEGREE STRUCTURES [*]

FREDERIK HARWATH [a], LUCAS HEIMBERG [b], AND NICOLE SCHWEIKARDT [c]

[a] Institut für Informatik, Goethe-Universität Frankfurt am Main
    e-mail address: harwath@cs.uni-frankfurt.de

[b,c] Institut für Informatik, Humboldt-Universität zu Berlin
    e-mail address: {lucas.heimberg, schweika}@informatik.hu-berlin.de

ABSTRACT. We provide elementary algorithms for two preservation theorems for first-order sentences (FO) on the class $\mathfrak{C}_d$ of all finite structures of degree at most d: For each FO-sentence that is preserved under extensions (homomorphisms) on $\mathfrak{C}_d$, a $\mathfrak{C}_d$-equivalent existential (existential-positive) FO-sentence can be constructed in 5-fold (4-fold) exponential time. This is complemented by lower bounds showing that a 3-fold exponential blow-up of the computed existential (existential-positive) sentence is unavoidable. Both algorithms can be extended (while maintaining the upper and lower bounds on their time complexity) to input first-order sentences with modulo m counting quantifiers (FO+MOD$_m$).

Furthermore, we show that for an input FO-formula, a $\mathfrak{C}_d$-equivalent Feferman-Vaught decomposition can be computed in 3-fold exponential time. We also provide a matching lower bound.

1. INTRODUCTION

Classical *preservation theorems* studied in model theory relate syntactic restrictions of formulas with structural properties of the classes of structures defined. For example, the *Loś-Tarski theorem* states that a first-order sentence is preserved under extensions on the class of all structures if, and only if, it is equivalent, on this class, to an existential first-order sentence. The *homomorphism preservation theorem* states that a first-order sentence is preserved under homomorphisms on the class of all structures if, and only if, it is equivalent, on this class, to an existential-positive first-order sentence.

In the last decade, variants of both theorems have been obtained, where the class of all structures is replaced by restricted classes that meet certain requirements. For example, [Ros08, ADK06, Daw10] obtained that the homomorphism preservation theorem holds for the class of all finite structures, as well as for the classes of all finite structures of degree at most d or of treewidth at most k, and, in general, for *quasi-wide* classes of structures that

are closed under taking substructures and disjoint unions (this includes classes of bounded expansion and classes that locally exclude minors). While the Łoś-Tarski theorem is known to fail on the class of all finite structures, in [ADG08] it was shown to hold for various classes of structures, including the class of all finite structures of degree at most d, the class of all finite structures of treewidth at most k, and all *wide* classes of structures that are closed under taking substructures and disjoint unions.

For most of these results, it is known that the equivalent existential or existential-positive sentence may be non-elementarily larger than the corresponding first-order sentence [DGKS07]. A notable exception affects the Łoś-Tarski theorem for the class of acyclic finite structures of degree at most d, for which [DGKS07] obtained a 5-fold exponential upper bound on the size of the equivalent existential first-order sentence.

The present paper's first main result (Theorem 3.1) generalises the latter in three ways: (1) We show that the 5-fold exponential upper bound for the Łoś-Tarski theorem holds for *every* class $\mathfrak{C}$ of structures of degree at most d that is closed under taking induced substructures and disjoint unions (this includes, e.g., the class of all finite structures of degree at most d). (2) We provide an *algorithmic* version of the theorem, showing that for a given first-order sentence, the existential sentence can be constructed in 5-fold exponential time. (3) Our algorithm also works for input sentences of the extension FO+MOD$_m$ of first-order logic with modulo m counting quantifiers. The main ingredient of our proof is a new, technically challenging upper bound on the size of *minimal* models of sentences that are preserved under extensions on $\mathfrak{C}$ (Theorem 3.5).

Our second main result (Theorem 3.2) provides an algorithmic version of the homomorphism preservation theorem over any class $\mathfrak{C}$ of structures of bounded degree that is closed under taking induced substructures and disjoint unions, and that is decidable in 1-fold exponential time (e.g., the class of all finite structures of degree at most d). Specifically, we show that for a given FO+MOD$_m$-sentence that is preserved under homomorphisms on $\mathfrak{C}$, an equivalent existential-positive first-order sentence can be constructed in 4-fold exponential time. The proof, again, relies on a new upper bound on the size of minimal models (Theorem 3.12).

Two counterexamples (Theorem 3.14 and Theorem 3.15) show that the closure properties of the classes of structures considered, required by our preservation theorems, are indeed necessary.

We complement our preservation theorems by lower bounds (Theorem 5.4 and Theorem 5.3), providing a sequence of first-order sentences that are preserved under extensions (homomorphisms) for which the smallest equivalent existential (respectively, existential-positive) sentences are 3-fold exponentially larger. Both lower bound proofs use particular encodings of numbers by binary trees introduced in [HKS13].

Our third main result deals with *Feferman-Vaught decompositions* of first-order formulas. The classical Feferman-Vaught theorem states that for certain forms of compositions of structures, the theory of a structure composed from simpler structures is determined by the theories of the simpler structures. This applies, for example, to disjoint sums and direct products (also known as cartesian products or as tensor products) of structures (cf., e.g., [Hod93]). Feferman-Vaught-like theorems find application in results about the *decidability* of theories, as well as in results about *model checking* and *satisfiability checking* [Mak04, GJL15]. Another use of Feferman-Vaught decompositions is within the proof of *Gaifman's theorem* [Gai82], which is an important tool for inexpressibility results as well as for so-called algorithmic meta-theorems [Kre11].

Algorithmic versions of decomposition theorems à la Feferman-Vaught are typically of the following form (cf., [Mak04, GJL15]): A given first-order sentence φ that shall be evaluated in the disjoint sum or the direct product $\mathcal{A}$ of s structures $\mathcal{A}_1, \ldots, \mathcal{A}_s$, can be transformed into a finite set Δ of formulas and a propositional formula β whose propositions are tests of the form "the i-th structure $\mathcal{A}_i$ satisfies the j-th formula in Δ", such that $\mathcal{A}$ is a model of φ iff β is true. It is known that the Feferman-Vaught decomposition (Δ, β) may be non-elementarily larger than φ [DGKS07]. Our third main result (Theorem 4.2 and Corollary 4.9) shows that for any class $\mathfrak{C}$ of structures of degree at most d, such Feferman-Vaught decompositions for disjoint sums and direct products can be computed in 3-fold exponential time. This is complemented by a matching lower bound (Theorem 5.8). Our lower bound proof, again, relies on encodings of numbers by binary trees, now along with a method of [GJL15]. Our algorithm produces a set Δ of so-called *Hanf-formulas* and relies on a result of [BK12] that transforms the given sentence φ into Hanf normal form.

The rest of this paper is structured as follows: Section 2 fixes the basic notations. Section 3 presents our algorithms concerning preservation theorems. Section 4 presents our results concerning Feferman-Vaught decompositions. Section 5 contains lower bounds that complement our results of Section 3 and Section 4. Section 6 gives a short conclusion and provides directions for future work.

2. Preliminaries

We write $\mathbb{Z}$ for the set of integers and $\mathbb{N}$ for the set of non-negative integers. For all $m, n \in \mathbb{N}$ with $m \leqslant n$, we denote the set $\{i \in \mathbb{N} : m \leqslant i \leqslant n\}$ by $[m, n]$. The set of non-negative real numbers is denoted by $\mathbb{R}_{\geqslant 0}$. For $r > 0$, by $\log(r)$ (respectively, $\log_b(r)$) we denote the logarithm of r with respect to base 2 (respectively, base b, for $b \geqslant 2$).

If f is a function from $\mathbb{N}$ to $\mathbb{R}_{\geqslant 0}$, then $\operatorname{poly}(f(n))$ denotes the class of all functions $g \colon \mathbb{N} \to \mathbb{R}_{\geqslant 0}$ for which there is a number $c > 0$ such that $g(n) \leqslant (f(n))^c$ is true for all sufficiently large $n \in \mathbb{N}$.

We say that a function f from $\mathbb{N}$ to $\mathbb{R}_{\geqslant 0}$ is *at most k-fold exponential*, for some $k \geqslant 1$, if there exists a number $c > 0$ such that for all sufficiently large $n \in \mathbb{N}$ we have $f(n) \leqslant T(k, n^c)$, where $T(k, m)$ is a tower of 2s of height k with an m on top (i.e., $T(1, m) = 2^m$ and $T(k+1, m) = 2^{T(k,m)}$ for all $k \geqslant 1$ and $m \geqslant 0$).

For a function $f \colon A \to B$ and a subset $A' \subseteq A$, we denote by $f(A')$ the set $\{f(a) : a \in A'\}$. For $n \in \mathbb{N}$, we write $\overline{x}$ to denote the tuple $(x_1, \ldots, x_n)$. For a tuple $\overline{x} \in A^n$ we write $f(\overline{x})$ to denote the tuple $(f(x_1), \ldots, f(x_n))$. Sometimes we treat tuples as if they were sets; e.g., $a \in \overline{x}$ means $a \in \{x_1, \ldots, x_n\}$.

2.1. **Structures and formulas.** A *signature* σ is a finite set of relation and constant symbols. Associated with every relation symbol R is a positive integer $\operatorname{ar}(R)$ called the *arity* of R. The *size* $\|\sigma\|$ of σ is the number of its constant symbols plus the sum of the arities of its relation symbols. A signature σ is called *relational* if it does not contain any constant symbol.

A σ-*structure* $\mathcal{A}$ consists of a non-empty set A called the *universe* of $\mathcal{A}$, a relation $R^{\mathcal{A}} \subseteq A^{\operatorname{ar}(R)}$ for each relation symbol $R \in \sigma$, and an element $c^{\mathcal{A}} \in A$ for each constant symbol $c \in \sigma$. The size $\|\mathcal{A}\|$ of $\mathcal{A}$ is the size of a reasonable representation of $\mathcal{A}$ as a binary string (cf., e.g., [EF99, Lib04]); in particular, $\|\mathcal{A}\| \in O(|A|^{\|\sigma\|})$.

For a relational signature σ and σ-structures $\mathcal{A}$ and $\mathcal{B}$, we say that $\mathcal{B}$ is a *substructure* of $\mathcal{A}$ if $B \subseteq A$ and $R^{\mathcal{B}} \subseteq R^{\mathcal{A}}$ for each $R \in \sigma$. The structure $\mathcal{B}$ is an *induced substructure* of $\mathcal{A}$ if $\mathcal{B}$ is a substructure of $\mathcal{A}$ and $R^{\mathcal{B}} = R^{\mathcal{A}} \cap B^{\mathrm{ar}(R)}$ for each $R \in \sigma$. We then say that $\mathcal{B}$ is the substructure of $\mathcal{A}$ *induced* by the set B.

For every non-empty set B such that $A \cap B \neq \emptyset$, we write $\mathcal{A}[B]$ to denote the substructure of $\mathcal{A}$ induced by $A \cap B$. Furthermore, if $A \setminus B \neq \emptyset$ then $\mathcal{A} - B$ is the induced substructure $\mathcal{A}[A \setminus B]$ of $\mathcal{A}$ obtained by deleting all elements from B.

We use the standard notation concerning first-order logic and extensions thereof, cf. [EF99, Lib04]. By $\mathrm{qr}(\varphi)$ we denote the *quantifier rank* of φ, i.e., the maximum nesting depth of quantifiers occurring in a formula φ. By $\mathrm{free}(\varphi)$ we denote the set of all free variables of φ. A *sentence* is a formula φ with $\mathrm{free}(\varphi) = \emptyset$. A σ-structure $\mathcal{A}$ is called a *model* of a sentence φ if φ is satisfied in $\mathcal{A}$.

We write $\varphi(\overline{x})$, for $\overline{x} = (x_1, \ldots, x_n)$ with $n \geqslant 0$, to indicate that $\mathrm{free}(\varphi) \subseteq \{x_1, \ldots, x_n\}$. If $\mathcal{A}$ is a σ-structure and $\overline{a} = (a_1, \ldots, a_n) \in A^n$, we write $\mathcal{A} \models \varphi[\overline{a}]$ to indicate that the formula $\varphi(\overline{x})$ is satisfied in $\mathcal{A}$ when interpreting the free occurrences of the variables $x_1, \ldots, x_n$ with the elements $a_1, \ldots, a_n$. We write $\varphi(\mathcal{A})$ to denote the set of all tuples $\overline{a} \in A^n$ such that $\mathcal{A} \models \varphi[\overline{a}]$. For a class $\mathfrak{C}$ of structures, two formulas $\varphi(\overline{x})$ and $\psi(\overline{x})$ of signature σ are called *equivalent on $\mathfrak{C}$* (for short: *$\mathfrak{C}$-equivalent*) if for all σ-structures $\mathcal{A} \in \mathfrak{C}$ we have $\varphi(\mathcal{A}) = \psi(\mathcal{A})$.

By $\mathrm{FO}(\sigma)$ we denote the class of all first-order formulas of signature σ. The extension of $\mathrm{FO}(\sigma)$ by *modulo counting quantifiers* is defined as follows: Let m be an integer such that $m \geqslant 2$. We write $\exists^{0 \bmod m}$ to denote the *modulo m counting quantifier*. A formula of the form $\exists^{0 \bmod m} y\, \psi(\overline{x}, y)$ is satisfied by a σ-structure $\mathcal{A}$ and an interpretation $\overline{a}$ of the variables $\overline{x}$ if, and only if, the number of elements $b \in A$ such that $\mathcal{A} \models \psi[\overline{a}, b]$ is a multiple of m. For a fixed number m we write $\mathrm{FO}+\mathrm{MOD}_m(\sigma)$ to denote the extension of $\mathrm{FO}(\sigma)$ with modulo m counting quantifiers. The quantifier rank $\mathrm{qr}(\varphi)$ of an $\mathrm{FO}+\mathrm{MOD}_m$-formula φ is defined as the maximum nesting depth of *all* quantifiers (i.e., first-order quantifiers and modulo counting quantifiers).

The size $\|\varphi\|$ of an $\mathrm{FO}+\mathrm{MOD}_m(\sigma)$-formula φ is its length when viewed as a word over the alphabet $\sigma \cup \{=\} \cup \{\exists^{0 \bmod m}, \exists, \forall, \neg, \wedge, \vee, \rightarrow, \leftrightarrow, (,)\} \cup \{,\} \cup \mathit{Var}$, where Var is a countable set of variable symbols.

2.2. **Gaifman graph.** For a σ-structure $\mathcal{A}$, we write $G_{\mathcal{A}}$ to denote the Gaifman graph of $\mathcal{A}$, i.e., the undirected, loop-free graph with vertex set A and an edge between two distinct vertices $a, b \in A$ iff there exists an $R \in \sigma$ and a tuple $(a_1, \ldots, a_{\mathrm{ar}(R)}) \in R^{\mathcal{A}}$ such that $a, b \in \{a_1, \ldots, a_{\mathrm{ar}(R)}\}$.

Given a σ-structure $\mathcal{A}$ and two elements $a, b \in A$ that are connected in the Gaifman graph $G_{\mathcal{A}}$, the *distance $\mathit{dist}^{\mathcal{A}}(a, b)$* between a and b is the minimal length (i.e., the number of edges) of a path from a to b in $G_{\mathcal{A}}$. For $a, b \in A$ that are not connected in $G_{\mathcal{A}}$ we let $\mathit{dist}^{\mathcal{A}}(a, b) := \infty$.

For $r \geqslant 0$ and $a \in A$, the *r-neighbourhood* of a in $\mathcal{A}$ is the set
$$N_r^{\mathcal{A}}(a) \ := \ \{b \in A : \mathit{dist}^{\mathcal{A}}(a, b) \leqslant r\}.$$

The r-neighbourhood $N_r^{\mathcal{A}}(W)$ of a set $W \subseteq A$ is the union of the r-neighbourhoods $N_r^{\mathcal{A}}(a)$ for all $a \in W$. For a tuple $\overline{a} = (a_1, \ldots, a_n)$, we write $N_r^{\mathcal{A}}(\overline{a})$ instead of $N_r^{\mathcal{A}}(\{a_1, \ldots, a_n\})$.

2.3. **Bounded structures.** The *degree* of a σ-structure $\mathcal{A}$ is the degree of its Gaifman graph $G_\mathcal{A}$. Let $\nu \colon \mathbb{N} \to \mathbb{N}$ be a function. A σ-structure $\mathcal{A}$ is ν-*bounded* if $|N_r^\mathcal{A}(a)| \leq \nu(r)$ for all $r \geqslant 0$ and all $a \in A$. Clearly, if $\mathcal{A}$ is ν-bounded, then it has degree at most $\nu(1)-1$. On the other hand, if $\mathcal{A}$ has degree at most d, then $\mathcal{A}$ is ν_d-bounded for $\nu_d \colon \mathbb{N} \to \mathbb{N}$ with $\nu_d(r) = 1 + d \cdot \sum_{i=0}^{r-1}(d-1)^i$. Thus, $\mathcal{A}$ has degree at most d iff it is ν_d-bounded. Note that ν_d is at most 1-fold exponential.

We will restrict attention to at most 1-fold exponential functions $\nu \colon \mathbb{N} \to \mathbb{N}$ that are *strictly increasing*. This is reasonable, since then $(r+1)$-neighbourhoods may contain more elements than r-neighbourhoods, and it excludes pathological cases where ν-boundedness of a structure implies that the structure is a disjoint union of finite structures whose size is bounded by a constant depending on ν.

Let σ be a finite relational signature, let $\mathcal{A}$ be a σ-structure, let $n \geqslant 1$, and let $\overline{a} \in A^n$. Let $c_1, \ldots, c_n$ be distinct constant symbols.

For $r \geqslant 0$, the *r-sphere around $\overline{a}$* is defined as the $\sigma \cup \{c_1, \ldots, c_n\}$-structure $\mathcal{N}_r^\mathcal{A}(\overline{a}) := \big(\mathcal{A}[N_r^\mathcal{A}(\overline{a})], \overline{a} \big)$, where the constant symbols $c_1, \ldots, c_n$ are interpreted by the elements $a_1, \ldots, a_n$.

An *r-sphere with n centres* is a $\sigma \cup \{c_1, \ldots, c_n\}$-structure $\tau = (\mathcal{B}, \overline{b})$ with $\overline{b} \in B^n$ and universe $B = N_r^\mathcal{B}(\overline{b})$. We say that *$\tau$ is realised by $\overline{a}$ in $\mathcal{A}$* iff $\mathcal{N}_r^\mathcal{A}(\overline{a})$ is isomorphic to τ. By $\tau(\mathcal{A})$ we denote the set of all $\overline{a} \in A^n$ that realise τ in $\mathcal{A}$.

Note that a ν-bounded r-sphere τ with n centres contains at most $n \cdot \nu(r)$ elements. Thus, there is an FO(σ)-formula $sph_\tau(\overline{x})$ of size $(n \cdot \nu(r))^{O(\|\sigma\|)}$ such that for all σ-structures $\mathcal{A}$ we have $sph_\tau(\mathcal{A}) = \tau(\mathcal{A})$.

Unless otherwise indicated, we assume an r-sphere to have only one centre. Up to isomorphism, the number of σ-structures with exactly n elements is[1] at most $2^{n^a} \cdot n^c$, where a is the sum of the arities of the relation symbols in σ and c is the number of constant symbols in σ. Hence we can bound the number of non-isomorphic r-spheres that can be realised in ν-bounded σ-structures from above by

$$2^{\nu(r)^{\|\sigma\|+1}}.$$

2.4. **Disjoint unions.** For a relational signature σ and each $s \geqslant 1$, the *disjoint union* $\mathcal{A}_1 \mathbin{\dot\cup} \cdots \mathbin{\dot\cup} \mathcal{A}_s$ of σ-structures $\mathcal{A}_1, \ldots, \mathcal{A}_s$ is a structure $\mathcal{A}$ that is defined (up to isomorphism) as follows: Let A be a set of size $|A_1| + \cdots + |A_s|$ and let, for each $i \in [1, s]$, $f_i \colon A_i \to A$ be an injective function such that $f_1(A_1), \ldots, f_s(A_s)$ is a partition of A. Now, $\mathcal{A}$ is the σ-structure with universe A where, for each $R \in \sigma$, the relation $R^\mathcal{A}$ is the union of the sets $\{ f_i(\overline{a}) : \overline{a} \in R^{\mathcal{A}_i} \}$ for all $i \in [1, s]$. The *mapping of $\mathcal{A}$* is the function $\pi \colon A \to (A_1 \cup \cdots \cup A_s)$ with $\pi(a) = b$ where b is chosen such that $b \in A_i$ for some $i \in [1, s]$ and $f_i(b) = a$.

If the universes of $\mathcal{A}_1, \ldots, \mathcal{A}_s$ are pairwise disjoint, we let $A := A_1 \cup \cdots \cup A_s$ and we let f_i be the identity on A_i for each $i \in [1, s]$, and π the identity on A.

3. PRESERVATION THEOREMS

Throughout this section, σ will always denote a finite relational signature and $\mathfrak{C}$ will always denote a class of σ-structures. A σ-structure $\mathcal{B}$ is an *extension* of a σ-structure $\mathcal{A}$ if $\mathcal{A}$ is an *induced* substructure of $\mathcal{B}$. A sentence φ is *preserved under extensions on $\mathfrak{C}$* if for each

---

[1] We will usually omit brackets and shortly write 2^{n^a} for $2^{(n^a)}$.

model $\mathcal{A} \in \mathfrak{C}$ of φ and every extension $\mathcal{B} \in \mathfrak{C}$ of $\mathcal{A}$, $\mathcal{B}$ is also a model of φ. An *existential* FO(σ)-*formula* has the form $\exists x_1 \cdots \exists x_n \varphi$, where φ is quantifier-free. It is straightforward to see that every existential FO(σ)-sentence is preserved under extensions on arbitrary classes of σ-structures.

A *homomorphism* of σ-structures $\mathcal{A}$ and $\mathcal{B}$ is a mapping $h \colon A \to B$ such that for each relation symbol $R \in \sigma$ with $r := \mathrm{ar}(R)$ and all tuples $(a_1, \ldots, a_r) \in A^r$, if $(a_1, \ldots, a_r) \in R^{\mathcal{A}}$ then $(h(a_1), \ldots, h(a_r)) \in R^{\mathcal{B}}$. A sentence φ is *preserved under homomorphisms on* $\mathfrak{C}$ if for each model $\mathcal{A} \in \mathfrak{C}$ of φ and each structure $\mathcal{B} \in \mathfrak{C}$ for which a homomorphism h from $\mathcal{A}$ to $\mathcal{B}$ exists, $\mathcal{B}$ is also a model of φ. An *existential-positive* FO(σ)-*formula* is an existential FO(σ)-formula that does not contain any of the symbols $\neg$, $\to$, $\leftrightarrow$. For convenience, we will say that also *false* is an existential-positive FO(σ)-sentence (that is not satisfied by any σ-structure). It is straightforward to see that every existential-positive FO(σ)-sentence is preserved under homomorphisms on arbitrary classes of σ-structures.

For the remainder of this section, let $\nu \colon \mathbb{N} \to \mathbb{N}$ be a fixed time-constructible strictly increasing function that is at most 1-fold exponential. Recall that the number of non-isomorphic r-spheres (with one centre) that can be realised in ν-bounded σ-structures is bounded from above by $2^{\nu(r)^{O(\|\sigma\|)}}$. Henceforth, we will abbreviate this expression by $S_\nu(r, \|\sigma\|)$, i.e.,

$$ S_\nu(r, \|\sigma\|) \quad := \quad 2^{\nu(r)^{O(\|\sigma\|)}}. $$

Thus, $S_\nu(\cdot, \cdot)$ is the class of all functions $f \colon \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ for which there exists a $c > 0$ such that $f(r, s) \leqslant 2^{\nu(r)^{c \cdot s}}$ for all sufficiently large $r, s \geqslant 1$.

In this section, we explore the complexity of constructing existential (respectively, existential-positive) FO(σ)-sentences for FO+MOD$_m$(σ)-sentences that are preserved under extensions (respectively, homomorphisms) on classes of ν-bounded σ-structures that are closed under disjoint unions and closed under induced substructures (respectively, closed under disjoint unions, closed under induced substructures, and decidable in 1-fold exponential time). It is straightforward to see that the class $\mathfrak{C}_d$ of all finite σ-structures of degree at most d, for any fixed $d \geqslant 0$, meets all these requirements. Similarly, as we assume that ν is time-constructible and at most 1-fold exponential, also the class of all ν-bounded structures is easily seen to be decidable in 1-fold exponential time, as well as closed under taking induced substructures and disjoint unions.

In Subsection 3.5 we present two examples of classes of structures, which show that the closure properties required by our constructions are indeed necessary.

3.1. **Summary of this section's main results.**
Table 1 summarises the time complexity of our algorithms (depending on the size of an input sentence) on the class of all ν-bounded structures. The summary differentiates between functions ν with either exponential or polynomial growth.

The precise statement of this section's first main result reads as follows; a proof is given in Subsection 3.3 below.

Theorem 3.1. *Let $\mathfrak{C}_\nu$ be a class of ν-bounded σ-structures that is closed under disjoint unions and induced substructures. There is an algorithm that, given an FO+MOD$_m$(σ)-sentence φ of quantifier rank $q \geqslant 0$ as input, constructs in time*

$$ \|\varphi\| \cdot S_\nu(S_\nu(3^q, \|\sigma\|), \|\sigma\|)^{(\log m)^2} \qquad \left(= \quad \|\varphi\| \cdot 2^{(\log m)^2 \cdot \nu(2^{\nu(3^q)^{O(\|\sigma\|)}})^{O(\|\sigma\|)}} \right) \qquad (3.1) $$

	extensions	homomorphisms
exponential ν	5-exp	4-exp
polynomial ν	3-exp	3-exp

Table 1: The time complexity of our algorithms. For each fixed $d \geqslant 3$, the class $\mathfrak{C}_d$ of all finite structures of degree $\leqslant d$ is a case of "exponential ν". The class $\mathfrak{C}_2$ of all finite structures of degree $\leqslant 2$ is a case of "polynomial ν".

an existential $\mathrm{FO}(\sigma)$-*sentence* ψ *such that the following holds:*
If φ *is preserved under extensions on* $\mathfrak{C}_\nu$, *then* φ *and* ψ *are equivalent on* $\mathfrak{C}_\nu$.

Consequently, if the function ν is exponential, φ is an $\mathrm{FO}+\mathrm{MOD}_m(\sigma)$-sentence, and σ consists of exactly the relation symbols that occur in φ, then the algorithm uses time 5-fold exponential in the size of φ. In particular, if $\nu = \nu_d$ for a fixed $d \geqslant 3$, and $\mathfrak{C}_\nu$ is the class of all finite structures of degree $\leqslant d$, then our algorithm uses time at most

$$\|\varphi\| \cdot 2^{(\log m)^2 \cdot d^{2^{d^{2^{O(q + \log \|\sigma\|)}}}}} \qquad \leqslant \qquad 2^{(\log m)^2 \cdot d^{2^{d^{2^{O(\|\varphi\|)}}}}} \tag{3.2}$$

when given a first-order sentence φ with modulo m counting quantifiers, quantifier rank q, and signature σ. Note that the constant suppressed by the O-notation does not depend on the particular signature σ. Furthermore, if φ does not contain any modulo counting quantifier, we can assume m to be 2, and hence $(\log m)^2 = 1$.

On the other hand, if ν is polynomial (e.g., $\nu = \nu_2$ and $\mathfrak{C}_\nu$ is the class of all finite structures of degree $\leqslant 2$), then expression (3.1) simplifies to the 3-fold exponential expression

$$\|\varphi\| \cdot 2^{(\log m)^2 \cdot 2^{2^{O(q \cdot \|\sigma\|)}}} \qquad \leqslant \qquad 2^{(\log m)^2 \cdot 2^{2^{O(\|\varphi\|^2)}}}. \tag{3.3}$$

This section's second main result reads as follows; a proof is given in Subsection 3.4 below.

Theorem 3.2. *Let* $\mathfrak{C}_\nu$ *be a class of* ν-*bounded* σ-*structures that is closed under disjoint unions and induced substructures and decidable in time* $t(n)$ *for some function* $t : \mathbb{N} \to \mathbb{N}$. *There is an algorithm which, given an input* $\mathrm{FO}+\mathrm{MOD}_m(\sigma)$-*sentence* φ *of quantifier rank* $q \geqslant 0$, *constructs in time*

$$2^{\|\varphi\| \cdot S_\nu(2 \cdot 3^q, \|\sigma\|)} \cdot t(S_\nu(2 \cdot 3^q, \|\sigma\|)) \qquad \left(= \quad 2^{\|\varphi\| \cdot 2^{\nu(2 \cdot 3^q)^{O(\|\sigma\|)}}} \cdot t(2^{\nu(2 \cdot 3^q)^{O(\|\sigma\|)}}) \right) \tag{3.4}$$

an existential-positive $\mathrm{FO}(\sigma)$-*sentence* ψ *of size* $2^{S_\nu(2 \cdot 3^q, \|\sigma\|)}$ *such that the following holds: If* φ *is preserved under homomorphisms on* $\mathfrak{C}_\nu$, *then* φ *and* ψ *are* $\mathfrak{C}_\nu$-*equivalent.*

Consequently, if the functions ν and t are 1-fold exponential, then the algorithm uses time 4-fold exponential in the size of the input sentence φ. In particular, if $\nu = \nu_d$ for a fixed $d \geqslant 3$ and $\mathfrak{C}_\nu$ is the class of all finite structures of degree $\leqslant d$, the algorithm uses time

$$2^{\|\varphi\| \cdot 2^{d^{2^{O(q + \log \|\sigma\|)}}}} \qquad \leqslant \qquad 2^{2^{d^{2^{O(\|\varphi\|)}}}}.$$

If ν is polynomial and t is at most 1-fold exponential (this is, in particular, true if $\mathfrak{C}_\nu$ is the class of all finite structures of degree $\leqslant 2$), the algorithm's running time is

$$2^{\|\varphi\| \cdot 2^{2^{O(q \cdot \|\sigma\|)}}} \qquad \leqslant \qquad 2^{2^{2^{O(\|\varphi\|^2)}}},$$

which is 3-fold exponential in the size of the input sentence φ.

3.2. **Introductory notes on the proofs of Theorem 3.1 and Theorem 3.2.** Consider a class $\mathfrak{C}$ of σ-structures and an FO+MOD$_m(\sigma)$-sentence φ. A σ-structure $\mathcal{A}$ is a $\mathfrak{C}$-*minimal model of* φ if $\mathcal{A} \in \mathfrak{C}$, $\mathcal{A}$ is a model of φ, and there is no proper induced substructure $\mathcal{B} \in \mathfrak{C}$ of $\mathcal{A}$ that is a model of φ.

Let $\mathfrak{C}_\nu$ be a class of ν-bounded σ-structures that is closed under disjoint unions and closed under induced substructures. The main combinatorial parts (Theorem 3.5 and Theorem 3.12) of the proofs of Theorem 3.1 and Theorem 3.2 provide upper bounds on the size of $\mathfrak{C}_\nu$-minimal models for FO+MOD$_m(\sigma)$-sentences that are preserved under extensions (respectively, homomorphisms) on $\mathfrak{C}_\nu$.

The proofs of both upper bounds on the size of the $\mathfrak{C}_\nu$-minimal models proceed as follows: Assume that φ is an FO+MOD$_m(\sigma)$-sentence that is preserved under extensions (homomorphisms) on $\mathfrak{C}_\nu$. We show that for every $\mathfrak{C}_\nu$-minimal model whose size exceeds the upper bound, a proper induced subtructure $\mathcal{A}' \in \mathfrak{C}_\nu$ of $\mathcal{A}$ can be constructed that is also a model of φ, which is a contradiction to the minimality of $\mathcal{A}$.

In both cases, we use the following generalisation of Hanf's theorem (see e.g. [EF99, Lib04]) by Nurmonen [Nur00] to sentences with modulo counting quantifiers. Two σ-structures $\mathcal{A}$ and $\mathcal{B}$ are (m,q)-*equivalent* ($\mathcal{A} \equiv_m^q \mathcal{B}$, for short) if they satisfy the same FO+MOD$_m(\sigma)$-sentences of quantifier rank at most q.

Theorem 3.3 (Theorem 3.4 in [Nur00]). *Let $\mathcal{A}, \mathcal{B}$ be σ-structures. Let $m \geqslant 2$ and $q \geqslant 0$. Suppose that for some $e \geqslant 0$ each 3^q-neighbourhood of an element in $\mathcal{A}$ or $\mathcal{B}$ has less than e elements and that for each 3^q-sphere τ (with one centre), $|\tau(\mathcal{A})|$ and $|\tau(\mathcal{B})|$ are congruent modulo m and either*

$$|\tau(\mathcal{A})| = |\tau(\mathcal{B})| \qquad or \qquad \big(\quad |\tau(\mathcal{A})| \geqslant t \quad and \quad |\tau(\mathcal{B})| \geqslant t \quad \big) \; ,$$

where $t := q \cdot e + 1$. Then $\mathcal{A} \equiv_m^q \mathcal{B}$. $\qquad \square$

The proof of Theorem 3.5 employs a novel inductive construction that constructs a sequence of structures from $\mathfrak{C}_\nu$ that alternates between proper induced substructures and disjoint extensions of $\mathcal{A}$ and finally stops with two consecutive (m,q)-equivalent structures, where $q \geqslant 0$ is the quantifier rank of φ.

Our proof of Theorem 3.12 is an adaptation of a result by Ajtai and Gurevich (Lemma 7.1 in [AG94]) where we use Nurmonen's theorem (Theorem 3.3) instead of Gaifman's theorem.

Finally, the upper bounds on the size of $\mathfrak{C}_\nu$-minimal models are used as an input to algorithms that compute a $\mathfrak{C}_\nu$-equivalent existential (existential-positive) FO(σ)-sentence for an input FO+MOD$_m(\sigma)$-sentence φ that is preserved under extensions (homomorphisms) on $\mathfrak{C}_\nu$ (see Lemma 3.8 and Lemma 3.13).

The construction for existential sentences generalises Lemma 8.4 in [DGKS07] to first-order sentences with modulo m counting quantifiers. Here, the handling of the modulo m counting quantifier requires an inductive construction (rather than a straightforward brute-force approach) to ensure the desired time complexity.

The construction for existential-positive sentences uses an algorithmic version of the Chandra-Merlin Theorem [CM77], which requires an additional assumption on the decidability of $\mathfrak{C}_\nu$.

3.3. **Preservation under extensions: Proof of Theorem 3.1.** This subsection is devoted to the proof of Theorem 3.1.

A set of elements B in a σ-structure $\mathcal{A}$ is *r-scattered*, for an $r \geqslant 0$, if for all distinct elements $a, b \in B$ we have $N_r^\mathcal{A}(a) \cap N_r^\mathcal{A}(b) = \emptyset$. We will make use of the following easy fact.

Lemma 3.4. *Let $\mathcal{A}$ be a ν-bounded σ-structure and let $m, r \geqslant 0$. If $|A| > (m-1) \cdot \nu(2r)$, then there exists an r-scattered subset of A of cardinality m.*

Proof. We show that $|A| \leqslant (m-1) \cdot \nu(2r)$, if there is no r-scattered subset of cardinality m in $\mathcal{A}$.

Choose a number $n < m$ such that there is an r-scattered subset B of cardinality n in $\mathcal{A}$, but no r-scattered subset of cardinality greater than n. Every element a of A has to be contained in the $2r$-neighbourhood of B (for otherwise, $B \cup \{a\}$ would be an r-scattered subset of $\mathcal{A}$ of cardinality $n+1$). Therefore, $|A| = |N_{2r}^\mathcal{A}(B)| \leqslant (m-1) \cdot \nu(2r)$. $\square$

The main combinatorial contribution of this subsection is an upper bound on the size of $\mathfrak{C}_\nu$-minimal models for $\text{FO}+\text{MOD}_m(\sigma)$-sentences, which is provided by the following theorem:

Theorem 3.5. *Let $\mathfrak{C}_\nu$ be a class of ν-bounded σ-structures that is closed under disjoint unions and induced substructures. Let $m \geqslant 1$ and let φ be an $\text{FO}+\text{MOD}_m(\sigma)$-sentence of quantifier rank $q \geqslant 0$ that is preserved under extensions on $\mathfrak{C}_\nu$. There is a number*

$$N_\nu(m, q, \|\sigma\|) \quad \in \quad m \cdot S_\nu(S_\nu(3^q, \|\sigma\|), \|\sigma\|) \qquad \left(= \quad m \cdot 2^{\nu(2^{\nu(3^q)^{O(\|\sigma\|)}})^{O(\|\sigma\|)}} \right)$$

such that every $\mathfrak{C}_\nu$-minimal model of φ has size at most $N_\nu(m, q, \|\sigma\|)$.

Proof. Let $r := 3^q$ and let $s \in S_\nu(r, \|\sigma\|)$ be the number of isomorphism types of r-spheres (with one centre) that are realised in σ-structures in $\mathfrak{C}_\nu$. Let $R := 2sr$ and let $S \in S_\nu(R, \|\sigma\|)$ be the number of isomorphism types of R-spheres (with one centre) that are realised in σ-structures in $\mathfrak{C}_\nu$. Finally, let $t := q \cdot (\nu(r)+1) + 1$ be the threshold from Theorem 3.3 for ν-bounded σ-structures and for quantifier rank q.

Let $\mathcal{A}$ be a $\mathfrak{C}_\nu$-minimal model of φ. Towards a contradiction, assume that the universe A of $\mathcal{A}$ has cardinality greater than $(2Stm-1) \cdot \nu(2R)$. Then, by Lemma 3.4, there exists an R-scattered subset of A of cardinality $2Stm$. Because there are at most S different R-spheres realised in $\mathcal{A}$, there is an R-scattered set X' of $2tm$ elements in A that realise the same R-sphere.

An r-sphere τ is *frequent in a structure $\mathcal{A}$* if $|\tau(\mathcal{A})| \geqslant t$. Otherwise, it is *rare in $\mathcal{A}$*. Note that each r-sphere realised by an element from the $(R-r)$-neighbourhood of X' is frequent in $\mathcal{A}$, because it occurs at least $2tm \geqslant t$ times in $\mathcal{A}$.

Let X be a subset of X' of cardinality tm. Since

$$\mathcal{N}_R^\mathcal{A}(a) \cong \mathcal{N}_R^\mathcal{A}(b) \quad \text{for each } a \in X \text{ and each } b \in X' \setminus X, \tag{3.5}$$

the following holds for each r-sphere τ:

$$\text{If} \quad |\tau(\mathcal{A}) \cap N_{R-r}^\mathcal{A}(X)| \geqslant t, \quad \text{also} \quad |\tau(\mathcal{A} - N_{R-r}^\mathcal{A}(X))| \geqslant t.$$
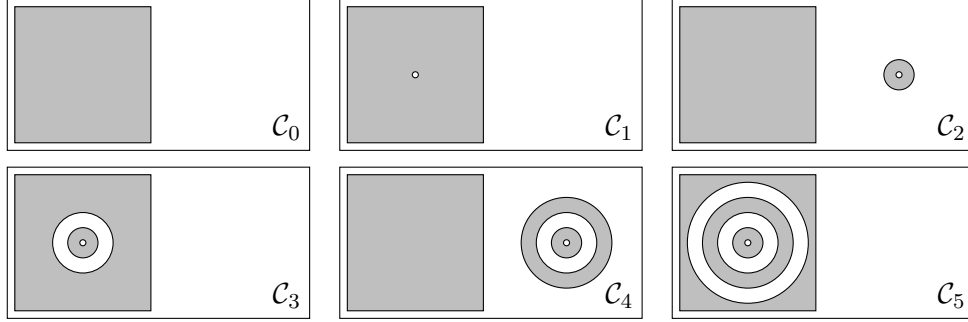
Figure 1: The first six elements of the sequence $(\mathcal{C}_i)_{i\geqslant 0}$. Note that for all even i, $\mathcal{C}_i$ is a disjoint extension of $\mathcal{A}$ while for each odd i, $\mathcal{C}_i$ is a proper induced substructure of $\mathcal{A}$.

I.e., every r-sphere τ that is realised by at least t elements of the $(R-r)$-neighbourhood of X is still frequent in the substructure of $\mathcal{A}$ induced by deleting the $(R-r)$-neighbourhood of X.

Consider the following sequences $(\mathcal{C}_i)_{i\geqslant 0}$ and $(\mathcal{D}_{2i})_{i>0}$ of σ-structures in $\mathfrak{C}_\nu$. Let $\mathcal{C}_0 := \mathcal{A}$ and $\mathcal{C}_1 := \mathcal{A} - X$. For each even $i > 1$, let $\mathcal{D}_i := \mathcal{C}_{i-1}[N^{\mathcal{A}}_{2(i-1)r}(X)]$. Note that for odd i, $\mathcal{D}_i$ is neither defined nor required. For each $i \geqslant 2$, let

$$
\mathcal{C}_i \;:=\; \begin{cases} \mathcal{A} \mathbin{\dot\cup} \mathcal{D}'_i & \text{if } i \text{ is even,} \\ \left(\mathcal{A} - N^{\mathcal{A}}_{2(i-1)r}(X)\right) \mathbin{\dot\cup} \mathcal{D}_{i-1} & \text{if } i \text{ is odd,} \end{cases}
$$

where $\mathcal{D}'_i$ is a structure isomorphic to $\mathcal{D}_i$ whose universe is disjoint to A. For all even i, $\mathcal{C}_i$ is a disjoint extension of $\mathcal{A}$, and for all odd i, $\mathcal{C}_i$ is a proper induced substructure of $\mathcal{A}$.

Let τ be an r-sphere. Recall that $|X| = tm$ and that the R-spheres around the elements of X are disjoint and isomorphic. This implies that, for each $i \in [1, s-1]$, there is a number $k \in \mathbb{Z}$ such that $|\tau(\mathcal{A} - N^{\mathcal{A}}_{2(i-1)r}(X))| = |\tau(\mathcal{A})| + ktm$ and therefore $|\tau(\mathcal{A} - N^{\mathcal{A}}_{2(i-1)r}(X))| \equiv |\tau(\mathcal{A})| \mod m$. Furthermore, since $|X| \equiv 0 \mod m$, it is straightforward to see that $|\tau(\mathcal{D}'_i)| = |\tau(\mathcal{D}_i)| \equiv 0 \mod m$ for all even $i \leqslant s$.

This immediately proves the following Claim 1.

Claim 3.6. For all $i < s$ and every r-sphere τ, $|\tau(\mathcal{C}_i)|$ and $|\tau(\mathcal{C}_{i+1})|$ are congruent modulo m.

A proof of the following Claim 2 is deferred to the end of the proof of Theorem 3.5.

Claim 3.7. The following holds for all $i < s$ and every r-sphere τ:

(a) If τ is frequent in $\mathcal{C}_i$, it is also frequent in $\mathcal{C}_{i+1}$.
(b) If τ is rare in $\mathcal{C}_i$ and rare in $\mathcal{C}_{i+1}$, then $|\tau(\mathcal{C}_i)| = |\tau(\mathcal{C}_{i+1})|$.

While every r-sphere that is frequent in $\mathcal{C}_i$ is also frequent in $\mathcal{C}_{i+1}$, the opposite is not necessarily true: There may be r-spheres that are rare in $\mathcal{C}_i$ but that occur frequently in $\mathcal{C}_{i+1}$. However, since there are only s pairwise non-isomorphic r-spheres in σ-structures in $\mathfrak{C}_\nu$, and $\mathcal{C}_0$ already contains frequent r-spheres, Claim 3.7 (a) implies that there has to be an $i < s$ such that all frequent r-spheres of $\mathcal{C}_{i+1}$ are frequent already in $\mathcal{C}_i$. Thus, for this particular i we know that any r-sphere is either frequent in $\mathcal{C}_{i+1}$ and in $\mathcal{C}_i$ or it is rare

in $\mathcal{C}_{i+1}$ and in $\mathcal{C}_i$. Hence, with Claim 3.7 (b) it follows that for every r-sphere τ, either $|\tau(\mathcal{C}_i)| = |\tau(\mathcal{C}_{i+1})|$ or τ is frequent in $\mathcal{C}_i$ and $\mathcal{C}_{i+1}$.

Together with Claim 1, we can conclude from Theorem 3.3 that $\mathcal{C}_i \equiv_m^q \mathcal{C}_{i+1}$. Therefore, by using Claim 2, the proof of Theorem 3.5 can be completed as follows: In case that i is even, we let $\mathcal{B} := \mathcal{C}_i$ and $\mathcal{A}' := \mathcal{C}_{i+1}$; and in case that i is odd, we let $\mathcal{B} := \mathcal{C}_{i+1}$ and $\mathcal{A}' := \mathcal{C}_i$. Since $\mathcal{B}$ is a disjoint extension of $\mathcal{A}$ with an induced substructure of $\mathcal{A}$ (and hence belongs to $\mathfrak{C}_\nu$, since $\mathfrak{C}_\nu$ is closed under disjoint unions and induced substructures), $\mathcal{A} \models \varphi$, and φ is preserved under extensions on $\mathfrak{C}_\nu$, we obtain that $\mathcal{B} \models \varphi$. Because $\mathcal{A}' \equiv_m^q \mathcal{B}$ and φ has quantifier rank q, we know that $\mathcal{A}' \models \varphi$. Since $\mathcal{A}'$ is a proper induced substructure of $\mathcal{A}$ and $\mathfrak{C}_\nu$ is closed under induced substructures, we have $\mathcal{A}' \in \mathfrak{C}_\nu$. However, this is a contradiction to the assumption that $\mathcal{A}$ is a $\mathfrak{C}_\nu$-minimal model of φ.

Therefore, the size $|A|$ of the universe of $\mathcal{A}$ is at most $(2Stm-1) \cdot \nu(2R)$. We let $N_\nu(m, q, \|\sigma\|) := (2Stm-1) \cdot \nu(2R)$. To obtain an upper bound on $N_\nu(m, q, \|\sigma\|)$, recall that $S \in S_\nu(R, \|\sigma\|)$ and $R \in 2r \cdot S_\nu(r, \|\sigma\|)$. Therefore, since $S_\nu(R, \|\sigma\|)$ is an abbreviation for the expression $2^{\nu(R)^{O(\|\sigma\|)}}$, we know that

$$S \quad \in \quad S_\nu(2r \cdot 2^{\nu(r)^{O(\|\sigma\|)}}, \|\sigma\|) \quad \subseteq \quad S_\nu(S_\nu(r, \|\sigma\|), \|\sigma\|). \tag{3.6}$$

The latter inclusion is correct since ν is strictly increasing. Similarly, we have that

$$\nu(2R) \quad = \quad \nu(2 \cdot 2r \cdot 2^{\nu(r)^{O(\|\sigma\|)}}) \quad \subseteq \quad \nu(S_\nu(r, \|\sigma\|)). \tag{3.7}$$

Using (3.6) and considering that $t \in O(q \cdot \nu(r))$, we conclude that

$$2Stm \quad = \quad O(2^{\nu(S_\nu(r, \|\sigma\|))^{O(\|\sigma\|)}} \cdot q \cdot \nu(r) \cdot m) \quad \subseteq \quad S_\nu(S_\nu(r, \|\sigma\|), \|\sigma\|) \cdot q \cdot m. \tag{3.8}$$

Putting (3.8) and (3.7) together, we obtain that

$$(2Stm-1) \cdot \nu(2R) \quad \subseteq \quad S_\nu(S_\nu(r, \|\sigma\|), \|\sigma\|) \cdot \nu(S_\nu(r, \|\sigma\|)) \cdot q \cdot m$$
$$\subseteq \quad S_\nu(S_\nu(r, \|\sigma\|), \|\sigma\|) \cdot q \cdot m.$$

Therefore, recalling that $r := 3^q$, we know that

$$N_\nu(m, q, \|\sigma\|) \quad \in \quad S_\nu(S_\nu(3^q, \|\sigma\|), \|\sigma\|) \cdot q \cdot m \quad \subseteq \quad m \cdot S_\nu(S_\nu(3^q, \|\sigma\|), \|\sigma\|).$$

All that remains to be done to finish the proof of Theorem 3.5 is to prove Claim 2.

Proof of Claim 2. Observe that, for all $i, j \leqslant s$,

$$\mathcal{C}_i[A \setminus N_{R-2r}^{\mathcal{A}}(X)] \quad \cong \quad \mathcal{C}_j[A \setminus N_{R-2r}^{\mathcal{A}}(X)]. \tag{3.9}$$

Let $i < s$. For the proof of Claim 2 (b) let τ be an r-sphere that is rare in $\mathcal{C}_i$ and $\mathcal{C}_{i+1}$. Since X is an R-scattered set of size $\geqslant t$, the rareness of τ implies that $\tau(\mathcal{C}_i)$ and $\tau(\mathcal{C}_{i+1})$ are subsets of $A \setminus N_{R-r}^{\mathcal{A}}(X)$. Hence, (3.9) implies that $|\tau(\mathcal{C}_i)| = |\tau(\mathcal{C}_{i+1})|$. This proves Claim 2 (b).

For the proof of Claim 2 (a), we distinguish between even and odd i.

Even i: Recall that $\mathcal{C}_0 = \mathcal{A}$ and that for each even $i \geqslant 2$, $\mathcal{C}_i = \mathcal{A} \dot{\cup} \mathcal{D}_i'$. Let τ be an r-sphere that is frequent in $\mathcal{C}_i$.
If τ is realised in $\mathcal{C}_i$ by an element in $N_{R-r}^{\mathcal{A}}(X)$ then, since X is R-scattered and of cardinality at least t, we know that $\tau(\mathcal{A}) \cap N_{R-r}^{\mathcal{A}}(X)$ contains at least t elements. Furthermore, we obtain by Observation (3.5) that also $\tau(\mathcal{A}) \cap N_{R-r}^{\mathcal{A}}(X' \setminus X)$ contains at least t elements. Since $X \subseteq X'$, the set X' is R-scattered, and $2ir < R$, we obtain that $|\tau(\mathcal{A} - N_{2ir}^{\mathcal{A}}(X))| \geqslant t$. Thus, τ is also frequent in $\mathcal{C}_{i+1}$.
If τ is realised in $\mathcal{C}_i$ by an element in $\mathcal{D}_i'$, for $i \geqslant 2$, then, since $\mathcal{D}_i' \cong \mathcal{D}_i$, τ is realised in $\mathcal{C}_{i+1}$

also by an element in $D_i \subseteq N_{R-r}^{\mathcal{A}}(X)$. Thus, τ is also frequent in $\mathcal{C}_{i+1}$.

Otherwise, we know that $|\tau(\mathcal{A}) \setminus N_{R-r}^{\mathcal{A}}(X)| \geqslant t$. Then it follows from Observation (3.9) that τ is also frequent in $\mathcal{C}_{i+1}$.

Odd i: Recall that $\mathcal{C}_1 = \mathcal{A} - X$ and that, for each odd $i \geqslant 3$, $\mathcal{C}_i = (\mathcal{A} - N_{2(i-1)r}^{\mathcal{A}}(X)) \dot{\cup} \mathcal{D}_{i-1}$. Let τ be an r-sphere that is frequent in $\mathcal{C}_i$.

If τ is realised in $\mathcal{C}_i$ by an element in $N_{2(i-1)r+r}^{\mathcal{A}}(X)$ then, since X is R-scattered and of cardinality $\geqslant t$, the set $\tau(\mathcal{C}_i) \cap N_{2(i-1)r+r}^{\mathcal{A}}(X)$ contains at least t elements. Since $\mathcal{D}_{i+1}$ is defined as $\mathcal{C}_i[N_{2(i-1)r+2r}^{\mathcal{A}}(X)]$, we obtain $|\tau(\mathcal{D}_{i+1})| \geqslant t$. Furthermore, since $\mathcal{D}_{i+1}' \cong \mathcal{D}_{i+1}$, we have $|\tau(\mathcal{D}_{i+1}')| \geqslant t$. Consequently, since $\mathcal{C}_{i+1} = \mathcal{A} \dot{\cup} \mathcal{D}_{i+1}'$, also $|\tau(\mathcal{C}_{i+1})| \geqslant t$. Thus, τ is frequent in $\mathcal{C}_{i+1}$.

Otherwise, we know that $\tau(\mathcal{C}_i) \cap (A \setminus N_{2(i-1)r+r}^{\mathcal{A}}(X))$ has cardinality $\geqslant t$. Since the r-sphere of each $a \in A \setminus N_{2(i-1)r+r}^{\mathcal{A}}(X)$ in $\mathcal{A} - N_{2(i-1)r}^{\mathcal{A}}(X)$ is isomorphic to its r-sphere in $\mathcal{A}$, τ is also frequent in $\mathcal{C}_{i+1} = \mathcal{A} \dot{\cup} \mathcal{D}_{i+1}'$.

This concludes the proof Claim 2 and Theorem 3.5. $\square$

For proving Theorem 3.1, it remains to do the following: for a given $\mathrm{FO}{+}\mathrm{MOD}_m(\sigma)$-sentence ψ that is preserved under extensions on $\mathfrak{C}_\nu$ and for an upper bound on the size of its $\mathfrak{C}_\nu$-minimal models (obtained from Theorem 3.5), construct an existential $\mathrm{FO}(\sigma)$-sentence that is $\mathfrak{C}_\nu$-equivalent to ψ. This is achieved by the following lemma which is a generalisation of Lemma 8.4 in [DGKS07] to sentences with modulo counting quantifiers.

Lemma 3.8. *Let $\mathfrak{C}$ be a class of σ-structures that is closed under induced substructures. There is an algorithm which, given a number $N \geqslant 0$ and an $\mathrm{FO}{+}\mathrm{MOD}_m(\sigma)$-sentence ψ of quantifier rank $q \geqslant 0$, constructs in time*

$$\|\psi\| \cdot N^{O(q \cdot \log m)}$$

an existential $\mathrm{FO}(\sigma)$-sentence ψ_N such that the following holds: If ψ is preserved under extensions on $\mathfrak{C}$ and every $\mathfrak{C}$-minimal model of ψ has at most N elements, then ψ_N is $\mathfrak{C}$-equivalent to ψ.

Furthermore, the constant suppressed by the O-notation does not depend on the signature σ. If ψ does not contain any modulo counting quantifier, the construction only takes time $O(\|\psi\| \cdot N^q)$.

The key ingredient for the proof of Lemma 3.8 is contained in the following lemma. Here, an *enumeration of a set A* is a tuple $(a_1, \ldots, a_M) \in A^M$ of length $M = |A|$ that contains each element of A exactly once (i.e., $A = \{a_1, \ldots, a_M\}$).

Lemma 3.9. *Let $k \geqslant 0$ and let ψ be an $\mathrm{FO}{+}\mathrm{MOD}_m(\sigma)$-formula with variables among $x_1, \ldots, x_k$. Let $M \geqslant 1$ and let $s \colon [1, k] \to [1, M]$. There is a quantifier-free $\mathrm{FO}(\sigma)$-formula $(\psi)_{M,s}$ with variables among $y_1, \ldots, y_M$ such that for each σ-structure $\mathcal{A}$ with exactly M elements and each enumeration $(a_1, \ldots, a_M)$ of A, the following equivalence holds:*

$$(\mathcal{A}, a_{s(1)}, \ldots, a_{s(k)}) \quad \models \quad \psi(x_1, \ldots, x_k)$$
$$\textit{iff} \qquad (\mathcal{A}, a_1, \ldots, a_M) \quad \models \quad (\psi)_{M,s}(y_1, \ldots, y_M).$$

Furthermore, if ψ does not contain any modulo counting quantifier, the above equivalence holds more generally for each σ-structure $\mathcal{A}$ with at most M elements and each tuple $(a_1, \ldots, a_M) \in A^M$ that contains each element of A at least once.

The formula, $(\psi)_{M,s}$ can be constructed in time

$$\|\psi\| \cdot M^{O(q \cdot \log m)},$$

and the constant suppressed by the O-notation does not depend on the signature σ.
If ψ does not contain any modulo counting quantifier, the construction of $(\psi)_{M,s}$ only takes time in

$$O(\|\psi\| \cdot M^q).$$

Before presenting the proof of Lemma 3.9, we first show how to use Lemma 3.9 for proving Lemma 3.8.

Proof of Lemma 3.8 using Lemma 3.9.
Choose k such that the variables occurring in ψ are among $x_1, \ldots, x_k$. For applying Lemma 3.9 for any $M \geqslant 1$, we let s be an arbitrary function from $[1, k]$ to $[1, M]$. For every $M \in [1, N]$ we apply Lemma 3.9 and let

$$\varphi_M(y_1, \ldots, y_M) \; := \; \bigwedge_{1 \leqslant i < j \leqslant M} \neg\, y_i{=}y_j \quad \wedge \quad (\psi)_{M,s}(y_1, \ldots, y_M).$$

Furthermore, we let

$$\psi_N \; := \; \bigvee_{M=1}^{N} \exists y_1 \cdots \exists y_M \; \varphi_M(y_1, \ldots, y_M).$$

Suppose that $\mathcal{A}' \in \mathfrak{C}$ is a model of ψ. Since ψ is preserved under extensions on $\mathfrak{C}$ and every $\mathfrak{C}$-minimal model of ψ has at most N elements, there is an induced substructure $\mathcal{A} \in \mathfrak{C}$ of $\mathcal{A}'$ with a universe of exactly $M \leqslant N$ pairwise distinct elements $a_1, \ldots, a_M$, such that $\mathcal{A} \models \psi$. Therefore, by Lemma 3.9, $(\mathcal{A}, a_1, \ldots, a_M)$ is a model of $\varphi_M(y_1, \ldots, y_M)$. Hence, $\mathcal{A} \models \psi_N$, and since ψ_N is existential, also $\mathcal{A}' \models \psi_N$.

On the other hand, suppose that $\mathcal{A}' \in \mathfrak{C}$ is a model of ψ_N. Then, there is an $M \in [1, N]$ and a substructure $\mathcal{A} \in \mathfrak{C}$ of $\mathcal{A}'$, induced by pairwise distinct elements $a_1, \ldots, a_M$ from A', such that $(\mathcal{A}, a_1, \ldots, a_M)$ is a model of $(\psi)_{M,s}(y_1, \ldots, y_M)$. Hence, $\mathcal{A} \models \psi$ and, since $\mathcal{A}'$ is an extension of $\mathcal{A}$, also $\mathcal{A}' \models \psi$.

Thus, ψ_N is $\mathfrak{C}$-equivalent to ψ.

By Lemma 3.9, $(\psi)_{M,s}$ can be constructed in time $\|\psi\| \cdot N^{O(q \cdot \log m)}$. Hence, ψ_N can be constructed in time

$$N \cdot (N^2 + \|\psi\| \cdot N^{O(q \cdot \log m)}) \quad \subseteq \quad \|\psi\| \cdot N^{O(q \cdot \log m)}.$$

Furthermore, if ψ does not contain any modulo counting quantifier, it suffices to let

$$\psi_N \; := \; \exists y_1 \cdots \exists y_N \; (\psi)_{N,s}(y_1, \ldots, y_N).$$

Here, the construction of $(\psi)_{N,s}$ and ψ_N takes time only $O(\|\psi\| \cdot N^q)$. $\qquad\square$

Proof of Lemma 3.9.
The proof proceeds by induction on the shape of ψ.

- For the induction start, let ψ be an atomic formula with variables among $x_1, \ldots, x_k$. In this case, $(\psi)_{M,s}$ is the formula obtained by replacing all occurrences of variables x_i in ψ by $y_{s(i)}$. It is easy to check that $(\psi)_{M,s}$ satisfies the condition of Lemma 3.9. Assume that $\psi = R(x_{i_1}, \ldots, x_{i_\ell})$ for a relation symbol $R \in \sigma$ of arity $\ell \geqslant 1$ and for $i_1, \ldots, i_\ell \in [1, k]$. For each σ-structure $\mathcal{A}$ with at most M elements and each tuple $(a_1, \ldots, a_M) \in A^M$ that contains each element of A at least once, the following equivalence holds:

$$(\mathcal{A}, a_{s(1)}, \ldots, a_{s(k)}) \models R(x_{i_1}, \ldots, x_{i_\ell})$$

iff $(a_{s(i_1)}, \ldots, a_{s(i_\ell)}) \in R^{\mathcal{A}}$

iff $(\mathcal{A}, a_1, \ldots, a_N) \models R(y_{s(i_1)}, \ldots, y_{s(i_\ell)})$.

For $\psi = (x_i = x_j)$ the argumentation is analogous.

- If $\psi = \neg\psi'$ then $(\psi)_{M,s} := \neg(\psi')_{M,s}$.
  Similarly, for each Boolean connective $\star \in \{\wedge, \vee, \rightarrow, \leftrightarrow\}$, we let

$$(\psi' \star \psi'')_{M,s} \quad := \quad (\psi')_{M,s} \star (\psi'')_{M,s}.$$

In each of these cases it is easy to check that the constructed formula satisfies the lemma's condition.

- If $\psi = \exists x_i\, \psi'$ then

$$(\psi)_{M,s} \quad := \quad \bigvee_{j=1}^{M} (\psi')_{M,s[i \to j]},$$

where $s[i \to j]$ is the function s' that agrees with s on all values except i, and $s'(i) = j$. Note that $\|(\psi)_{M,s}\| \in O(M \cdot \|(\psi')_{M,s'}\|)$. For each σ-structure $\mathcal{A}$ with at most M elements and each tuple $(a_1, \ldots, a_M) \in A^M$ that contains each element of A at least once, the following equivalence holds:

$$(\mathcal{A}, a_{s(1)}, \ldots, a_{s(k)}) \models \exists x_i\, \psi'(x_1, \ldots, x_k)$$

iff there is a $j \in [1, M]$ such that for $s' := s[i \to j]$,
  $(\mathcal{A}, a_{s'(1)}, \ldots, a_{s'(k)}) \models \psi'(x_1, \ldots, x_k)$

iff there is a $j \in [1, M]$ such that $(\mathcal{A}, a_1, \ldots, a_M) \models (\psi')_{M,s[i \to j]}(y_1, \ldots, y_M)$

iff $(\mathcal{A}, a_1, \ldots, a_M) \models (\psi)_{M,s}(y_1, \ldots, y_M)$.

- Accordingly, if $\psi = \forall x_i\, \psi'$ then

$$(\psi)_{M,s} \quad := \quad \bigwedge_{j=1}^{M} (\psi')_{M,s[i \to j]}.$$

- Finally, assume that $\psi = \exists^{0 \bmod m} x_i\, \psi'$. While it is straightforward to write, in an analogous manner to the two latter cases, a formula with the desired semantics but with size exponential in M, we employ an inductive construction whose size is linear in M.

  More precisely, we let $(\psi)_{M,s} := \gamma_0^{1,M}$, where $\gamma_0^{1,M}$ is the formula provided by the following claim:

Claim 3.10. For all $j, j' \in [1, M]$ with $j \leqslant j'$ and each $p \in [0, m-1]$, there is a quantifier-free formula $\gamma_p^{j,j'}(y_1, \ldots, y_M)$ such that for each σ-structure $\mathcal{A}$ with exactly M elements and each enumeration $(a_1, \ldots, a_M)$ of A, the following equivalence holds:

$$(\mathcal{A}, a_1, \ldots, a_M) \quad \models \quad \gamma_p^{j,j'}(y_1, \ldots, y_M)$$

if, and only if, the number of indices $\ell \in [j, j']$ such that

$$(\mathcal{A}, a_1, \ldots, a_M) \quad \models \quad (\psi')_{M, s[i \to \ell]}(y_1, \ldots, y_M)$$

is congruent to p modulo m.

Furthermore, the formula $\gamma_p^{j,j'}$ is of size

$$\|(\psi')_{M,s}\| \cdot (j'-j+1)^{O(\log m)},$$

where the constant suppressed by the O-notation does not depend on the signature σ.

To show that $\gamma_0^{1,M}$ satisfies the requirements of Lemma 3.9, let $\mathcal{A}$ be a σ-structure with exactly M elements and let $(a_1, \ldots, a_M)$ be an enumeration of A. Then, due to Claim 3.10, the following equivalence holds:

$$(\mathcal{A}, a_{s(1)}, \ldots, a_{s(k)}) \models \exists^{0 \bmod m} x_i \, \psi'(x_1, \ldots, x_k)$$

iff   the number of indices $\ell \in [1, M]$, such that for $s' := s[i \to \ell]$,
$(\mathcal{A}, a_{s'(1)}, \ldots, a_{s'(k)}) \models \psi'(x_1, \ldots, x_k)$, is congruent to 0 modulo m

iff   the number of indices $\ell \in [1, M]$, such that $(\mathcal{A}, a_1, \ldots, a_M) \models (\psi')_{M, s[i \to \ell]}(y_1, \ldots, y_M)$,
is congruent to 0 modulo m
(recall that ψ' and $(\psi')_{M, s[i \to \ell]}$ satisfy the assumptions of Lemma 3.9)

iff $(\mathcal{A}, a_1, \ldots, a_M) \models \gamma_0^{1,M}(y_1, \ldots, y_M)$.

By Claim 3, the size of $\gamma_0^{1,M}$ is in $\|(\psi')_{M,s}\| \cdot M^{O(\log m)}$. It remains to prove Claim 3.10.

Proof of Claim 3.10. We start by defining, for each $j \in [1, M]$ and each $p \in [0, m-1]$, the formula $\gamma_p^{j,j}$: We let $\gamma_0^{j,j} := \neg(\psi')_{M, s[i \to j]}$ and $\gamma_1^{j,j} := (\psi')_{M, s[i \to j]}$. For each $p > 1$, we let $\gamma_p^{j,j}$ be an unsatisfiable formula, e.g., $\neg y_1 = y_1$. By definition, these formulas satisfy Claim 3.10.

For all $j, j' \in [1, M]$ with $j < j'$ and each $p \in [0, m-1]$, we let

$$\gamma_p^{j,j'} := \bigvee_{(p_1, p_2) \in P} \left(\gamma_{p_1}^{j,h} \wedge \gamma_{p_2}^{h+1,j'} \right),$$

where $h := \left\lfloor \frac{j+j'}{2} \right\rfloor$ and P is the set of all tuples $(p_1, p_2) \in [0, m-1] \times [0, m-1]$ such that $p_1 + p_2$ is congruent to p modulo m. Note that $|P| = m$.

For each σ-structure $\mathcal{A}$ with exactly M elements and each enumeration $(a_1, \ldots, a_M)$ of A, the following equivalence holds:

$$(\mathcal{A}, a_1, \ldots, a_M) \models \gamma_p^{j,j'}(y_1, \ldots, y_M)$$

iff   there are $p_1, p_2 \in [0, m-1]$ with $p_1 + p_2 \equiv p \mod m$ such that
$(\mathcal{A}, a_1, \ldots, a_M) \models \gamma_{p_1}^{j,h}(y_1, \ldots, y_M)$ and $(\mathcal{A}, a_1, \ldots, a_M) \models \gamma_{p_2}^{h+1,j'}(y_1, \ldots, y_M)$

iff there are $p_1, p_2 \in [0, m-1]$ such that $p_1 + p_2 \equiv p \mod m$ and the number of indices ℓ in the intervals $[j, h]$ and $[h+1, j']$, such that $(\mathcal{A}, a_1, \dots, a_M) \models (\psi')_{M, s[i \to \ell]}(y_1, \dots, y_M)$, is congruent to p_1 modulo m and congruent to p_2 modulo m, respectively,

iff the number of indices $\ell \in [j, j']$, such that $(\mathcal{A}, a_1, \dots, a_M) \models (\psi')_{M, s[i \to \ell]}(y_1, \dots, y_M)$, is congruent to p modulo m.

It remains to show that for all $j, j' \in [1, M]$ and $p \in [0, m-1]$,
$$\|\gamma_p^{j,j'}\| \quad \in \quad \|(\psi')_{M,s}\| \cdot (j'-j+1)^{O(\log m)}.$$

Note that the number of occurrences of $(\psi')_{M,s}$ in $\gamma_p^{j,j'}$ in the "leaves" of the inductive definition of $\gamma_p^{j,j'}$ is in $m^{O(\log(j'-j+1))}$. Moreover, taking into account the "inner nodes" and the "leaves" of the inductive definition of $\gamma_p^{j,j'}$ separately, we can bound the size of $\gamma_p^{j,j'}$ from above by
$$\|\gamma_p^{j,j'}\| \quad \in \quad T(j'-j+1) \quad + \quad \|(\psi')_{M,s}\| \cdot m^{O(\log(j'-j+1))}, \tag{3.10}$$
where the function $T \colon \mathbb{N}_{\geqslant 1} \to \mathbb{N}_{\geqslant 1}$ is defined inductively by $T(1) := 1$ and, for all $n > 1$,
$$T(n) \quad := \quad a \cdot T(\lceil n/2 \rceil) \tag{3.11}$$
for a number $a \geqslant 1$ of size $O(m)$. For an upper bound on the growth of $T(n)$ we apply a variation of the well-known Master Theorem, which reads as follows:

Theorem 3.11 (adapted from Theorem 4.1 of [CLRS09])**.** *Let $a \geqslant 1$ and let $b > 1$. Let $T \colon \mathbb{N}_{\geqslant 1} \to \mathbb{N}$ be defined by $T(1) := 1$ and, for all $n > 1$, $T(n) := a \cdot T(\lceil n/b \rceil)$. Then, $T(n) \in \Theta(n^{\log_b a})$.* $\qquad\square$

By choosing the number a as in (3.11) and with $b := 2$, we can conclude that
$$T(n) \quad \in \quad \Theta(n^{\log a}) \quad \subseteq \quad n^{O(\log m)}.$$
Replacing $T(j'-j+1)$ by $(j'-j+1)^{O(\log m)}$ in (3.10) we obtain that
$$\begin{aligned}
\|\gamma_p^{j,j'}\| \quad &\in \quad (j'-j+1)^{O(\log m)} \; + \; \|(\psi')_{M,s}\| \cdot m^{O(\log(j'-j+1))} \\
&\subseteq \quad \|(\psi')_{M,s}\| \cdot \big((j'-j+1)^{O(\log m)} \; + \; m^{O(\log(j'-j+1))}\big) \\
&\subseteq \quad \|(\psi')_{M,s}\| \cdot (j'-j+1)^{O(\log m)}
\end{aligned}$$
This completes the proof of Claim 3.10. $\qquad\square$

For a bound on the size of $(\psi)_{M,s}$, note that the only size increasing steps in the inductive translation are the ones for the quantifiers, which increase the size of the formula by a factor of M, for first-order quantifiers, and by a factor in $M^{O(\log m)}$, for modulo counting quantifiers. It follows that $\|(\psi)_{M,s}\| \in \|\psi\| \cdot M^{O(q \cdot \log m)}$, where q is the quantifier rank of ψ. If ψ does not contain any modulo counting quantifier, $\|(\psi)_{M,s}\| \in O(\|\psi\| \cdot N^q)$. It is easy to see that the inductive translation can be also carried out in time $O(\|(\psi)_{M,s}\|)$. This completes the proof of Lemma 3.9. $\qquad\square$

Theorem 3.1 is now obtained by a straightforward combination of Lemma 3.8 and Theorem 3.5.

Proof of Theorem 3.1. Let $m \geqslant 2$ and let φ be an FO+MOD$_m(\sigma)$-sentence of quantifier rank $q \geqslant 0$ that is preserved under extensions on $\mathfrak{C}_\nu$. Let

$$N \quad := \quad N_\nu(m, q, \|\sigma\|) \quad \in \quad m \cdot S_\nu(S_\nu(3^q, \|\sigma\|), \|\sigma\|) \quad = \quad m \cdot 2^{\nu(S_\nu(3^q, \|\sigma\|))^{O(\|\sigma\|)}}$$

be the upper bound on the size of $\mathfrak{C}_\nu$-minimal models of φ, obtained from Theorem 3.5.

By Lemma 3.8 there is an algorithm that constructs an existential FO(σ)-sentence that is $\mathfrak{C}_\nu$-equivalent to φ in time

$$
\begin{aligned}
\|\varphi\| \cdot N^{O(q \cdot \log m)} \quad &= \quad \|\varphi\| \cdot \left(m \cdot 2^{\nu(S_\nu(3^q, \|\sigma\|))^{O(\|\sigma\|)}} \right)^{O(q \cdot \log m)} \\[2mm]
&\subseteq \quad \|\varphi\| \cdot 2^{O((\log m)^2 \cdot q \cdot \nu(S_\nu(3^q, \|\sigma\|))^{O(\|\sigma\|)})} \\[2mm]
&\subseteq \quad \|\varphi\| \cdot \left(2^{\nu(S_\nu(3^q, \|\sigma\|))^{O(\|\sigma\|)}} \right)^{(\log m)^2} \\[2mm]
&= \quad \|\varphi\| \cdot S_\nu(S_\nu(3^q, \|\sigma\|), \|\sigma\|)^{(\log m)^2}.
\end{aligned}
$$

Here we make again use of the assumption that ν is strictly increasing. This concludes the proof of Theorem 3.1. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $\square$

3.4. **Preservation under homomorphisms: Proof of Theorem 3.2.** The combinatorial essence of the proof of Theorem 3.2 is contained in the following theorem.

Theorem 3.12. *Let $\mathfrak{C}_\nu$ be a class of ν-bounded σ-structures that is closed under disjoint unions and induced substructures. Let φ be an FO+MOD$_m(\sigma)$-sentence of quantifier rank $q \geqslant 0$ that is preserved under homomorphisms on $\mathfrak{C}_\nu$. There is a number*

$$N_\nu(q, \|\sigma\|) \quad \in \quad S_\nu(2 \cdot 3^q, \|\sigma\|) \qquad \left(= \quad 2^{\nu(2 \cdot 3^q)^{O(\|\sigma\|)}} \right)$$

such that every $\mathfrak{C}_\nu$-minimal model of φ has size at most $N_\nu(q, \|\sigma\|)$.

Proof. The proof is similar to the proof of Lemma 7.1 in [AG94]. However, it does not rely on Gaifman's theorem but uses Nurmonen's generalisation of Hanf's theorem, stated in Theorem 3.3. Towards applying Theorem 3.3, we let $r := 3^q$, let $t := q \cdot (\nu(r)+1)+1$, and let $s \in S_\nu(2r, \|\sigma\|)$ be the number of non-isomorphic $2r$-spheres (with one centre) realised in σ-structures in $\mathfrak{C}_\nu$.

Let $\mathcal{A}$ be a $\mathfrak{C}_\nu$-minimal model of φ. Towards a contradiction, assume that $|A| > s \cdot \nu(4r)$. By Lemma 3.4, $\mathcal{A}$ contains a $2r$-scattered set of size $s+1$. Thus, since there are at most s non-isomorphic $2r$-spheres realised in $\mathcal{A}$, there must be two elements $a_1, a_2 \in A$ with disjoint and isomorphic $2r$-neighbourhoods.

Let $\mathcal{A}' := \mathcal{A} - \{a_1\}$. Clearly, the r-spheres of elements in $A \setminus N_r^{\mathcal{A}}(a_1)$ are the same in $\mathcal{A}$ and in $\mathcal{A}'$. But the r-sphere of an element in $N_r^{\mathcal{A}}(a_1)$ might change when moving from $\mathcal{A}$ to $\mathcal{A}'$. However, by our choice of a_1 and a_2 we know that every r-sphere that is realised in $\mathcal{A}$ is also realised in $\mathcal{A}'$ (for elements outside the r-neighbourhood of a_1 this is obvious; and for elements $a_1' \in N_r^{\mathcal{A}}(a_1)$, the r-sphere of a_1' in $\mathcal{A}$ is realised in $\mathcal{A}'$ by the corresponding element a_2' in the r-neighbourhood of a_2).

Now let $\mathcal{B}'$ be the disjoint union of $t \cdot m$ copies of $\mathcal{A}'$, and let $\mathcal{B}$ be the disjoint union of $\mathcal{B}'$ and of $t \cdot m$ copies of $\mathcal{A}$.

By construction, every r-sphere that is realised in $\mathcal{B}$ is also realised in $\mathcal{B}'$, and vice versa. Furthermore, the number of realisations of any r-sphere in $\mathcal{B}$ or $\mathcal{B}'$ is a multiple of $t \cdot m$. In particular, $\mathcal{B}$ and $\mathcal{B}'$ satisfy the assumption of Theorem 3.3, and therefore we have $\mathcal{B} \equiv^q_m \mathcal{B}'$. Thus, $\mathcal{B} \models \varphi$ iff $\mathcal{B}' \models \varphi$.

Furthermore, since $\mathfrak{C}_\nu$ is closed under taking induced substructures and disjoint unions, we know that $\mathcal{B}'$ and $\mathcal{B}$ belong to $\mathfrak{C}_\nu$. Obviously, there is a homomorphism that maps $\mathcal{A}$ to one of the copies of $\mathcal{A}$ in $\mathcal{B}$. Since φ is preserved under homomorphisms on $\mathfrak{C}_\nu$ and $\mathcal{A} \models \varphi$, we thus have $\mathcal{B} \models \varphi$, and hence also $\mathcal{B}' \models \varphi$.

Recall that $\mathcal{B}'$ is a disjoint union of copies of $\mathcal{A}'$. By mapping each element of each copy of $\mathcal{A}'$ to the corresponding element in $\mathcal{A}'$, we obtain a homomorphism from $\mathcal{B}'$ to $\mathcal{A}'$. Hence, since φ is preserved under homomorphisms and $\mathcal{B}' \models \varphi$, we obtain that also $\mathcal{A}' \models \varphi$. This, however, contradicts our assumption that $\mathcal{A}$ is a $\mathfrak{C}_\nu$-*minimal* model of φ. Therefore, $|A| \leqslant s \cdot \nu(4r)$. We obtain that

$$N_\nu(q, \|\sigma\|) \;:=\; s \cdot \nu(4r) \quad \in \quad S_\nu(2 \cdot 3^q, \|\sigma\|).$$

This concludes the proof of Theorem 3.12. $\qquad\square$

In the following lemma, we construct existential-positive $\mathrm{FO}(\sigma)$-sentences for $\mathrm{FO+MOD}_m(\sigma)$-sentences that are preserved under homomorphisms. The proof is an algorithmic version of the proof of Theorem 3.1 in [ADK06].

Lemma 3.13. *Let $\mathfrak{C}$ be a class of σ-structures that is decidable in time $t(n)$ for an arbitrary function $t \colon \mathbb{N} \to \mathbb{N}$. There is an algorithm that, on input of a number $N \geqslant 0$ and an $\mathrm{FO+MOD}_m(\sigma)$-sentence φ of quantifier rank $q \geqslant 0$ constructs in time*

$$2^{N^{O(\|\sigma\|)}} \cdot \left(N^{O(\|\sigma\| \cdot \|\varphi\|)} + t(O(N^{\|\sigma\|})) \right)$$

an existential-positive $\mathrm{FO}(\sigma)$-sentence ψ of size $2^{N^{O(\|\sigma\|)}}$, such that the following holds: If φ is preserved under homomorphisms on $\mathfrak{C}$ and every $\mathfrak{C}$-minimal model of φ has at most N elements, then ψ is $\mathfrak{C}$-equivalent to φ.

Proof. For each finite σ-structure $\mathcal{A}$ let $\gamma_\mathcal{A}$ be the *canonical conjunctive query* associated with $\mathcal{A}$. I.e., $\gamma_\mathcal{A} = \exists x_1 \cdots \exists x_{|A|}\, \theta_\mathcal{A}$ where $x_1, \ldots, x_{|A|}$ are variables representing the elements $a_1, \ldots, a_{|A|}$ of $\mathcal{A}$'s universe A, and $\theta_\mathcal{A}$ is the conjunction of all atoms of the form $R(x_{i_1}, \ldots, x_{i_r})$ where $R \in \sigma$, $r = \mathrm{ar}(R)$, $i_1, \ldots, i_r \in \{1, \ldots, |A|\}$, and $(a_{i_1}, \ldots, a_{i_r}) \in R^\mathcal{A}$. The well-known *Chandra-Merlin Theorem* [CM77] states that for any σ-structure $\mathcal{B}$, there is a homomorphism from $\mathcal{A}$ to $\mathcal{B}$ if, and only if, $\mathcal{B} \models \gamma_\mathcal{A}$. Clearly, given $\mathcal{A}$, the sentence $\gamma_\mathcal{A}$ can be constructed in time $O(\|\mathcal{A}\|) = O(|A|^{\|\sigma\|})$.

On input of φ and N, the lemma's algorithm

(1) computes the set M that consists all models of φ
with universe $\{1, \ldots, n\}$, for $n \leqslant N$, that belong to $\mathfrak{C}$,

(2) if $M = \emptyset$, it outputs the formula $\psi := \textit{false}$,

(3) if $M \neq \emptyset$, it outputs the formula $\psi := \bigvee_{\mathcal{A} \in M} \gamma_\mathcal{A}$.

Clearly, $|M|$ is bounded by the number of all σ-structures with universe $\{1, \ldots, n\}$ and $n \leqslant N$. Hence, the size of ψ is in

$$O(|M| \cdot N^{\|\sigma\|}) \quad \subseteq \quad 2^{N^{O(\|\sigma\|)}}$$

Obviously, ψ is an existential-positive $\mathrm{FO}(\sigma)$-sentence.

Before giving details on the algorithm's step 1 and its running time, let us first show that ψ is $\mathfrak{C}$-equivalent to φ, provided that φ is preserved under homomorphisms on $\mathfrak{C}$ and that every $\mathfrak{C}$-minimal model of φ has at most N elements. To this end, let $\mathcal{B}$ be an arbitrary σ-structure in $\mathfrak{C}$.

If $\mathcal{B} \models \psi$, then there is an $\mathcal{A} \in M$ such that $\mathcal{B} \models \gamma_{\mathcal{A}}$. Due to the Chandra-Merlin Theorem, there is a homomorphism from $\mathcal{A}$ to $\mathcal{B}$. As $\mathcal{A} \in \mathfrak{C}$ and $\mathcal{A} \models \varphi$, and since φ is preserved under homomorphisms on $\mathfrak{C}$, we obtain that $\mathcal{B} \models \varphi$.

On the other hand, if $\mathcal{B} \models \varphi$, then let $\mathcal{A}$ be a minimal induced substructure of $\mathcal{B}$ such that $\mathcal{A} \in \mathfrak{C}$ and $\mathcal{A} \models \varphi$. I.e., $\mathcal{A}$ is a $\mathfrak{C}$-minimal model of φ. By assumption, N is an upper bound on the size of the universe of $\mathcal{A}$. Thus, by our choice of M, the set M contains a structure $\mathcal{A}'$ that is isomorphic to $\mathcal{A}$. Since $\mathcal{A}$ is a substructure of $\mathcal{B}$, the particular choice of the formula $\gamma_{\mathcal{A}'}$ implies that $\mathcal{B} \models \gamma_{\mathcal{A}'}$. Since $\mathcal{A}' \in M$, we obtain that also $\mathcal{B} \models \psi$. In summary, this shows that ψ is $\mathfrak{C}$-equivalent to φ.

Let us now turn to the algorithm's step 1 and the analysis of its time complexity. To compute M, the algorithm enumerates all σ-structures $\mathcal{A}$ with $A = \{1, \ldots, n\}$ and $n \leqslant N$, and checks for each such $\mathcal{A}$ whether $\mathcal{A} \models \varphi$ and $\mathcal{A} \in \mathfrak{C}$.

By assumption, the question whether $\mathcal{A} \in \mathfrak{C}$ can be answered within time $t(\|\mathcal{A}\|) \in t(O(N^{\|\sigma\|}))$. Using the naive model checking algorithm for $\mathrm{FO}+\mathrm{MOD}_m$, the question whether $\mathcal{A} \models \varphi$ can be answered within time $\|\mathcal{A}\|^{O(\|\varphi\|)} \subseteq N^{O(\|\sigma\| \cdot \|\varphi\|)}$. Since $2^{N^{O(\|\sigma\|)}}$ is an upper bound on the number of σ-structures with universe $\{1, \ldots, n\}$ and $n \leqslant N$, the entire computation of M takes time at most

$$2^{N^{O(\|\sigma\|)}} \cdot \left(N^{O(\|\sigma\| \cdot \|\varphi\|)} + t(O(N^{\|\sigma\|})) \right).$$

This completes the proof of Lemma 3.13. $\qquad\square$

Theorem 3.2 is now obtained by a straightforward combination of Lemma 3.13 and Theorem 3.12 (in the analogous way as Theorem 3.1 was obtained by combining Lemma 3.8 with Theorem 3.5).

3.5. **Closure properties.** Our Theorems 3.1 and 3.2 require that the considered classes be closed under disjoint unions and induced substructures. In this section we provide simple examples which show that these closure properties are indeed necessary. Both examples use graphs that are directed paths where some endpoints are colored green. These are represented as structures over the signature $\sigma := \{E, G\}$ as usual. That is, the binary relation symbol E is interpreted by the edge relation and the unary relation symbol G is interpreted by the set of green vertices. In the following, a vertex is a *left endpoint* or a *right endpoint* if it is has, respectively, no ingoing edge or no outgoing edge. An *endpoint* is either a left or a right endpoint. For $n \geqslant 1$, a directed path on n vertices where exactly the endpoints are colored green will be denoted by $\mathcal{P}_n$ and a directed path on $2n+1$ vertices where just the central vertex is colored green will be denoted by $\mathcal{P}_n^C$.

Theorem 3.14. *There is a class $\mathfrak{C}_1$ of σ-structures of degree at most 2 that is closed under substructures but not under disjoint unions, and there is an $\mathrm{FO}(\sigma)$-sentence φ that is preserved under extensions and homomorphisms on $\mathfrak{C}_1$, but that has no $\mathfrak{C}_1$-equivalent existential $\mathrm{FO}(\sigma)$-sentence.*

Proof. Let $\mathfrak{C}_1$ be the class that contains a σ-structure if there is an $n \geqslant 1$ such that the structure is isomorphic to a substructure of $\mathcal{P}_n$. By construction, $\mathfrak{C}_1$ is closed under substructures. It is not closed under disjoint unions, since e.g. the disjoint union of two copies of $\mathcal{P}_n$ is not a substructure of any $\mathcal{P}_m$.

There is an obvious $\mathrm{FO}(\sigma)$-sentence φ that is satisfied by a σ-structure $\mathcal{A}$ iff $|A| \geqslant 3$ and all endpoints of $\mathcal{A}$ are green. The models of φ that belong to $\mathfrak{C}_1$ are exactly the structures that are isomorphic to $\mathcal{P}_n$ for some $n \geqslant 3$, because each proper substructure $\mathcal{A}$ of $\mathcal{P}_n$ contains an endpoint that is not green. The sentence φ is preserved under homomorphisms (and hence also under extensions) on $\mathfrak{C}_1$ for the trivial reason that the only structure in $\mathfrak{C}_1$ to which there is an homomorphism from $\mathcal{P}_n$ is $\mathcal{P}_n$ itself.

It remains to show that on $\mathfrak{C}_1$ the formula φ is not equivalent to an existential sentence. Assume towards a contradiction that φ is $\mathfrak{C}_1$-equivalent to an existential sentence $\psi :=$ $\exists x_1 \cdots \exists x_k \, \xi$, where $k \geqslant 1$ and where ξ is quantifier-free. In particular, $\mathcal{P}_{k+3} \models \psi$ so that there are $a_1, \ldots, a_k \in P_{k+3}$ for which $(\mathcal{P}_{k+3}, a_1, \ldots, a_k) \models \xi$. Let $\mathcal{P}$ be the substructure of $\mathcal{P}_{k+3}$ induced by $\{a_1, \ldots, a_k\}$. Clearly, $\mathcal{P} \cong \mathcal{P}_{k+3}[\{a_1, \ldots, a_k\}]$. Thus $(\mathcal{P}, a_1, \ldots, a_k) \models \xi$ and so $\mathcal{P} \models \psi$. On the other hand, $\mathcal{P}$ contains at least one endpoint that is not colored green. Therefore, $\mathcal{P} \not\models \varphi$. This contradicts our assumption that φ and ψ are equivalent. $\square$

Theorem 3.15. *There is a class $\mathfrak{C}_2$ of σ-structures of degree at most 2 that is closed under disjoint unions but not under induced substructures, and there is an $\mathrm{FO}(\sigma)$-sentence that is preserved under extensions and homomorphisms on $\mathfrak{C}_2$, but that has no $\mathfrak{C}_2$-equivalent existential $\mathrm{FO}(\sigma)$-sentence.*

Proof. Let $\mathfrak{C}_2$ be the class of all σ-structures that are disjoint unions of structures that are isomorphic to $\mathcal{P}_n$ or $\mathcal{P}_n^C$, for possibly different lengths $n \geqslant 1$. By construction, $\mathfrak{C}_2$ is closed under disjoint unions. It is not closed under induced substructures, since e.g. $\mathcal{P}_3$ has an isolated vertex that is not colored green as an induced substructure, but such graphs do not belong to $\mathfrak{C}_2$.

There is an obvious $\mathrm{FO}(\sigma)$-sentence φ that is satisfied by a σ-structure $\mathcal{A}$ iff it contains a green endpoint. A structure belonging to $\mathfrak{C}_2$ satisfies φ iff it contains a copy of $\mathcal{P}_n$ for some $n \geqslant 1$, since the $\mathcal{P}_n^C$ do not contain green endpoints. The sentence φ is preserved under homomorphisms (and hence also under extensions) on $\mathfrak{C}_2$ because $\mathcal{P}_n$ cannot be mapped homomorphically to a $\mathcal{P}_m^C$ for any m whatsoever, due to the two green endpoints.

It remains to show that φ is not $\mathfrak{C}_2$-equivalent to an existential sentence. Assume to the contrary that φ is $\mathfrak{C}_2$-equivalent to a sentence $\psi := \exists x_1 \cdots \exists x_k \, \xi$, where $k \geqslant 1$ and where ξ is quantifier-free. Then $\mathcal{P}_{k+1} \models \psi$, i.e. there are $a_1, \ldots, a_k \in P_{k+1}$ for which $(\mathcal{P}_{k+1}, a_1, \ldots, a_k) \models \xi$. Let $M := \{a_1, \ldots, a_k\}$. We partition M into sets L, R, where L is the (possibly empty) set of all vertices from M that belong to the connected component, in $\mathcal{P}_{k+1}[M]$, of the left endpoint of $\mathcal{P}_{k+1}$, and $R := M \setminus L$. Clearly, L and R are disconnected in $\mathcal{P}_{k+1}[M]$. In particular, the set L cannot contain the right endpoint of $\mathcal{P}_{k+1}$.

Suppose that L is empty and thus, $R = M$. There is a set $M' \subseteq P_k^C$ of vertices ("left of" the central green vertex) in $\mathcal{P}_k^C$ such that $\mathcal{P}_k^C[M'] \cong \mathcal{P}_{k+1}[M]$. Hence, $\mathcal{P}_k^C \models \psi$. But clearly $\mathcal{P}_k^C[M'] \not\models \varphi$. This is a contradiction.

In the case that L is not empty, the path $\mathcal{P}_k^C$ contains induced substructures that are isomorphic to $\mathcal{P}_{k+1}[L]$ and $\mathcal{P}_{k+1}[R]$, respectively ("right of" and "left of" the central green vertex). Let $\mathcal{A}$ be the disjoint union of two copies $\mathcal{A}_1$ and $\mathcal{A}_2$ of $\mathcal{P}_k^C$. We map the elements of L and R to corresponding elements of $\mathcal{A}_1$ and $\mathcal{A}_2$, respectively. Now let M' be the image

of M under this mapping. It is easy to verify that $\mathcal{A}[M'] \cong \mathcal{P}_{k+1}[M]$. Hence $\mathcal{A} \models \psi$. But clearly $\mathcal{A} \not\models \varphi$. This is a contradiction. $\qquad\square$

4. Feferman-Vaught decompositions

Throughout this section, let $\nu \colon \mathbb{N} \to \mathbb{N}$ be a fixed time-constructible strictly increasing function and let $\mathfrak{C}_\nu$ be an arbitrary class of ν-bounded σ-structures for a finite relational signature σ.

Let $(P_i)_{i \geqslant 1}$ be a sequence of unary relation symbols that are not already contained in σ. For every $s \geqslant 1$, by σ_s we denote the signature $\sigma \cup \{P_1, \ldots, P_s\}$.

Recall that a disjoint union $\mathcal{B} = \mathcal{A}_1 \,\dot\cup\, \cdots \,\dot\cup\, \mathcal{A}_s$ of σ-structures $\mathcal{A}_1, \ldots, \mathcal{A}_s$ involves injective functions $f_i \colon A_i \to B$ for each $i \in [1, s]$ such that $f_1(A_1), \ldots, f_s(A_s)$ is a partition of B. The *disjoint sum* $\mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_s$ *of* $\mathcal{A}_1, \ldots, \mathcal{A}_s$ is a σ_s-structure $\mathcal{A}$ that expands the disjoint union of $\mathcal{A}_1, \ldots, \mathcal{A}_s$ by the unary relations $P_i^{\mathcal{A}} := f_i(A_i)$, for all $i \in [1, s]$. Clearly, $P_1^{\mathcal{A}}, \ldots, P_s^{\mathcal{A}}$ is a partition of A, and for all $a \in P_i^{\mathcal{A}}$ and $b \in P_j^{\mathcal{A}}$ with $i, j \in [1, s]$ and $i \neq j$, there is no edge between a and b in the Gaifman graph of $\mathcal{A}$.

In Subsection 4.2 we provide an algorithm that computes, for every $s \geqslant 1$ and each $\mathrm{FO}(\sigma_s)$-formula φ, a *Feferman-Vaught decomposition* of φ, that is, a decomposition into a Boolean combination of $\mathrm{FO}(\sigma)$-formulas which is equivalent to φ on each disjoint sum of s structures in $\mathfrak{C}_\nu$ (for a precise definition, see Subsection 4.1). For functions ν of exponential growth, this algorithm has 3-fold exponential time complexity in terms of the input formula; for polynomial ν, the time complexity is 2-fold exponential. In Subsection 4.3 we show how to extend the algorithm to products of structures obtained by applying *transductions* to disjoint sums, e.g., to *direct products* of structures as well as to *cartesian products* and *strong products* of graphs.

In Section 5.3 we show that the algorithm's time complexity is basically optimal: For structures of degree 3, a 3-fold exponential blow-up of the decomposition in terms of the size of the input formula is unavoidable, and for structures of degree 2 there is still a 2-fold exponential blow-up.

4.1. **Disjoint decompositions.** Before presenting this section's main results, we give a precise definition of the decompositions constructed by our algorithm. These decompositions are a special case of so-called *reduction sequences* [Mak04]. They give conditions for the validity of an $\mathrm{FO}(\sigma_s)$-formula in a disjoint sum of structures in terms of a Boolean combination of $\mathrm{FO}(\sigma)$-formulas that speak about the component structures of the disjoint sum.

Let $s \geqslant 1$ and let $\overline{x}$ be a tuple of $n \geqslant 0$ variables. For each $i \in [1, s]$, let Δ_i be a finite set of $\mathrm{FO}(\sigma)$-formulas δ with $\mathrm{free}(\delta) \subseteq \overline{x}$ and let β be a propositional formula with variables from the set $X_D := \{X_{i,\delta} \colon i \in [1, s], \delta \in \Delta_i\}$. The tuple $D = (\Delta_1, \ldots, \Delta_s, \beta)$ is an *s-reduction sequence over $\overline{x}$* (for short: *reduction sequence*). The size $\|D\|$ of D is defined as $\|\beta\| + \sum_{i=1}^{s} \sum_{\delta \in \Delta_i} \|\delta\|$, where $\|\beta\|$ is the size of the propositional formula β when viewed as a word over the alphabet $\{\neg, \wedge, \vee, \rightarrow, \leftrightarrow, (,)\} \cup X_D$.

Let $\mathcal{A}_1, \ldots, \mathcal{A}_s$ be σ-structures and let $\overline{a}$ be a tuple $(a_1, \ldots, a_n)$ from $(A_1 \cup \cdots \cup A_s)^n$. We say that $(\mathcal{A}_1, \ldots, \mathcal{A}_s, \overline{a})$ is a *model of the reduction sequence D*, in symbols:

$(\mathcal{A}_1, \ldots, \mathcal{A}_s, \overline{a}) \models D$, iff $\mu \models \beta$, where $\mu \colon X_D \to \{0, 1\}$ is the truth assignment such that for each $i \in [1, s]$ and $\delta \in \Delta_i$,

$$\mu(X_{i,\delta}) \ := \ \begin{cases} 1, & \text{if } \mathrm{free}(\delta) \subseteq \overline{x}_i \ \text{ and } \ (\mathcal{A}_i, \overline{a}_i) \models \delta(\overline{x}_i), \\ 0, & \text{otherwise.} \end{cases}$$

Here, $\overline{a}_i$ is the subsequence of $\overline{a}$ induced by all $a_j \in \overline{a} \cap A_i$ and $\overline{x}_i$ is the subsequence of $\overline{x}$ induced by all x_j such that $a_j \in \overline{a} \cap A_i$.

Let $\mathfrak{C}$ be a class of σ-structures and let $\varphi(\overline{x})$ be an $\mathrm{FO}(\sigma_s)$-formula. An s-reduction sequence D over $\overline{x}$ is an s-*disjoint decomposition for* $\varphi(\overline{x})$ *on* $\mathfrak{C}$ (for short: *disjoint decomposition for* $\varphi(\overline{x})$) if for every s-disjoint sum $\mathcal{A} = \mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_s$ of structures $\mathcal{A}_1, \ldots, \mathcal{A}_s \in \mathfrak{C}$ and all tuples $\overline{a} \in A^n$,

$$(\mathcal{A}, \overline{a}) \quad \models \quad \varphi(\overline{x}) \qquad \text{iff} \qquad (\mathcal{A}_1, \ldots, \mathcal{A}_s, \pi(\overline{a})) \quad \models \quad D.$$

Here, π is the *mapping* of the disjoint sum $\mathcal{A}$, i.e., the mapping of the underlying disjoint union of $\mathcal{A}$'s component structures.

Intuitively, an s-disjoint decomposition for an $\mathrm{FO}(\sigma_s)$-formula $\varphi(\overline{x})$ is a Boolean combination of $\mathrm{FO}(\sigma)$-formulas from sets Δ_i, $i \in [1, s]$. This Boolean combination is equivalent to φ on every s-disjoint sum $\mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_s$ and, for each $i \in [1, s]$, every $\mathrm{FO}(\sigma)$-formula from Δ_i is only interpreted over the component $\mathcal{A}_i$ and with its free variables assigned to elements from $\mathcal{A}_i$.

4.2. An upper bound.

This section's main result (Theorem 4.2) provides a 3-fold exponential algorithm that computes a disjoint decomposition for an input $\mathrm{FO}(\sigma_s)$-formula $\varphi(\overline{x})$ on $\mathfrak{C}_\nu$. The algorithm proceeds as follows: First, $\varphi(\overline{x})$ is turned in 3-fold exponential time into a Boolean combination $\varphi^H(\overline{x})$ of so-called *Hanf-formulas* (which will be defined below) such that the formula $\varphi^H(\overline{x})$ is equivalent to $\varphi(\overline{x})$ on disjoint sums of structures from $\mathfrak{C}_\nu$. To achieve this, we apply an algorithm by Bollig and Kuske [BK12]. In a second step, for each of the Hanf-formulas occurring in $\varphi^H(\overline{x})$, a disjoint decomposition is computed in linear time (see Lemma 4.3). Finally, these disjoint decompositions are combined into a disjoint decomposition $D = (\Delta_1, \ldots, \Delta_s, \beta)$ for $\varphi(\overline{x})$ on $\mathfrak{C}_\nu$. In particular, also the formulas in Δ_i are *Hanf-formulas*.

A *Hanf-formula* with $n \geqslant 0$ free variables $\overline{x}$ is a formula of the form $\exists^{\geqslant k} y \ sph_\tau(\overline{x}, y)$ where τ is the isomorphism type of a finite r-sphere (for an $r \geqslant 0$) with $n+1$ centres.

Here, for a number $k \geqslant 1$ and a formula $\varphi(\overline{x}, y)$ we write

$$\exists^{\geqslant k} y \ \varphi(\overline{x}, y)$$

as a shorthand for the formula

$$\exists y_1 \cdots \exists y_k \ \Big(\bigwedge_{1 \leqslant i < j \leqslant k} \neg y_i = y_j \ \wedge \ \forall y \ \big(\bigvee_{1 \leqslant i \leqslant k} y = y_i \ \to \varphi(\overline{x}, y) \big) \Big).$$

Note that, given k, y and φ, this formula can be constructed in time $O(k^2 + \|\varphi\|)$.

In [BK12], Bollig and Kuske provided a 3-fold exponential algorithm that transforms a given $\mathrm{FO}(\sigma')$-formula ψ, for a finite relational signature σ', into a formula in Hanf normal form that is equivalent to ψ on all σ'-structures of degree at most d (for $d \geqslant 1$). Here, we apply their result in the slightly more general setting of ν-bounded structures. In this setting, their proof yields the following (a proof can be found in the full version of [HKS13]).

Theorem 4.1 ([BK12, HKS13]). *Let σ' be a finite relational signature, let $\nu\colon \mathbb{N} \to \mathbb{N}$ be a time-constructible and strictly increasing function, and let $\mathfrak{C}'_\nu$ be the class of all ν-bounded σ'-structures. There is an algorithm which transforms an input $\mathrm{FO}(\sigma')$-formula φ of quantifier rank $q \geqslant 0$ in time*

$$2^{(\|\varphi\|\cdot\nu(4^q))^{O(\|\sigma'\|)}} \tag{4.1}$$

into a $\mathfrak{C}'_\nu$-equivalent formula φ^H in Hanf normal form.

Moreover, each Hanf-formula occurring in φ^H is of the form $\exists^{\geqslant k} y\; sph_\tau(\overline{x}, y)$, where k is at most $\|\varphi\| \cdot (q{+}1) \cdot \nu(4^q)$, τ is a ν-bounded r-sphere of radius $r \leqslant 4^q$, and $|\overline{x}| \leqslant |free(\varphi)|$. $\square$

In the rest of this section we abbreviate the upper bound $2^{(\|\varphi\|\cdot\nu(4^q))^{O(\|\sigma'\|)}}$ on the time complexity for the construction of Hanf normal forms from Theorem 4.1 by the expression $H_\nu(\|\varphi\|, 4^q, \|\sigma'\|)$.

With these preparations, this section's main result can now be stated as follows:

Theorem 4.2. *Let $s \geqslant 1$. There is an algorithm which, given an input $\mathrm{FO}(\sigma_s)$-formula $\varphi(\overline{x})$ of quantifier rank $q \geqslant 0$, constructs in time*

$$H_\nu(\|\varphi\|, 4^q, \|\sigma\| + s) \qquad \left(= \quad 2^{(\|\varphi\|\cdot\nu(4^q))^{O(\|\sigma\|+s)}} \right) \tag{4.2}$$

a disjoint decomposition $(\Delta_1, \ldots, \Delta_s, \beta)$ for $\varphi(\overline{x})$ on $\mathfrak{C}_\nu$. Furthermore, the sets $\Delta_1, \ldots, \Delta_s$ contain only Hanf-formulas over the signature σ.

Thus, if the function ν is exponential and φ is an $\mathrm{FO}(\sigma_s)$-formula, where σ_s consists of exactly the relation symbols that occur in φ, then φ can be decomposed in 3-fold exponential time — e.g., if $\nu = \nu_d$ for $d \geqslant 3$, then φ can be decomposed in time

$$2^{d^{2^{O(\|\varphi\|)}}}.$$

If ν is polynomial, then the transformation requires only 2-fold exponential time, i.e. time

$$2^{2^{O(\|\varphi\|^2)}}.$$

The remainder of this subsection is devoted to the proof of Theorem 4.2. The first step of the algorithm employs the algorithm of Theorem 4.1 to transform an $\mathrm{FO}(\sigma)$-formula φ into a $\mathfrak{C}_\nu$-equivalent formula φ^H in Hanf normal form.

For the second step of our algorithm, the following lemma provides an algorithm which constructs, given a Hanf-formula, a disjoint decomposition for it; not only on ν-bounded σ-structures but on *all* σ-structures.

Lemma 4.3. *Let $s \geqslant 1$. There is an algorithm which, given an input Hanf-formula $\psi(\overline{x})$ of the form $\exists^{\geqslant k} x_{n+1}\; sph_\tau(\overline{x}, x_{n+1})$ with $n \geqslant 0$ free variables, where $k \geqslant 1$ and, for an $r \geqslant 0$, τ is an r-sphere with $n{+}1$ centres over the signature σ_s, constructs in time*

$$O(s + \|\psi\|)$$

a disjoint decomposition $(\Delta_1, \ldots, \Delta_s, \beta)$ for $\psi(\overline{x})$ on the class of all σ-structures.

Furthermore, the sets $\Delta_1, \ldots, \Delta_s$ contain only Hanf-formulas over σ, and each Δ_i consists of at most one formula.

Proof. Let $n, r \geqslant 0$, $k \geqslant 1$, let $\overline{x} = (x_1, \ldots, x_n)$, and let $\tau = (\mathcal{T}, \overline{c})$ be an r-sphere with $n+1$ centres $\overline{c} = (c_1, \ldots, c_{n+1})$, i.e., $\mathcal{T}$ is a σ_s-structure whose elements have distance at most r to $\overline{c}$.

Recall that $\mathcal{T}$ is an induced substructure of an s-disjoint sum iff

(1) every element in the universe of $\mathcal{T}$ is contained in exactly one of the sets $P_1^{\mathcal{T}}, \ldots, P_s^{\mathcal{T}}$, and

(2) if there is an edge between two nodes a, b in the Gaifman graph of $\mathcal{T}$, then there is an $i \in [1, s]$ such that $a, b \in P_i^{\mathcal{T}}$.

We distinguish whether or not $\mathcal{T}$ is an induced substructure of an s-disjoint sum.

Case 1: $\mathcal{T}$ is not an induced substructure of an s-disjoint sum. Then, there is no s-disjoint sum $\mathcal{A}$ with a tuple $\overline{a} \in A^{n+1}$ such that $\mathcal{N}_r^{\mathcal{A}}(\overline{a}) \cong \tau$. Therefore, the algorithm can output an arbitrary unsatisfiable s-disjoint decomposition.

Case 2: $\mathcal{T}$ is an induced substructure of an s-disjoint sum. In the following, we first describe the construction of an s-disjoint decomposition for $\psi(\overline{x})$. Afterwards we show the correctness of this construction and give an analysis of its running time.

Since $\mathcal{T}$ is an induced substructure of an s-disjoint sum, for each $i \in [1, n+1]$, there is a $j \in [1, s]$ such that c_i as well as all elements in $N_r^{\mathcal{T}}(c_i)$ belong to $P_j^{\mathcal{T}}$. For each $i \in [1, s]$, let $\overline{c}_i$ be the subsequence of $\overline{c}$ induced by all $c_j \in \overline{c} \cap P_i^{\mathcal{T}}$, and let $\overline{x}_i$ be the subsequence of $\overline{x}$ induced by all x_j such that $c_j \in \overline{c} \cap P_i^{\mathcal{T}}$. W.l.o.g., $c_{n+1} \in P_s^{\mathcal{T}}$.

The algorithm decomposes $\psi(\overline{x})$ into a disjoint decomposition $D := (\Delta_1, \ldots, \Delta_s, \beta)$, which is defined as follows:

- For each $i \in [1, s]$ where $\overline{c} \cap P_i^{\mathcal{T}} = \emptyset$, let $\Delta_i := \emptyset$.
- For each $i \in [1, s]$ where $\overline{c} \cap P_i^{\mathcal{T}} \neq \emptyset$, let $\Delta_i := \{\delta_i(\overline{x}_i)\}$, where the FO$(\sigma)$-formula δ_i will be defined later on.
- β is the propositional formula

$$\bigwedge \{ X_{i, \delta_i} : i \in [1, s] \text{ such that } \overline{c} \cap P_i^{\mathcal{T}} \neq \emptyset \}$$

It remains to define the formulas $\delta_i(\overline{x}_i)$. Let $\mathcal{T}_i := \mathcal{N}_r^{\mathcal{T}}(\overline{c}_i)$ and let τ_i be the r-sphere $(\mathcal{T}_i, \overline{c}_i)$. Let $\mathcal{T}_i' := (\mathcal{T}_i)_{|\sigma}$ be the σ-reduct of $\mathcal{T}_i$, and let τ_i' be the r-sphere $(\mathcal{T}_i', \overline{c}_i)$. Recall that we assume that $c_{n+1} \in P_s^{\mathcal{T}}$. For $i = s$ we let

$$\delta_s(\overline{x}_s) \ := \ \exists^{\geqslant k} x_{n+1} \ sph_{\tau_s'}(\overline{x}_s, x_{n+1}).$$

For each $i \in [1, s-1]$, i.e., for each i where $c_{n+1} \notin P_i^{\mathcal{T}}$, we define the formula $\delta_i(\overline{x}_i)$ by

$$\delta_i(\overline{x}_i) \ := \ \exists^{\geqslant 1} x_{n+1} \ sph_{(\mathcal{T}_i', \overline{c}_i c)}(\overline{x}_i, x_{n+1}),$$

where c is an arbitrary element from the tuple $\overline{c}_i$. Note that $\delta_i(\overline{x}_i)$ is a Hanf-formula of signature σ and equivalent to the formula $sph_{\tau_i'}(\overline{x}_i)$.

Claim 4.4. $(\Delta_1, \ldots, \Delta_s, \beta)$ is a disjoint decomposition for $\psi(\overline{x})$ on the class of all σ-structures.

Proof of Claim 4.4. Let $\mathcal{A} = \mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_s$ be the disjoint sum of σ-structures $\mathcal{A}_1, \ldots, \mathcal{A}_s$, let π be the mapping of $\mathcal{A}$, and let $\overline{a} = (a_1, \ldots, a_n) \in A^n$. For every $i \in [1, s]$, let $\overline{a}_i$ be the

subsequence of $\overline{a}$ induced by all $a_j \in \overline{a} \cap P_i^{\mathcal{A}}$. By definition of a disjoint decomposition, we have to show that

$$(\mathcal{A}, \overline{a}) \quad \models \quad \exists^{\geqslant k} x_{n+1} \ sph_{(\mathcal{T}, \overline{c})}(\overline{x}, x_{n+1}) \qquad (= \psi(\overline{x}))$$

iff $(\mathcal{A}, \overline{a})$ is a model of $(\Delta_1, \ldots, \Delta_s, \beta)$. That is, iff for all $i \in [1, s]$ with $\overline{c} \cap P_i^{\mathcal{T}} \neq \emptyset$, the conditions (a) and (b) hold, where

(a) Each free variable of δ_i is interpreted by an element from $P_i^{\mathcal{T}}$, i.e., by definition of δ_i, for each $j \in [1, n]$, $N_r^{\mathcal{A}}(a_j) \subseteq P_i^{\mathcal{A}}$ iff $c_j \in P_i^{\mathcal{T}}$.
(b) $\mathcal{A}_i \models \delta_i[\pi(\overline{a}_i)]$.

Suppose that $(\mathcal{A}, \overline{a})$ satisfies $\psi(\overline{x})$ and let $i \in [1, s]$ such that $\overline{c} \cap P_i^{\mathcal{T}} \neq \emptyset$. Thus, for every $j \in [1, n]$, $\mathcal{N}_r^{\mathcal{A}}(a_j) \cong \mathcal{N}_r^{\mathcal{T}}(c_j)$. Since $\mathcal{T}$ is an induced substructure of a disjoint sum, this implies Condition (a). Furthermore, there are at least k elements $a \in A$ such that $\mathcal{N}_r^{\mathcal{A}}(a) \cong \mathcal{N}_r^{\mathcal{T}}(c_{n+1})$. Since $c_{n+1} \in P_s^{\mathcal{T}}$ it follows that $N_r^{\mathcal{A}}(a)$ is a subset of $P_s^{\mathcal{A}}$.

Recall that $\mathcal{T}_i'$ is the σ-reduct of $\mathcal{T}_i$ and that $\mathcal{N}_r^{\mathcal{A}_i|\sigma}(\overline{a}_i)$ is isomorphic to $\mathcal{N}_r^{\mathcal{A}_i}(\pi(\overline{a}_i))$. Hence, also Condition (b) holds as $\mathcal{A}_i$ satisfies $\delta_i[\pi(\overline{a}_i)]$ by definition of the corresponding formula $\delta_i(\overline{x}_i)$.

For the other direction, suppose that the conditions (a) and (b) hold for all $i \in [1, s]$ with $\overline{c} \cap P_i^{\mathcal{T}} \neq \emptyset$. Then, for each $i \in [1, s-1]$, Condition (b) and the definition of δ_i imply that $\mathcal{N}_r^{\mathcal{A}_i}(\pi(\overline{a}_i)) \cong \tau_i'$. Furthermore, by Condition (b) and the definition of δ_s, there are at least k elements $a \in A_s$ such that $\mathcal{N}_r^{\mathcal{A}_s}(\pi(\overline{a}_s a)) \cong \tau_s'$.

Recall that $\mathcal{A}$ is the disjoint sum of $\mathcal{A}_1, \ldots, \mathcal{A}_s$. Therefore, the r-neighbourhoods of all tuples $\overline{a}_i$ and $\overline{a}_j$, for distinct $i, j \in [1, s]$, are disjoint and there are no edges in the Gaifman graph of $\mathcal{A}$ between them. Hence, for each tuple $\overline{a}$ and each element a as chosen above, $\mathcal{A} \models sph_\tau[\overline{a}, a]$. Consequently, $(\mathcal{A}, \overline{a})$ satisfies $\psi(\overline{x})$. This completes the proof of Claim 4.4. $\qquad\square$

Claim 4.5. The disjoint decomposition $(\Delta_1, \ldots, \Delta_s, \beta)$ can be computed in time $O(s + \|\psi\|)$.

Proof of Claim 4.5. Let $\psi(\overline{x}) := \exists^{\geqslant k} x_{n+1} \ sph_\tau(\overline{x}, x_{n+1})$ be a Hanf-formula of signature σ_s, where $k \geqslant 1$, $n, r \geqslant 0$, and $\tau = (\mathcal{T}, \overline{c})$ is an r-sphere with $n+1$ centres. On input of $\psi(\overline{x})$ our algorithm proceeds as follows:

(1) Decide if $\mathcal{T}$ is an induced substructure of an s-disjoint sum. If yes, continue to Step (2). Otherwise, return an unsatisfiable s-disjoint decomposition.
(2) Construct the propositional formula β.
(3) For each $i \in [1, s]$, construct $\delta_i(\overline{x}_i)$.

For Step (1), note that the structure $\mathcal{T}$ can easily be reconstructed from the formula sph_τ in time $O(\|sph_\tau\|)$. The algorithm performs two passes over the structure $\mathcal{T}$. In the first pass, it remembers occurrences of the elements of the universe of $\mathcal{T}$ in the sets $P_1^{\mathcal{T}}, \ldots, P_s^{\mathcal{T}}$. This way it verifies that every element of $\mathcal{T}$ is in exactly one of these sets. If this is the case, the algorithm continues with a pass over the other relations of $\mathcal{T}$ and checks that no two elements from different sets $P_1^{\mathcal{T}}, \ldots, P_s^{\mathcal{T}}$ occur in a tuple of these relations. Both passes take time $O(\|\mathcal{T}\|) \subseteq O(\|sph_\tau\|)$.

Step (2) takes time $O(s)$, since the mapping of the constants $\overline{c}$ to the sets $P_1^{\mathcal{T}}, \ldots, P_s^{\mathcal{T}}$ is already known as a byproduct of Step (1).

In Step (3), the mapping of the elements of $\mathcal{T}$ to the sets $P_1^{\mathcal{T}}, \ldots, P_s^{\mathcal{T}}$, gathered in Step (1), helps with constructing the formulas $sph_{\tau_i'}$, for each $i \in [1, s]$, within a single pass over the formula sph_τ. Therefore, the construction of all the formulas $\delta_i(\overline{x}_i)$ takes time $O(\|\psi\|)$.

Altogether, the algorithm takes time $O(s + \|\psi\|)$. This completes the proof of Claim 4.5 and Lemma 4.3. $\qquad\square$

By using the algorithm for the construction of Hanf normal forms by Bollig and Kuske (Theorem 4.1) and Lemma 4.3, we can now prove Theorem 4.2.

Proof of Theorem 4.2. Given an input $FO(\sigma_s)$-formula $\varphi(\overline{x})$ of quantifier rank $q \geqslant 0$ and with $n \geqslant 0$ free variables, our algorithm proceeds as follows:

(1) We employ Theorem 4.1 to transform $\varphi(\overline{x})$ into a formula $\varphi^H(\overline{x})$ in Hanf normal form that is equivalent to $\varphi(\overline{x})$ on the class of all ν-bounded σ_s-structures. This takes time $H_\nu(\|\varphi\|, 4^q, \|\sigma\| + s)$. Furthermore, the Hanf-formulas occurring in this formula are of the form $\exists^{\geqslant k} x_{n+1}\, sph_\tau(\overline{x}, x_{n+1})$, where k is at most $\|\varphi\| \cdot (q+1) \cdot \nu(4^q)$, and τ is a ν-bounded r-sphere of radius $r \leqslant 4^q$ with at most $n+1$ centres. Since $q+2 \leqslant \|\varphi\|$, we have that

$$k+1 \quad \leqslant \quad \|\varphi\| \cdot (q+1) \cdot \nu(4^q) + 1 \quad \leqslant \quad \|\varphi\|^2 \cdot \nu(4^q).$$

(2) For a suitable $L \geqslant 1$, let $\psi_1(\overline{x}), \ldots, \psi_L(\overline{x})$ be the Hanf-formulas occurring in $\varphi^H(\overline{x})$. For each $\ell \in [1, L]$, we use the construction described in the proof of Lemma 4.3 to compute a disjoint decomposition $(\Delta_{1,\ell}, \ldots, \Delta_{s,\ell}, \beta_\ell)$ for the Hanf-formula $\psi_\ell(\overline{x})$. Note that, for each $\ell \in [1, L]$ and each $i \in [1, s]$, $\Delta_{i,\ell}$ is either empty or consists of just a single Hanf-formula $\delta_{i,\ell}$.

For each Hanf-formula $\exists^{\geqslant k} x_{n+1}\, sph_\tau(\overline{x}, x_{n+1})$ occurring in $\varphi^H(\overline{x})$, this takes time

$$O(s + k^2 + \|sph_\tau\|) \quad \subseteq \quad (\|\varphi\| \cdot \nu(4^q))^{O(\|\sigma\| + s)}.$$

Since $L < \|\varphi^H\|$ and $\|\varphi^H\| \in H_\nu(\|\varphi\|, 4^q, \|\sigma\| + s)$, Step (2) takes time

$$L \cdot (\|\varphi\| \cdot \nu(4^q))^{O(\|\sigma\| + s)} \quad \subseteq \quad H_\nu(\|\varphi\|, 4^q, \|\sigma\| + s).$$

(3) We output the disjoint decomposition $D := (\Delta_1, \ldots, \Delta_s, \beta)$, where $\Delta_i := \Delta_{i,1} \cup \cdots \cup \Delta_{i,L}$ for each $i \in [1, s]$, and the propositional formula β is obtained from $\varphi^H(\overline{x})$ by replacing each Hanf-formula $\psi_\ell(\overline{x})$, for every $\ell \in [1, L]$, by the propositional formula β_ℓ. Again, this takes time $H_\nu(\|\varphi\|, 4^q, \|\sigma\| + s)$.

A ltogether, the running time of the algorithm is bounded by $H_\nu(\|\varphi\|, 4^q, \|\sigma\| + s)$. It is straightforward to verify that D is indeed a disjoint decomposition for $\varphi(\overline{x})$ on $\mathfrak{C}_\nu$, that is, for all structures $\mathcal{A}_1, \ldots, \mathcal{A}_s \in \mathfrak{C}_\nu$, their disjoint sum $\mathcal{A} = \mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_s$ with mapping π, and all $\overline{a} \in A^n$, we have $(\mathcal{A}, \overline{a}) \models \varphi(\overline{x})$ iff $(\mathcal{A}_1, \ldots, \mathcal{A}_s, \pi(\overline{a})) \models D$. This completes the proof of Theorem 4.2. $\qquad\square$

4.3. **Generalised decompositions.** In this section, we turn our attention from decompositions on disjoint sums of structures to decompositions speaking about more general products of structures. We show how to transfer the algorithm of Theorem 4.2 to an algorithm that produces decompositions on composite structures obtained by applying *transductions*[2] to disjoint sums.

Subsection 4.3.1 provides the necessary background on transductions. Subsection 4.3.2 presents this section's main result, Corollary 4.8, which lifts Theorem 4.2 from disjoint sums to more general decompositions defined via transductions. Subsection 4.3.3 presents

---

[2]also knows as *first-order interpretations*, cf. [EF99]

an application of the result to the particular case of direct products (also called tensor products or cartesian products).

4.3.1. *Transductions.* We consider two fixed finite relational signatures σ and τ.

Let $t \geqslant 1$ and let $\theta(x_1, \ldots, x_t)$ be an FO(σ)-formula with t free variables. Furthermore, assume that, for each $R \in \tau$ of arity $r := \mathrm{ar}(R)$, $\theta_R(\overline{y}_1 \ldots \overline{y}_r)$ is an FO(σ)-formula with $r \cdot t$ free variables from tuples $\overline{y}_i := (y_{i,1}, \ldots, y_{i,t})$, for each $i \in [1, r]$. Then, the tuple $(\theta, (\theta_R)_{R \in \tau})$ is called a *t-transduction from τ to σ* (or *transduction*, if the parameters are given by the context).

Assume now that $\Theta = (\theta, (\theta_R)_{R \in \tau})$ is a t-transduction from τ to σ as described above. For every σ-structure $\mathcal{A} = (A, (R^{\mathcal{A}})_{R \in \sigma})$, the *application $\Theta(\mathcal{A})$ of the transduction Θ to $\mathcal{A}$* is a τ-structure $\mathcal{B} = (B, (R^{\mathcal{B}})_{R \in \tau})$ whose universe B consists of exactly the tuples $(a_1, \ldots, a_t) \in A^t$ such that $\mathcal{A} \models \theta[a_1, \ldots, a_t]$, and where each relation symbol $R \in \tau$ with arity r is interpreted by the set of all tuples $(\overline{a}_1, \ldots, \overline{a}_r)$ in B^r, such that $\mathcal{A} \models \theta_R[\overline{a}_1 \ldots \overline{a}_r]$.

On the other hand, a t-transduction Θ from τ to σ can be applied to FO(τ)-formulas. For every FO(τ)-formula $\varphi(\overline{x})$ with $n \geqslant 0$ free variables $\overline{x} = (x_1, \ldots, x_n)$, the *application $\Theta(\varphi)$ of the transduction Θ to $\varphi(\overline{x})$* is an FO(σ)-formula $\psi(\overline{x}_1 \ldots \overline{x}_n)$ with $n{\cdot}t$ new free variables from variable tuples $\overline{x}_i := (x_{i,1}, \ldots, x_{i,t})$, $i \in [1, n]$, that is defined inductively as follows:

- If $\varphi(\overline{x}) = R(x_1, \ldots, x_r)$, for a relation symbol $R \in \tau$ with arity r, then

$$\Theta(\varphi)(\overline{x}_1 \ldots \overline{x}_r) \ := \ \bigwedge_{i=1}^{r} \theta(\overline{x}_i) \ \wedge \ \theta_R(\overline{x}_1 \ldots \overline{x}_r).$$

- If φ is of the form $x_1 = x_2$ then

$$\Theta(\varphi)(\overline{x}_1 \overline{x}_2) \ := \ \theta(\overline{x}_1) \ \wedge \ \bigwedge_{j=1}^{t} x_{1,j} = x_{2,j}.$$

- For the Boolean connectives the translation distributes, i.e. $\Theta(\neg\varphi) := \neg\Theta(\varphi)$ and, for each $\star \in \{\wedge, \vee, \rightarrow, \leftrightarrow\}$, we have $\Theta(\varphi \star \psi) := \Theta(\varphi) \star \Theta(\psi)$.
- If $\varphi(\overline{x}) = \exists y\, \psi(\overline{x}, y)$ then let $\overline{y} = (y_1, \ldots, y_t)$ be a tuple of new variables and let

$$\Theta(\varphi)(\overline{x}_1 \ldots \overline{x}_n) \ := \ \exists \overline{y} \, \big(\theta(\overline{y}) \ \wedge \ \Theta(\psi)(\overline{x}_1 \ldots \overline{x}_n \overline{y}) \big).$$

- Finally, for $\varphi(\overline{x}) = \forall y\, \psi(\overline{x}, y)$, let $\Theta(\varphi) := \Theta(\neg \exists y\, \neg\psi)$.

Obviously, the application of a fixed transduction to an input FO-formula can be carried out in time linear in the size of the input formula. The following lemma makes this precise and gives a bound on the quantifier rank of the resulting formula.

Lemma 4.6. *Let $t \geqslant 1$ and let $\Theta = (\theta, (\theta_R)_{R \in \tau})$ be a t-transduction from τ to σ. Let k and p be the maximum size and the maximum quantifier rank of the formulas θ and θ_R, for each $R \in \tau$. For every FO(τ)-formula φ, the FO(σ)-formula $\Theta(\varphi)$ can be constructed in time $O((t{+}k) \cdot \|\varphi\|)$. Furthermore, if q is the quantifier rank of φ then $\Theta(\varphi)$ has quantifier rank at most $t{\cdot}q + p$.*

Proof. Let φ be an FO(τ)-formula of quantifier rank $q \geqslant 0$. By induction on the shape of φ, we show that there is a suitable number $c \geqslant 1$, that has to be greater than some values obtained during the course of the induction, such that $\|\Theta(\varphi)\| \leqslant c \cdot (t{+}k) \cdot \|\varphi\|$ and $\mathrm{qr}(\Theta(\varphi)) \leqslant t \cdot q + p$.

- If $\varphi = R(x_1, \ldots, x_r)$, for a relation symbol $R \in \tau$ with arity r, then $\|\Theta(\varphi)\| \leqslant c \cdot k \cdot (r+1)$. Of course, $r+1 \leqslant \|\varphi\|$. Therefore, $\|\Theta(\varphi)\| \leqslant c \cdot (t+k) \cdot \|\varphi\|$. Furthermore, the quantifier rank of $\Theta(\varphi)$ is at most p.

- For a Boolean combination φ, the induction step is obvious.

- Assume that $\varphi = Qy\,\psi(\overline{x}, y)$ for a quantifier $Q \in \{\exists, \forall\}$. In this case, we have that $\|\Theta(\varphi)\| \leqslant c \cdot (t+k) + \|\Theta(\psi)\|$. Since $\|\Theta(\psi)\| \leqslant c \cdot (t+k) \cdot (\|\varphi\|-1)$, it follows that $\|\Theta(\varphi)\|$ is at most $c \cdot (t+k) \cdot \|\varphi\|$. Furthermore, by construction of $\Theta(\varphi)$, we have that $\mathrm{qr}(\Theta(\varphi))$ is at most $t + \max\{p, \mathrm{qr}(\Theta(\psi))\}$. Because $\mathrm{qr}(\Theta(\psi)) \leqslant t \cdot (q-1) + p$ it follows that $\mathrm{qr}(\Theta(\varphi)) \leqslant t \cdot q + p$.

Clearly, the formula $\Theta(\varphi)$ can be constructed in time $O((t+k) \cdot \|\varphi\|)$, that is linear in the size of φ. $\qquad\square$

The following proposition relates the application of transductions to structures and formulas to each other. The proof is an immediate consequence of the definition of $\Theta(\mathcal{A})$ and $\Theta(\varphi)$ (see, e.g., [EF99]).

Proposition 4.7. *Let $t \geqslant 1$ and let $\Theta = (\theta, (\theta_R)_{R \in \sigma})$ be a t-transduction from τ to σ. For every $\mathrm{FO}(\tau)$-formula $\varphi(\overline{x})$ with $n \geqslant 0$ free variables $\overline{x} = (x_1, \ldots, x_n)$, each σ-structure $\mathcal{A}$ and each tuple $(\overline{a}_1, \ldots, \overline{a}_n) \in (\theta(\mathcal{A}))^n$, the following is true:*

$$
\begin{aligned}
\bigl(\Theta(\mathcal{A}), \overline{a}_1, \ldots, \overline{a}_n\bigr) &\models \varphi(x_1, \ldots, x_n) \\
\textit{iff} \quad \bigl(\mathcal{A}, \overline{a}_1 \ldots \overline{a}_n\bigr) &\models \Theta(\varphi)(\overline{x}_1 \ldots \overline{x}_n),
\end{aligned}
$$

where, for each $i \in [1, n]$, $\overline{x}_i := (x_{i,1}, \ldots, x_{i,t})$. $\qquad\square$

Note here the difference between the expressions $(\overline{a}_1, \ldots, \overline{a}_n)$ and $\overline{a}_1 \ldots \overline{a}_n$, where $\overline{a}_i$ denotes an arbitrary tuple (for every $i \in [1, n]$): While $(\overline{a}_1, \ldots, \overline{a}_n)$ denotes the tuple of length n whose elements are exactly the tuples $\overline{a}_i$ for $i \in [1, n]$, the expression $\overline{a}_1 \ldots \overline{a}_n$ represents the concatenation of these tuples, i.e., a tuple of length $|\overline{a}_1| + \cdots + |\overline{a}_n|$.

4.3.2. *Decompositions obtained by transductions.* This section's main result is a corollary to Theorem 4.2 and reads as follows:

Corollary 4.8. *Let σ and τ be finite relational signatures and let $\mathfrak{D}_\nu$ be the class of all ν-bounded τ-structures. Let $s, t \geqslant 1$ and let $\Theta = (\theta, (\theta_R)_{R \in \tau})$ be a t-transduction from τ to σ_s. There is an algorithm which, given an input $\mathrm{FO}(\tau)$-formula $\varphi(\overline{x})$ of quantifier rank $q \geqslant 0$ and with $n \geqslant 0$ free variables, constructs in time*

$$
H_\nu(\|\varphi\| \cdot (t+k), 4^{t \cdot q+p}, \|\sigma\| + s) \qquad \left(= \quad 2^{(\|\varphi\| \cdot (t+k) \cdot \nu(4^{t \cdot q+p}))^{O(\|\sigma\|+s)}}\right)
$$

a reduction sequence $D = (\Delta_1, \ldots, \Delta_s, \beta)$ over $\overline{x}$, such that for all structures $\mathcal{A}_1, \ldots, \mathcal{A}_s \in \mathfrak{D}_\nu$ and every tuple $(\overline{a}_1, \ldots, \overline{a}_n)$ in $(\theta(\mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_s))^n$, the following is true:

$$
\begin{aligned}
\bigl(\Theta(\mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_s), (\overline{a}_1, \ldots, \overline{a}_n)\bigr) &\models \varphi(\overline{x}) \\
\textit{iff} \quad \bigl(\mathcal{A}_1, \ldots, \mathcal{A}_s, \pi(\overline{a}_1 \ldots \overline{a}_n)\bigr) &\models D,
\end{aligned}
$$

where π is the mapping of the disjoint sum $\mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_s$.

Here, k and p are the maximum size and maximum quantifier rank of the formulas θ and θ_R, for every $R \in \sigma$. Furthermore, the sets $\Delta_1, \ldots, \Delta_s$ consist of Hanf-formulas over σ.

Proof. Let $\varphi(\overline{x})$ be an FO(τ)-formula of quantifier rank $q \geqslant 0$ and with $n \geqslant 0$ free variables $\overline{x} = (x_1, \ldots, x_n)$. The algorithm proceeds in two steps:

(1) In a first step, the algorithm applies the transduction Θ to the input formula $\varphi(\overline{x})$ to obtain an FO(σ_s)-formula $\Theta(\varphi)(\overline{x}_1 \ldots \overline{x}_n)$ where, for each $i \in [1, n]$, $\overline{x}_i$ is a tuple $(x_{i,1}, \ldots, x_{i,t})$ of free variables. From Proposition 4.7 it follows that for all σ_s-structures $\mathcal{A}$ and all $(\overline{a}_1, \ldots, \overline{a}_n)$ in $(\theta(\mathcal{A}))^n$, we have that

$$\big(\Theta(\mathcal{A}), (\overline{a}_1, \ldots, \overline{a}_n)\big) \quad\models\quad \varphi(x_1, \ldots, x_n)$$
$$\text{iff} \quad \big(\mathcal{A}, \overline{a}_1 \ldots \overline{a}_n\big) \quad\models\quad \Theta(\varphi)(\overline{x}_1 \ldots \overline{x}_n).$$

By Lemma 4.6, this step requires time $O((t+k) \cdot \|\varphi\|)$, where k is the maximum size of the FO(σ_s)-formulas θ and θ_R, for every $R \in \tau$. Also, $\|\Theta(\varphi)\| \in O((t+k) \cdot \|\varphi\|)$ and $\mathrm{qr}(\Theta(\varphi)) \leqslant t \cdot q + p$, where p is the maximum quantifier rank of the FO(σ_s)-formulas θ and θ_R, for every $R \in \tau$.

(2) In the second step, the algorithm constructs a disjoint decomposition $D = (\Delta_1, \ldots, \Delta_s, \beta)$ for $\Theta(\varphi)(\overline{x}_1 \ldots \overline{x}_n)$ on $\mathfrak{D}_\nu$. I.e., for all $\mathcal{A}_1, \ldots, \mathcal{A}_s \in \mathfrak{D}_\nu$ and the disjoint sum $\mathcal{A}$ of $\mathcal{A}_1, \ldots, \mathcal{A}_s$ with its mapping π and for each tuple $\overline{a} \in A^{n \cdot t}$, we have

$$(\mathcal{A}, \overline{a}) \quad\models\quad \Theta(\varphi)(\overline{x}_1 \ldots \overline{x}_n) \quad\text{iff}\quad (\mathcal{A}_1, \ldots, \mathcal{A}_s, \pi(\overline{a})) \quad\models\quad D.$$

By Theorem 4.2, the construction of D from $\Theta(\varphi)$ takes time

$$H_\nu(\|\Theta(\varphi)\|, 4^{\mathrm{qr}(\Theta(\varphi))}, \|\sigma\| + s) \quad\subseteq\quad H_\nu(\|\varphi\| \cdot (t+k), 4^{t \cdot q + p}, \|\sigma\| + s)$$

It follows from Step (1) and Step (2) that, for every disjoint sum $\mathcal{A}$ of structures $\mathcal{A}_1, \ldots, \mathcal{A}_s \in \mathfrak{D}_\nu$ and all $(\overline{a}_1, \ldots, \overline{a}_n)$ in $(\theta(\mathcal{A}))^n$, it holds that

$$\big(\Theta(\mathcal{A}), (\overline{a}_1, \ldots, \overline{a}_n)\big) \quad\models\quad \varphi(\overline{x})$$
$$\text{iff} \quad (\mathcal{A}_1, \ldots, \mathcal{A}_s, \pi(\overline{a}_1 \ldots \overline{a}_n)) \quad\models\quad D$$

and that construction of D from $\varphi(\overline{x})$ altogether takes time

$$H_\nu(\|\varphi\| \cdot (t+k), 4^{t \cdot q + p}, \|\sigma\| + s).$$

This concludes the proof of Corollary 4.8. $\qquad\square$

4.3.3. *Decompositions on direct products.* We exemplify the application of Corollary 4.8 with the following result on *direct products* of structures. Let $s \geqslant 1$. For σ-structures $\mathcal{A}_1, \ldots, \mathcal{A}_s$, the direct product $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_s$ is the σ-structure $\mathcal{A} = (A, (R^{\mathcal{A}})_{R \in \sigma})$, where the universe A is the set $A_1 \times \cdots \times A_s$, and for each relation symbol $R \in \sigma$ of arity $r \geqslant 1$, the relation $R^{\mathcal{A}}$ is the set of all tuples $((a_{1,1}, \ldots, a_{1,s}), \ldots, (a_{n,1}, \ldots, a_{r,s}))$ in A^r such that, for each $i \in [1, s]$, the tuple $(a_{1,i}, \ldots, a_{r,i})$ belongs to the relation $R^{\mathcal{A}_i}$.

Corollary 4.9. *Let $s \geqslant 1$. There is an algorithm which, given an input FO(σ)-formula $\varphi(\overline{x})$ of quantifier rank $q \geqslant 0$ and with $n \geqslant 0$ free variables, constructs in time*

$$H_\nu(s \cdot \|\varphi\|, 4^{s \cdot q}, \|\sigma\| + s) \qquad \left(= \quad 2^{(s \cdot \|\varphi\| \cdot \nu(4^{s \cdot q}))^{O(\|\sigma\| + s)}}\right)$$

a reduction sequence $D = (\Delta_1, \ldots, \Delta_s, \beta)$ over $\overline{x}$, such that for all $\mathcal{A}_1, \ldots, \mathcal{A}_s \in \mathfrak{C}_\nu$ and all $(\overline{a}_1, \ldots, \overline{a}_n)$ in $(A_1 \times \cdots \times A_s)^n$,

$$\big(\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_s, (\overline{a}_1, \ldots, \overline{a}_n)\big) \quad\models\quad \varphi(\overline{x})$$
$$\text{iff} \quad (\mathcal{A}_1, \ldots, \mathcal{A}_s, \overline{a}_1 \ldots \overline{a}_n) \quad\models\quad D.$$

Furthermore, the sets $\Delta_1, \ldots, \Delta_s$ consist of Hanf-formulas over σ.

Proof. Consider the following s-transduction $\Theta := (\theta, (\theta_R)_{R \in \sigma})$ from σ to σ_s with

$$\theta(x_1, \ldots, x_s) := \bigwedge_{i=1}^{s} P_i(x_i)$$

and, for each $R \in \sigma$ of arity $r := \mathrm{ar}(R)$,

$$\theta_R(\overline{x}_1 \ldots \overline{x}_r) := \bigwedge_{i=1}^{s} R(x_{1,i}, \ldots, x_{r,i}),$$

where $\overline{x}_j = (x_{j,1}, \ldots, x_{j,s})$ for each $j \in [1, r]$. The formulas of Θ correspond to the definition of the (tensor) product of σ-structures. I.e., for all σ-structures $\mathcal{A}_1, \ldots, \mathcal{A}_s$ and the disjoint sum $\mathcal{A}$ of $\mathcal{A}_1, \ldots, \mathcal{A}_s$ with mapping π,

$$\Theta(\mathcal{A}) \quad \cong \quad \mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_s$$

via the isomorphism $f : \theta(\mathcal{A}) \to (A_1 \times \cdots \times A_s)$ defined by $f(\overline{a}) := \pi(\overline{a})$, for all $\overline{a} \in \theta(\mathcal{A})$.

Therefore, by Proposition 4.7, for every $\mathrm{FO}(\sigma)$-formula $\varphi(\overline{x})$ with a tuple $\overline{x}$ of $n \geqslant 0$ free variables $(x_1, \ldots, x_n)$, the following holds: Let $\mathcal{A}_1, \ldots, \mathcal{A}_s$ be σ-structures and let $\mathcal{A}$ be the disjoint sum of $\mathcal{A}_1, \ldots, \mathcal{A}_s$ with mapping π. Let $(\overline{a}_1, \ldots, \overline{a}_n) \in (\theta(\mathcal{A}))^n$. Then,

$$\begin{aligned}
&\big(\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_s, (\pi(\overline{a}_1), \ldots, \pi(\overline{a}_n))\big) && \models \varphi(\overline{x}) \\
\text{iff } &\big(\Theta(\mathcal{A}), (\overline{a}_1, \ldots, \overline{a}_n)\big) && \models \varphi(\overline{x}) \\
\text{iff } &\big(\mathcal{A}, \overline{a}_1 \ldots \overline{a}_n\big) && \models \Theta(\varphi)(\overline{x}_1 \ldots \overline{x}_n),
\end{aligned}$$

where $\overline{x}_i := (x_{i,1}, \ldots, x_{i,s})$ for each $i \in [1, n]$.

The proof concludes with an application of Corollary 4.8 to this transduction. Note that the formulas θ and θ_R, for each $R \in \sigma$, are quantifier-free. Hence, by Corollary 4.8, it takes time

$$H_\nu(s \cdot \|\varphi\|, 4^{s \cdot q}, \|\sigma\| + s)$$

to compute a reduction sequence $D = (\Delta_1, \ldots, \Delta_s, \beta)$ over $\overline{x}$, such that for each disjoint sum $\mathcal{A}$ of structures $\mathcal{A}_1, \ldots, \mathcal{A}_s \in \mathfrak{C}_\nu$ and all $(\overline{a}_1, \ldots, \overline{a}_n) \in (\theta(\mathcal{A}))^n$,

$$\begin{aligned}
&\big(\Theta(\mathcal{A}), (\overline{a}_1, \ldots, \overline{a}_n)\big) && \models \varphi(\overline{x}) \\
\text{iff } &\big(\mathcal{A}_1, \ldots, \mathcal{A}_s, \pi(\overline{a}_1 \ldots \overline{a}_n)\big) && \models D,
\end{aligned}$$

where π is the mapping of the disjoint sum $\mathcal{A}$. Observe that a tuple $\overline{a}$ is in $\theta(\mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_s)$ iff $\pi(\overline{a})$ is in $A_1 \times \cdots \times A_s$. Hence, for all $\mathcal{A}_1, \ldots, \mathcal{A}_s \in \mathfrak{C}_\nu$ and all $(\overline{a}_1, \ldots, \overline{a}_n) \in (A_1 \times \cdots \times A_s)^n$,

$$\begin{aligned}
&\big(\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_s, (\overline{a}_1, \ldots, \overline{a}_n)\big) && \models \varphi(\overline{x}) \\
\text{iff } &\big(\mathcal{A}_1, \ldots, \mathcal{A}_s, \overline{a}_1 \ldots \overline{a}_n\big) && \models D.
\end{aligned}$$

This concludes the proof of Corollary 4.9. $\qquad \square$

Note that Corollary 4.9 also holds when the direct product of structures is replaced by other graph theoretical products, e.g., the *cartesian product* and the *strong product* of graphs (cf. [HN04]).

5. Lower Bounds.

In this section, we prove lower bounds corresponding to the upper bounds of the main results of sections 3 and 4 concerning preservation theorems and Feferman-Vaught decompositions. A key ingredient to all our lower bounds is an encoding of large numbers by bounded degree trees that allows to compare numbers by using small FO-formulas In Subsection 5.1, we introduce this encoding. The remaining subsections 5.2 and 5.3 contain our lower bounds for preservation theorems and Feferman-Vaught decompositions, respectively.

5.1. **Binary tree encodings.** We recall an encoding of numbers by *binary trees*, i.e., trees where every node has at most two children, which was already used in [HKS13]. This binary tree encoding is an adaptation of an encoding of numbers by trees of unbounded degree from Chapter 10 in [FG06].

Consider the signature $\{E\}$ that consists of just a single binary relation symbol E. A *forest* is a disjoint union of finite directed rooted trees. The *height* of a forest $\mathcal{F}$ is the length of a longest directed path in $\mathcal{F}$. For each node a of a forest $\mathcal{F}$, $\mathcal{F}_a$ is the subtree of $\mathcal{F}$ induced by all nodes reachable by a directed path from a. For a tree $\mathcal{T}$ and a number $d \geqslant 0$, we write $\mathcal{T}[\leqslant d]$ to denote the subtree of $\mathcal{T}$ induced by all nodes of $\mathcal{T}$ that are reachable from the root of $\mathcal{T}$ by a directed path of length at most d. A tree $\mathcal{B}$ is a *complete binary tree* if all leaves of $\mathcal{B}$ have the same height and every non-leaf node has exactly two children.

For numbers $i, n \in \mathbb{N}$, we write $\mathrm{bit}(i, n)$ to denote the i-th bit in the binary representation of n. I. e., $\mathrm{bit}(i, n) = 1$ iff $\lfloor \frac{n}{2^i} \rfloor$ is odd.

We define the (non-elementary) function $\mathrm{Tower} \colon \mathbb{N} \to \mathbb{N}$ by $\mathrm{Tower}(0) := 1$ and

$$\mathrm{Tower}(h) \ := \ 2^{\mathrm{Tower}(h-1)} \quad \text{for all } h \geqslant 1.$$

I. e., $\mathrm{Tower}(h)$ is a tower of $2s$ of height h.

For each $h \geqslant -1$ and $i \in [0, \mathrm{Tower}(h+3)-1]$, we define inductively a set $\mathfrak{B}_h(i)$ of binary trees that (each) encode the (binary expansion of the) number i.

$h = -1$: For $i \in [0, \mathrm{Tower}(2)-1] = \{0, 1, 2, 3\}$, the set $\mathfrak{B}_{-1}(i)$ contains exactly the binary trees that are isomorphic to the binary tree $\mathcal{B}_{-1}(i)$ depicted in Figure 2.

$h \geqslant 0$: For $i \in [0, \mathrm{Tower}(h+3)-1]$, the set $\mathfrak{B}_h(i)$ consists of all binary trees $\mathcal{B}$ that satisfy each of the following properties:

- $\mathcal{B}[\leqslant \mathrm{Tower}(h+1)-1]$ is a complete binary tree of height $\mathrm{Tower}(h+1)-1$.

- For every $j \in [0, \mathrm{Tower}(h+2)-1]$ with $\mathrm{bit}(j, i) = 1$, there is a node b of height $\mathrm{Tower}(h+1)$ in $\mathcal{B}$ such that $\mathcal{B}_b \in \mathfrak{B}_{h-1}(j)$.

- For every node b of height $\mathrm{Tower}(h+1)$ in $\mathcal{B}$, there is a $j \in [0, \mathrm{Tower}(h+2)-1]$ such that $\mathcal{B}_b \in \mathfrak{B}_{h-1}(j)$ and $\mathrm{bit}(j, i) = 1$.
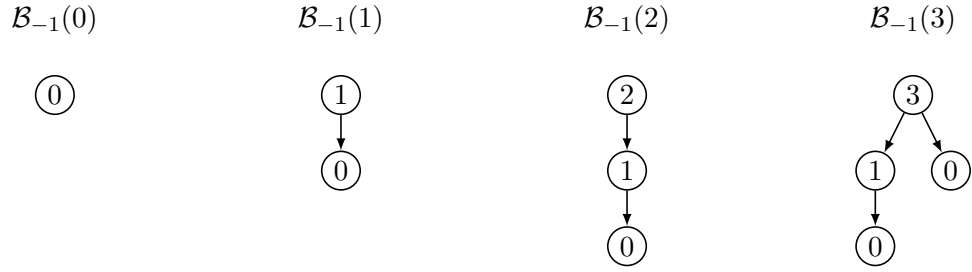
Figure 2: The binary trees $\mathcal{B}_{-1}(0)$, $\mathcal{B}_{-1}(1)$, $\mathcal{B}_{-1}(2)$, and $\mathcal{B}_{-1}(3)$ are tree encodings for the numbers 0, 1, 2, and 3, respectively, as defined in [FG06]. Note that the numbers depicted within the nodes are *not* part of the tree encoding; they are just indicated here to illustrate which number is encoded by the subtree starting at the respective node.

Figure 3: A binary tree from the set $\mathfrak{B}_1(42)$, i.e., a binary tree encoding of the number 42 with parameter 1. The numbers depicted within some of the nodes are *not* part of the tree encoding; they are just indicated here to illustrate which number is encoded by the subtree starting at the respective node.
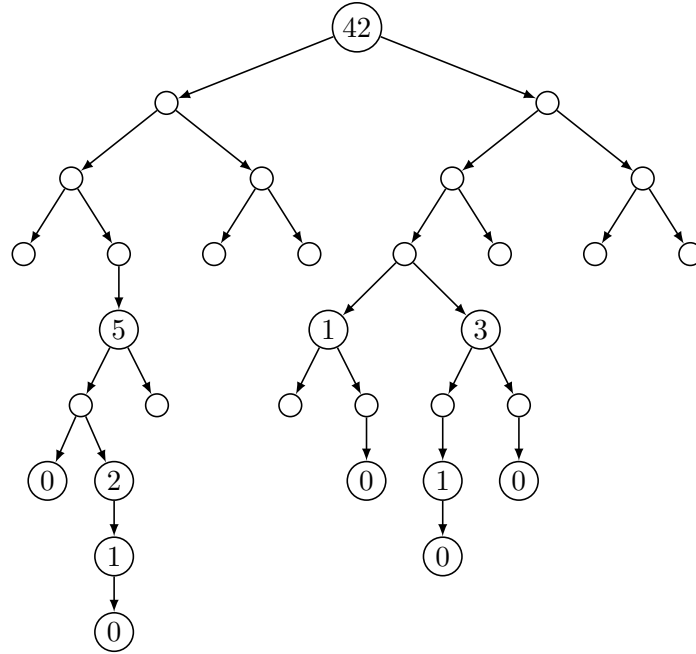
Each tree in $\mathfrak{B}_h(i)$ is called a *binary tree encoding of i with parameter h*. An example of a binary tree encoding of $i = 42$ with parameter $h = 1$ is depicted in Figure 3.

An induction on the parameter h shows that every number $i \in [0, \mathrm{Tower}(h+3)-1]$ has at least one binary tree encoding with parameter h. The following easy observation gives an upper bound on the height of binary tree encodings.

Lemma 5.1. *For each $h > 1$ and every $i \in [0, Tower(h+3) - 1]$, each binary tree in $\mathfrak{B}_h(i)$ has height less than $2 \cdot Tower(h+1)$.*

Proof. Let $h > 1$ and $i \in [0, Tower(h+3) - 1]$. By definition of $\mathfrak{B}_h(i)$ (see also Figure 3) it is easy to convince oneself that each binary tree in $\mathfrak{B}_h(i)$ has height less than

$$2 \; + \; \sum_{k=0}^{h} Tower(k+1) \quad < \quad 2 \cdot Tower(h+1).$$

The latter inequality can be shown by a straightforward induction.   □

An adaptation of Lemma 3.2 (see also Chapter 10 in [FG06]), Lemma 3.3, and Lemma 3.4 of [DGKS07] allows to express arithmetic relations between binary tree encodings of numbers by "small" FO(E)-formulas:

Lemma 5.2. *For $h \geqslant -1$, there are FO(E)-formulas $ench_h(x)$, $min_h(x)$, $eq_h(x,y)$, $less_h(x,y)$, $succ_h(x,y)$ and $max_h(x)$ of size $O(Tower(h))$ (for $h \geqslant 0$) such that for each binary forest $\mathcal{F}$ and all nodes $a, b \in F$,*

$$\begin{aligned}
\mathcal{F} \;&\models\; ench_h[a] &&\quad \textit{iff} \quad \mathcal{F}_a \in \mathfrak{B}_h(i) \\
& && \quad \textit{for an } i \in [0, Tower(h+3)-1],
\end{aligned}$$

and if there are $i, j \in [0, Tower(h+3)-1]$ such that $\mathcal{F}_a \in \mathfrak{B}_h(i)$ and $\mathcal{F}_b \in \mathfrak{B}_h(j)$, then

$$\begin{aligned}
\mathcal{F} \;&\models\; min_h[a] &&\quad \textit{iff} \quad i = 0, \\
\mathcal{F} \;&\models\; eq_h[a,b] &&\quad \textit{iff} \quad i = j, \\
\mathcal{F} \;&\models\; less_h[a,b] &&\quad \textit{iff} \quad i < j, \\
\mathcal{F} \;&\models\; succ_h[a,b] &&\quad \textit{iff} \quad i+1 = j, \quad \textit{and} \\
\mathcal{F} \;&\models\; max_h[a] &&\quad \textit{iff} \quad i = Tower(h+3) - 1.
\end{aligned}$$

Proof. For each $d \geqslant 0$, there is an FO(E)-formula $\delta_{\leqslant d}(x,y)$ of size $O(\log d)$ expressing in a binary forest that there is a directed path of length at most d from node x to node y. For $d = 0$ and $d = 1$, the formulas $\delta_{\leqslant 0}(x,y)$ and $\delta_{\leqslant 1}(x,y)$ can be chosen as $x{=}y$ and $x{=}y \vee E(x,y)$, respectively. For $d \geqslant 1$, define

$$\delta_{\leqslant 2d}(x,y) \;:=\; \exists z \forall x' \forall y' \Big(\big((x'{=}x \wedge y'{=}z) \; \vee \; (x'{=}z \wedge y'{=}y) \big) \; \rightarrow \; \delta_{\leqslant d}(x',y') \Big)$$

$$\delta_{\leqslant 2d+1}(x,y) \;:=\; \exists z \big(\delta_{\leqslant 1}(x,z) \; \wedge \; \delta_{\leqslant 2d}(z,y) \big).$$

It is easy to check that $\delta_{\leqslant d}$ has size in $O(\log d)$.

Furthermore, let $\delta_{=d}(x,y) := \delta_{\leqslant d}(x,y) \wedge \neg \delta_{\leqslant d-1}(x,y)$ be the FO(E)-formula expressing that there is a directed path of length exactly d from node x to node y. Of course, the size of $\delta_{=d}(x,y)$ is linear in $\delta_{\leqslant d}$ and therefore also in $O(\log d)$.

Consider for each $d \geqslant 0$ the following FO(E)-formula:

$$\begin{aligned}
\gamma_d(x) \;:=\; & \exists y \; \delta_{=d}(x,y) \; \wedge \\
& \forall y \; \big(\delta_{\leqslant d-1}(x,y) \; \rightarrow \\
& \qquad \exists z_0 \, \exists z_1 \; (E(y,z_0) \wedge E(y,z_1) \wedge \neg z_0{=}z_1) \big).
\end{aligned}$$

For each $d \geqslant 0$, the formula $\gamma_d(x)$ states that there is a node y that is reachable from x by a path of length exactly d and that every node y that is reachable from x by a path of length

less than d has exactly two children. Hence, for every binary forest $\mathcal{F}$ and each $a \in F$, we have that $\mathcal{F} \models \gamma_d[a]$ iff $\mathcal{F}_a[\leqslant d]$ is a complete binary tree of depth d. Since $\delta_{=d}$ and $\delta_{\leqslant d-1}$ have size in $O(\log d)$, also γ_d has size in $O(\log d)$.

Furthermore, for each $i \in \{0, 1, 2, 3\}$, it is straightforward to define a formula $enc_{-1,i}(x)$ that is satisfied if x is the root node of a binary tree that is isomorphic to the binary tree $\mathcal{B}_{-1}(i)$ depicted in Figure 2.

With this preparation, the $\mathrm{FO}(E)$-formula $enc_h(x)$ can be defined corresponding to the definition of binary tree encodings with parameter h. Let

$$enc_{-1}(x) \quad := \quad enc_{-1,0}(x) \ \lor \ enc_{-1,1}(x) \ \lor \ enc_{-1,2}(x) \ \lor \ enc_{-1,3}(x).$$

Clearly, the formula $enc_{-1}(x)$ is satisfied if x is the root node of a binary tree that is isomorphic to the binary tree $\mathcal{B}_{-1}(i)$, for one of the numbers $i \in \{0, 1, 2, 3\}$.

For $h \geqslant 0$, we let

$$enc_h(x) \quad := \quad \gamma_{\mathrm{Tower}(h+1)-1}(x) \ \land \ \forall y\big(\delta_{=\mathrm{Tower}(h+1)}(x, y) \ \rightarrow \ enc_{h-1}(y)\big).$$

For each $h \geqslant 0$, the size of the formulas $\gamma_{\mathrm{Tower}(h+1)-1}$ and $\delta_{=\mathrm{Tower}(h+1)}$ is linear in $\mathrm{Tower}(h)$. An easy induction shows that, for each $h \geqslant 0$,

$$\sum_{i=0}^{h} \mathrm{Tower}(i) \quad < \quad 2 \cdot \mathrm{Tower}(h).$$

Hence, for $h \geqslant 0$, the formula enc_h has size in $O(\mathrm{Tower}(h))$.

We let $min_{-1}(x) := enc_{-1,0}(x)$. For $h \geqslant 0$, we let $min_h(x) := \neg \exists y \ \delta_{=\mathrm{Tower}(h+1)}(x, y)$. Clearly, min_h is of size $O(\mathrm{Tower}(h))$.

The construction of the formulas eq_h, $less_h$, $succ_h$ and max_h are easy adaptations of Lemma 3.2 in [DGKS07] (Lemma 10.21 in [FG06]) and Lemma 3.4 in [DGKS07]. Their construction is best understood by keeping in mind the binary expansions of the numbers, encoded by binary trees.

Consider the formula

$$eq_{-1}(x, y) \quad := \quad \bigvee_{i=0}^{3} \Big(enc_{-1,i}(x) \ \land \ enc_{-1,i}(y) \Big).$$

Note that the formula $eq_{-1}(x, y)$ is satisfied if there is an $i \in \{0, 1, 2, 3\}$ such that x and y are root nodes of two binary trees isomorphic to the binary tree $\mathcal{B}_{-1}(i)$ depicted in Figure 2.

For each $h \geqslant 0$, we let

$$
\begin{aligned}
eq_h(x, y) \ := \ & \big(\exists x' \ \delta_{=\mathrm{Tower}(h+1)}(x, x') \ \leftrightarrow \ \exists y' \ \delta_{=\mathrm{Tower}(h+1)}(y, y')\big) \ \land \\
& \forall x' \ \Big(\delta_{=\mathrm{Tower}(h+1)}(x, x') \ \rightarrow \\
& \qquad \exists y' \ \Big(\delta_{=\mathrm{Tower}(h+1)}(y, y') \ \land \\
& \qquad\qquad \forall y'' \ \big(\delta_{=\mathrm{Tower}(h+1)}(y, y'') \ \rightarrow \\
& \qquad\qquad\qquad \exists x'' \ \big(\delta_{=\mathrm{Tower}(h+1)}(x, x'') \ \land \\
& \qquad\qquad\qquad\qquad \forall u \forall v \big(\big((u{=}x' \ \land \ v{=}y') \ \lor \\
& \qquad\qquad\qquad\qquad\qquad (u{=}x'' \ \land \ v{=}y'')\big) \\
& \qquad\qquad\qquad\qquad\qquad \rightarrow \ eq_{h-1}(u, v)\big)\big)\big)\big)\big)
\end{aligned}
$$

As for enc_h, it can easily be seen that for each $h \geqslant 0$, the size of eq_h is in $O(\mathrm{Tower}(h))$.

The following formula $less_{-1}(x, y)$ is satisfied if there are numbers $i, j \in \{0, \ldots, 3\}$ with $i < j$ such that x and y are root nodes of binary trees that are isomorphic to the binary trees $\mathcal{B}_{-1}(i)$ and $\mathcal{B}_{-1}(j)$, respectively:

$$less_{-1}(x, y) := \bigvee_{0 \leqslant i < j \leqslant 3} \Big(enc_{-1,i}(x) \ \wedge \ enc_{-1,j}(y) \Big).$$

For every $h \geqslant 0$, let

$$\begin{aligned} less_h(x, y) := \ &\exists y' \Big(\delta_{=\mathrm{Tower}(h+1)}(y, y') \wedge \\ &\forall x' \ \big(\delta_{=\mathrm{Tower}(h+1)}(x, x') \ \rightarrow \ \neg eq_{h-1}(x', y') \big) \wedge \\ &\forall x'' \ \big((\delta_{=\mathrm{Tower}(h+1)}(x, x'') \ \wedge \ less_{h-1}(y', x'')) \\ &\qquad \rightarrow \ \exists y'' \ (\delta_{=\mathrm{Tower}(h+1)}(y, y'') \ \wedge \ eq_{h-1}(y'', x'')) \big) \Big) \end{aligned}$$

There is a number $c > 0$ such that, for all $h \geqslant 0$,

$$\|less_h\| \quad \leqslant \quad c \ + \ \sum_{i=0}^{h} (c \ + \ 4 \cdot \|\delta_{=\mathrm{Tower}(i+1)}\| \ + \ 2 \cdot \|eq_{i-1}\|).$$

Since, for $i \geqslant 0$, the formulas $\delta_{=\mathrm{Tower}(i+1)}$ and eq_{i-1} have size in $O(\mathrm{Tower}(h))$, the formula $less_h$ has size in $O(\mathrm{Tower}(h))$.

Similarly to the formula $less_{-1}$, a formula $succ_{-1}$ can be chosen as

$$succ_{-1}(x, y) := \bigvee_{0 \leqslant i \leqslant 2} \Big(enc_{-1,i}(x) \ \wedge \ enc_{-1,i+1}(y) \Big).$$

For each $h \geqslant 0$, we let

$$\begin{aligned} succ_h(x, y) := \ &\exists y' \Big(\delta_{=\mathrm{Tower}(h+1)}(y, y') \\ &\wedge \ \forall y'' \big((\delta_{=\mathrm{Tower}(h+1)}(y, y'') \ \wedge \ \neg eq_{h-1}(y'', y')) \ \rightarrow \ less_{h-1}(y', y'') \big) \\ &\wedge \ \forall x' \big(\delta_{=\mathrm{Tower}(h+1)}(x, x') \ \rightarrow \ \neg eq_{h-1}(x', y') \big) \\ &\wedge \ \forall y'' \big((\delta_{=\mathrm{Tower}(h+1)}(y, y'') \ \wedge \ less_{h-1}(y', y'')) \\ &\qquad \rightarrow \ \exists x''(\delta_{=\mathrm{Tower}(h+1)}(x, x'') \ \wedge \ eq_{h-1}(x'', y'')) \big) \\ &\wedge \ \forall x'' \big((\delta_{=\mathrm{Tower}(h+1)}(x, x'') \ \wedge \ less_{h-1}(y', x'')) \\ &\qquad \rightarrow \ \exists y''(\delta_{=\mathrm{Tower}(h+1)}(y, y'') \ \wedge \ eq_{h-1}(y'', x'')) \big) \\ &\wedge \ \Big(\neg min_{h-1}(y') \\ &\qquad \rightarrow \ \big(\exists x' \ (\delta_{=\mathrm{Tower}(h+1)}(x, x') \ \wedge \ min_{h-1}(x')) \\ &\qquad\qquad \wedge \ \forall x' \ ((\delta_{=\mathrm{Tower}(h+1)}(x, x') \ \wedge \ less_{h-1}(x', y')) \\ &\qquad\qquad\qquad \rightarrow \ \exists z \ (succ_{h-1}(x', z) \\ &\qquad\qquad\qquad\qquad \wedge \ (z{=}y' \ \vee \ \delta_{=\mathrm{Tower}(h+1)}(x, z))))) \big) \Big) \Big) \end{aligned}$$

Recall that the formulas $\delta_{=\mathrm{Tower}(h+1)}$, eq_h, and $less_h$ have size in $O(\mathrm{Tower}(h))$. Hence, a simple induction shows that for $h \geqslant 0$, also $succ_h$ has size in $O(\mathrm{Tower}(h))$.

Finally, let $max_{-1}(x) := enc_{-1,3}(x)$; and for each $h \geqslant 0$, let

$$max_h(x) := \exists y \left(\delta_{=\text{Tower}(h+1)}(x,y) \wedge min_{h-1}(y) \right) \wedge$$

$$\forall y \left(\delta_{=\text{Tower}(h+1)}(x,y) \right.$$

$$\left. \rightarrow \left(max_{h-1}(y) \vee \exists z \left(\delta_{=\text{Tower}(h+1)}(x,z) \wedge succ_{h-1}(y,z) \right) \right) \right).$$

Again, a simple induction shows that, for each $h \geqslant 0$, the size of max_h is in $O(\text{Tower}(h))$. This concludes the proof of Lemma 5.2. $\qquad\square$

5.2. **Lower bounds for preservation theorems.** The upper bounds of Section 3 are complemented by the following two lower bounds for certain classes of finite acyclic structures of degree 3 and first-order sentences that are preserved under extensions (homomorphisms): We show that even under these restrictions, a 3-fold exponential blow-up in terms of the size of the input sentence is unavoidable when constructing the equivalent existential (existential-positive) first-order sentence.

Theorem 5.3. *Let $\sigma := \{S_0, S_1, V_0, V_1\}$, where S_0, S_1 are binary and V_0, V_1 are unary relation symbols, and let $\mathfrak{C}$ be the class of all finite ordered binary forests $\mathcal{F}$ of signature σ, where $V_0^{\mathcal{F}}$ and $V_1^{\mathcal{F}}$ may be arbitrary subsets of the universe. There is a real number $\epsilon > 0$ and a sequence $(\varphi_h)_{h>1}$ of $\text{FO}(\sigma)$-sentences of increasing size such that for each $h > 1$ the following holds:*

(1) *φ_h is preserved under extensions on $\mathfrak{C}$, and*

(2) *every existential $\text{FO}(\sigma)$-sentence that is equivalent to φ_h on $\mathfrak{C}$ has size at least*

$$2^{2^{2^{\epsilon \cdot \|\varphi_h\|}}}.$$

Theorem 5.4. *Let $\sigma' := \{S_0, S_1\} \cup \{V_M : M \subseteq \{0,1\}\}$, where S_0, S_1 are binary relation symbols and, for each $M \subseteq \{0,1\}$, V_M is a unary relation symbol. Let $\mathfrak{C}'$ be the class of all finite ordered binary forests $\mathcal{F}$ over σ', where $(V_M^{\mathcal{F}})_{M \subseteq \{0,1\}}$ is a partition of the universe. There is a real number $\epsilon > 0$ and a sequence $(\varphi_h')_{h>1}$ of $\text{FO}(\sigma')$-sentences of increasing size such that for each $h > 1$ the following holds:*

(1) *φ_h' is preserved under homomorphisms on $\mathfrak{C}'$, and*

(2) *every existential-positive $\text{FO}(\sigma')$-sentence that is equivalent to φ_h' on $\mathfrak{C}'$ has size at least*

$$2^{2^{2^{\epsilon \cdot \|\varphi_h'\|}}}.$$

In Theorem 5.3 and Theorem 5.4, an *ordered binary forest* of signature σ or σ' is a structure whose Gaifman graph is a forest and where the binary relation symbols S_0 and S_1 are interpreted as the left and right successor relation and every node is allowed to have at most one left successor and at most one right successor. An *ordered binary tree* is an ordered binary forest with only one connected component.

The proofs of Theorem 5.3 and Theorem 5.4, which can be found below, use the encoding of numbers by binary trees. The main challenge here is to find sequences of sentences that not only have large minimal models but are also preserved under extensions and homomorphisms, respectively. Towards this end, the auxiliary unary relation symbols in σ and σ' are introduced to interpret binary tree encodings in ordered binary forests. Both proofs rely on the following observation:

Lemma 5.5. *Let σ be a relational signature and let $\mathfrak{C}$ be a class of σ-structures that is closed under induced substructures. For each $\mathrm{FO}(\sigma)$-sentence φ and each $N \geqslant 1$ the following holds: If φ has a $\mathfrak{C}$-minimal model of size at least N, then every existential $\mathrm{FO}(\sigma)$-sentence that is $\mathfrak{C}$-equivalent to φ has size greater than N.*

Proof. Let φ be an $\mathrm{FO}(\sigma)$-sentence and let $\mathcal{A}$ be a $\mathfrak{C}$-minimal model of φ with at least N elements. For contradiction, assume that ψ is an existential $\mathrm{FO}(\sigma)$-sentence of size at most N that is $\mathfrak{C}$-equivalent to φ. In particular, ψ has less than N quantifiers.

Since φ and ψ are $\mathfrak{C}$-equivalent, $\mathcal{A}$ is also a model of ψ. But since ψ is an existential sentence with less than N quantifiers, there is an induced substructure $\mathcal{B}$ of $\mathcal{A}$ of size less than N, such that $\mathcal{B}$ is a model of ψ. Since $\mathfrak{C}$ is closed under induced substructures, $\mathcal{B}$ belongs to $\mathfrak{C}$. And since φ and ψ are $\mathfrak{C}$-equivalent, $\mathcal{B}$ also is a model of φ. However, this contradicts the assumption that $\mathcal{A}$ is a $\mathfrak{C}$-*minimal* model of φ. $\square$

Recall the $\mathrm{FO}(E)$-formulas provided by Lemma 5.2, which define arithmetic relations between binary tree encodings of numbers. It is straightforward to use these formulas to construct, for each $h \geqslant 1$, an $\mathrm{FO}(E)$-sentence of size $O(\mathrm{Tower}(h))$ that has a minimal model of size at least $\mathrm{Tower}(h{+}3)$. However, it is far less obvious to construct sentences that, at the same time, are preserved under extensions (on finite binary forests). To achieve this, we interpret binary tree encodings in complete *ordered* binary forests of suitable height. This way, we make sure that no extension of such binary forests can modify the binary tree encodings.

Consider the signature $\sigma := \{S_0, S_1, V_0, V_1\}$ where S_0 and S_1 are binary relation symbols and V_0 and V_1 are unary relation symbols. Recall that in each ordered binary forest $\mathcal{F}$ over the signature σ, the unary relations $V_0^{\mathcal{F}}$ and $V_1^{\mathcal{F}}$ may be interpreted by arbitrary sets of nodes of $\mathcal{F}$. In the following, we denote the class of all finite ordered binary forests over σ by $\mathfrak{C}$. Note that every structure in $\mathfrak{C}$ has degree at most 3.

Complete ordered binary trees are defined in the obvious way. Also, we adapt the notions $\mathcal{T}[\leqslant d]$ and $\mathcal{F}_a$ from Section 5.1 from unordered binary trees $\mathcal{T}$ and forests $\mathcal{F}$ to *ordered* binary trees $\mathcal{B}$ and forests $\mathcal{F}$ (with $d \geqslant 0$ and $a \in F$) in the obvious way. Thus, $\mathcal{B}[\leqslant d]$ is the subtree of $\mathcal{B}$ induced by all nodes of $\mathcal{B}$ that are reachable from the root of $\mathcal{B}$ by a directed path of length at most d. And $\mathcal{F}_a$ is the subtree of $\mathcal{F}$ induced by all nodes reachable by a directed path from a.

Recall the definition of transductions from Section 4.3.1 and consider the following transduction $\Theta := (\theta, \theta_E)$ from $\{E\}$ to σ, defined by $\theta(x) := x{=}x$ and

$$\theta_E(x, y) := \bigvee_{i \in \{0,1\}} \big(S_i(x, y) \,\wedge\, V_i(x) \big).$$

The transduction Θ makes use of the unary relations V_0 and V_1 to interpret binary forests in ordered binary forests. More specifically, for an ordered binary forest $\mathcal{F}$, the structure $\Theta(\mathcal{F})$ is the (unordered) binary forest consisting of all nodes from $\mathcal{F}$ and, for all nodes $a, b \in F$, there is an edge from a to b iff b is the left successor of a in $\mathcal{F}$ and $a \in V_0^{\mathcal{F}}$, or b is the right successor of a in $\mathcal{F}$ and $a \in V_1^{\mathcal{F}}$.

Proof of Theorem 5.3. Recall that each extension of a structure $\mathcal{A}$ contains $\mathcal{A}$ as an *induced* substructure. Consider a complete ordered binary tree $\mathcal{A} \in \mathfrak{C}$ of height $d \geqslant 1$ and let a be its root node. Assume that $\Theta(\mathcal{A})_a$ is a binary tree of height at most $d{-}1$. Note that every leaf b of $\Theta(\mathcal{A})_a$ has a left and a right successor in $\mathcal{A}$ but is neither contained in $V_0^{\mathcal{A}}$ nor in

$V_1^{\mathcal{A}}$. Therefore, by construction of the transduction Θ, for each extension $\mathcal{B}$ of $\mathcal{A}$ in $\mathfrak{C}$ we have that $\Theta(\mathcal{A})_a$ and $\Theta(\mathcal{B})_a$ are isomorphic.

We will use this observation to protect binary tree encodings, interpreted in complete ordered binary trees, against modifications by extensions of the underlying structure. Recall that we know by Lemma 5.1 that for each $h > 1$ and every $i \in [0, \mathrm{Tower}(h{+}3){-}1]$, each binary tree from $\mathfrak{B}_h(i)$ has height less than $2 \cdot \mathrm{Tower}(h{+}1)$.

Similarly to the $\mathrm{FO}(E)$-formula $\gamma_d(x)$, defined in the proof of Lemma 5.2, for each $d \geqslant 0$, there is an $\mathrm{FO}(\sigma)$-formula $\gamma_d^<(x)$ of size $O(\log d)$ that is satisfied by a node a of an ordered binary forest $\mathcal{A}$ iff $\mathcal{A}_a[\leqslant d]$ is a complete ordered binary tree of height d.

In the following, the transduction Θ is applied to the $\mathrm{FO}(E)$-formulas provided by Lemma 5.2. Let $(\varphi_h)_{h>1}$ be the sequence of $\mathrm{FO}(\sigma)$-sentences that is defined, for each $h > 1$, by

$$\begin{aligned}
\varphi_h \;:=\;\; & \exists x \,\big(enc_h^<(x) \;\wedge\; \Theta(min_h)(x)\big) \;\wedge\; \\
& \forall x \,\big(enc_h^<(x) \;\rightarrow \\
& \qquad \big(\Theta(max_h)(x) \;\vee\; \exists y \,(enc_h^<(y) \;\wedge\; \Theta(succ_h)(x,y))\big)\big),
\end{aligned}$$

where

$$enc_h^<(x) \;:=\;\; \Theta(enc_h)(x) \;\wedge\; \gamma_{2\cdot\mathrm{Tower}(h+1)}^<(x).$$

The formula φ_h expresses that there is a node x that encodes the number 0 as a binary tree encoding with parameter h, and for each number $i < \mathrm{Tower}(h{+}3) - 1$ that is encoded by a node x, there also exists a node y that encodes the number $i + 1$.

For the fixed transduction Θ it follows from Lemma 4.6 that, for each $\mathrm{FO}(E)$-formula φ, the size of $\Theta(\varphi)$ is linear in the size of φ. Therefore, by Lemma 5.2 and since $\gamma_{2\cdot\mathrm{Tower}(h+1)}^<$ has size in $O(\mathrm{Tower}(h))$, the formula φ_h also has size in $O(\mathrm{Tower}(h))$.

For the remainder of the proof, we fix an arbitrary $h > 1$. The following claim follows from the choice of the sentence φ_h and from Lemma 5.5.

Claim 5.6. Every existential $\mathrm{FO}(\sigma)$-sentence that is equivalent to φ_h on $\mathfrak{C}$ has size at least $\mathrm{Tower}(h{+}3)$.

Proof of Claim 5.6. It is easy to see that there are structures in $\mathfrak{C}$ that satisfy φ_h. Furthermore, by definition of the subformulas of φ_h (see Lemma 5.2), each model $\mathcal{A} \in \mathfrak{C}$ of φ_h has to contain at least $\mathrm{Tower}(h{+}3)$ pairwise distinct nodes $a_0, \ldots, a_{\mathrm{Tower}(h+3)-1}$ such that, for each $i \in [0, \mathrm{Tower}(h{+}3){-}1]$, the binary tree $\Theta(\mathcal{A})_{a_i}$ is a binary tree encoding with parameter h of the number i, i.e., the binary tree $\Theta(\mathcal{A})_{a_i}$ belongs to the set $\mathfrak{B}_h(i)$.

Together with Lemma 5.5, this observation completes the proof of Claim 5.6. $\qquad\square$

Claim 5.7. φ_h is preserved under extensions on $\mathfrak{C}$.

Proof of Claim 5.7. Let $\mathcal{A} \in \mathfrak{C}$ be a model of φ_h. By definition of φ_h, there are pairwise distinct nodes $a_0, \ldots, a_{\mathrm{Tower}(h+3)-1}$ in $\mathcal{A}$ such that, for each $i \in [0, \mathrm{Tower}(h{+}3){-}1]$, the binary tree $\Theta(\mathcal{A})_{a_i}$ belongs to the set $\mathfrak{B}_h(i)$.

Let $\mathcal{B} \in \mathfrak{C}$ be an extension of $\mathcal{A}$. By construction of φ_h, for each $i \in [0, \mathrm{Tower}(h{+}3){-}1]$, the substructure $\mathcal{A}_{a_i}[\leqslant 2\cdot\mathrm{Tower}(h{+}1)]$ is a complete ordered binary tree. Therefore, $\Theta(\mathcal{A})_{a_i}$ and $\Theta(\mathcal{B})_{a_i}$ are isomorphic and thus, also $\Theta(\mathcal{B})_{a_i}$ belongs to the set $\mathfrak{B}_h(i)$. On the other hand, let b be a node from $\mathcal{B}$ such that $\mathcal{B} \models enc_h^<[b]$. Then, there is an $i \in [0, \mathrm{Tower}(h{+}3){-}1]$ such that $\Theta(\mathcal{B})_b$ belongs to $\mathfrak{B}_h(i)$ and hence, either $\mathcal{B} \models \Theta(max_h)[b]$ or $\mathcal{B} \models \Theta(succ_h)[b, a_{i+1}]$. Altogether, it follows that $\mathcal{B} \models \varphi_h$. This completes the proof of Claim 5.7. $\qquad\square$

Since $(\varphi_h)_{h>1}$ is a sequence of $\mathrm{FO}(\sigma)$-sentences of increasing size $O(\mathrm{Tower}(h))$, there is a real number $\epsilon > 0$ such that $\epsilon \cdot \|\varphi_h\| \leqslant \mathrm{Tower}(h)$, for every $h > 1$. By Claim 5.6, every existential $\mathrm{FO}(\sigma)$-sentence that is equivalent to φ_h on $\mathfrak{C}$ has size at least

$$\mathrm{Tower}(h{+}3) \quad = \quad 2^{2^{2^{\mathrm{Tower}(h)}}} \quad \geqslant \quad 2^{2^{2^{\epsilon \cdot \|\varphi_h\|}}}.$$

This concludes the proof of Theorem 5.3. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

For the proof of Theorem 5.4 we need to construct FO-sentences with large minimal models that are preserved under homomorphism. Let $\sigma' := \{S_0, S_1\} \cup \{V_M : M \subseteq \{0,1\}\}$ be a signature, where S_0 and S_1 are binary relation symbols and, for each $M \subseteq \{0,1\}$, V_M is a unary relation symbol. Consider the class $\mathfrak{C}'$ of all finite *ordered and colored binary forests* over σ', i.e., finite σ'-structures $\mathcal{A}$ where the binary relations $S_0^{\mathcal{A}}$ and $S_1^{\mathcal{A}}$ correspond to the left and right successor relation of a forest, every node is allowed to have at most one left and at most one right successor, and the unary relations $V_M^{\mathcal{A}}$, for $M \subseteq \{0,1\}$, are a partition of the universe of $\mathcal{A}$. Note that, for each $\mathcal{A} \in \mathfrak{C}'$ and each $\mathcal{B} \in \mathfrak{C}'$, if there is a homomorphism h from $\mathcal{A}$ to $\mathcal{B}$ then for every $a \in A$ and each $M \subseteq \{0,1\}$, it holds that $a \in V_M^{\mathcal{A}}$ iff $h(a) \in V_M^{\mathcal{B}}$.

We only sketch a proof of Theorem 5.4, which is very similar to the one of Theorem 5.3.

Let Θ' be the transduction (θ', θ'_E) from $\{E\}$ to σ' that is defined by $\theta'(x) := x{=}x$ and

$$\theta'_E(x,y) \;:=\; \bigvee_{i \in \{0,1\}} \left(S_i(x,y) \;\wedge\; \bigvee_{M \subseteq \{0,1\}, i \in M} V_M(x) \right).$$

Proof of Theorem 5.4. For colored and ordered binary trees and forests from $\mathfrak{C}'$, we make use of the same notation already introduced for ordered binary forests. Consider a complete colored and ordered binary tree $\mathcal{A} \in \mathfrak{C}'$ of height $d \geqslant 1$ and let a be its root node. Assume that $\Theta(\mathcal{A})_a$ is a binary tree of height at most $d{-}1$. Observe that every homomorphism h from $\mathcal{A}$ to a colored and ordered binary forest $\mathcal{B} \in \mathfrak{C}'$ is injective. Furthermore, for each $\mathcal{B} \in \mathfrak{C}'$ for which there is a homomorphism h from $\mathcal{A}$ to $\mathcal{B}$, the binary trees $\Theta(\mathcal{A})_a$ and $\Theta(\mathcal{B})_{h(a)}$ are isomorphic.

Define a sequence $(\varphi'_h)_{h>1}$ of $\mathrm{FO}(\sigma')$-formulas similar to the sequence $(\varphi_h)_{h>1}$ in the proof of Theorem 5.3, with the only modification being the application of the transduction Θ' instead of Θ. Of course, for each $h > 1$, also φ'_h has size $O(\mathrm{Tower}(h))$.

Let $h > 1$. Since every existential-positive sentence is an existential sentence and the class $\mathfrak{C}'$ is closed under induced substructures we can follow the lines of the proof of Theorem 5.3 to show that every existential-positive $\mathrm{FO}(\sigma)$-sentence that is equivalent to φ'_h on $\mathfrak{C}'$ has size at least $\mathrm{Tower}(h{+}3)$.

Similarly, using the observation above, some small adaptations to the proof of Theorem 5.3 suffice to show that φ'_h is preserved under homomorphisms on $\mathfrak{C}'$. This concludes the proof of Theorem 5.4. $\qquad\qquad\qquad\qquad\qquad\qquad\square$

5.3. **Lower bounds for Feferman-Vaught decompositions.** The following two lower bounds show that for structures of degree 3 or 2 the algorithm of Theorem 4.2 for the construction of reduction sequences for Feferman-Vaught decompositions is basically optimal. Recall that a *binary forest* is a disjoint union of directed trees of signature $\{E\}$, where every node has at most 2 children.

Theorem 5.8. *There is a real number $\epsilon > 0$ and a sequence of $\mathrm{FO}(E)$-sentences $(\varphi_h)_{h \geqslant 1}$ of increasing size such that, for every $h \geqslant 1$, every 2-disjoint decomposition for φ_h on finite binary forests has size at least*

$$2^{2^{2^{\epsilon \cdot \|\varphi_h\|}}}.$$

Let $\sigma' := \{S, L_0, L_1\}$ be the signature consisting of a binary relation symbol S and two unary relation symbols L_0 and L_1. A *labeled chain* is a finite σ'-structure $\mathcal{C}$ whose $\{S\}$-reduct is a chain of finite length, i.e., a finite directed path, and where the sets $L_0^{\mathcal{C}}$ and $L_1^{\mathcal{C}}$ are disjoint subsets of the universe of $\mathcal{C}$. The class of all σ'-structures that are disjoint unions of finitely many labeled chains is denoted by $\mathfrak{LC}$. Note that all structures in $\mathfrak{LC}$ have degree at most two.

Theorem 5.9. *There is a real number $\epsilon > 0$ and a sequence of $\mathrm{FO}(\sigma')$-sentences $(\varphi'_h)_{h \geqslant 1}$ of increasing size such that, for every $h \geqslant 1$, every 2-disjoint decomposition for φ'_h on $\mathfrak{LC}$ has size at least*

$$2^{2^{\epsilon \cdot \|\varphi'_h\|}}.$$

Both theorems are corollaries to a generalisation of Proposition 6.7 in [GJL15]. For proving our lower bounds, we use the following Lemma 5.10, along with suitable encodings of numbers by binary trees (for Theorem 5.8) and labeled paths (for Theorem 5.9). The lemma is proved by a simple counting argument and distills the combinatorial essence of the proof of Proposition 6.7 in [GJL15].

Lemma 5.10. *Let σ be a relational signature and let $\sigma_2 := \sigma \cup \{P_1, P_2\}$, where P_1 and P_2 are unary relation symbols that are not contained in σ. Let $\mathfrak{C}$ be a class of σ-structures, and let φ be an $\mathrm{FO}(\sigma_2)$-sentence. Let $H \geqslant 1$. If there are σ-structures $\mathcal{A}_0, \ldots, \mathcal{A}_{2^H-1} \in \mathfrak{C}$ such that for all $i, j \in [0, 2^H-1]$ it holds that*

$$\mathcal{A}_i \oplus \mathcal{A}_j \quad \models \quad \varphi \qquad \textit{iff} \qquad i = j,$$

then every 2-disjoint decomposition for φ on $\mathfrak{C}$ has size at least H.

Proof. Assume that $\mathcal{A}_0, \ldots, \mathcal{A}_{2^H-1}$ are structures from $\mathfrak{C}$ such that for all $i, j \in [0, 2^H-1]$,

$$\mathcal{A}_i \oplus \mathcal{A}_j \quad \models \quad \varphi \qquad \text{iff} \qquad i = j. \tag{5.1}$$

For contradiction, assume that there is a 2-disjoint decomposition $D := (\Delta_1, \Delta_2, \beta)$ for φ on $\mathfrak{C}$ of size less than H. Thus, Δ_1 and Δ_2 are finite sets of $\mathrm{FO}(\sigma)$-sentences and β is a propositional formula with variables from the set $X_D := \{X_{k,\delta} \colon k \in [1, 2], \delta \in \Delta_k\}$, and for all $i, j \in [0, 2^H-1]$ we have

$$\mathcal{A}_i \oplus \mathcal{A}_j \quad \models \quad \varphi \qquad \text{iff} \qquad \mu_{i,j} \quad \models \quad \beta, \tag{5.2}$$

where $\mu_{i,j} \colon X_D \to \{0, 1\}$ assigns variables of β such that for all $\delta_1 \in \Delta_1$,

$$\mu_{i,j}(X_{1,\delta_1}) = 1 \qquad \text{iff} \qquad \mathcal{A}_i \quad \models \quad \delta_1,$$

and for all $\delta_2 \in \Delta_2$,

$$\mu_{i,j}(X_{2,\delta_2}) = 1 \qquad \text{iff} \qquad \mathcal{A}_j \quad \models \quad \delta_2.$$

From (5.2) and (5.1) as just mentioned above, we know that for all $i, j \in [0, 2^H-1]$,

$$\mu_{i,j} \quad \models \quad \beta \qquad \text{iff} \qquad i = j. \tag{5.3}$$

Note that the number of variables of β is less than H, and hence the number of distinct variable assignments is less than 2^H. Thus, there exist $i, j \in [0, 2^H-1]$ with $i \neq j$ such that $\mu_{i,i} = \mu_{j,j}$. We let $\mu := \mu_{i,i}$. Clearly, due to (5.3), we have

$$\mu \quad \models \quad \beta. \tag{5.4}$$

Along the definition of $\mu_{i,j}$, and using the fact that $\mu = \mu_{i,i} = \mu_{j,j}$, it is straightforward to see that $\mu_{i,j} = \mu$. Thus, from (5.4) and (5.2) we obtain that $\mathcal{A}_i \oplus \mathcal{A}_j \models \varphi$. This, however, contradicts (5.1). $\qquad\square$

Using the above lemma, along with the binary tree encodings introduced in Subsection 5.1 we obtain a proof of Theorem 5.8. Let us remark that the FO-formulas we use within the proof do not even make use of the unary relation symbols P_1 and P_2.

Proof of Theorem 5.8. We use the FO(E)-formula $eq_h(x, y)$ from Lemma 5.2. For every $h \geqslant 1$, let

$$\varphi_h \quad := \quad \forall x \left(root(x) \; \rightarrow \; \exists y \left(root(y) \; \wedge \; eq_h(x, y) \; \wedge \; \neg x{=}y \right) \right).$$

The formula $root(x)$ is satisfied by exactly the root nodes of the trees in a forest; it can easily defined by $root(x) := \neg \exists y\, E(y, x)$. Because eq_h has size in $O(\mathrm{Tower}(h))$, also φ_h has size in $O(\mathrm{Tower}(h))$.

Let $h \geqslant 1$ and let $H := \mathrm{Tower}(h{+}3)$. For every $i \in [0, H{-}1]$, let $\mathcal{B}_{h,i}$ be a binary tree encoding of i with parameter h, i.e., let $\mathcal{B}_{h,i} \in \mathfrak{B}_h(i)$ (see Section 5.1 above for the definition of $\mathfrak{B}_h(i)$). Furthermore, for each $i \in [0, 2^H{-}1]$, let $\mathcal{A}_{h,i}$ be the disjoint union of all binary tree encodings $\mathcal{B}_{h,i'}$ for $i' \in [0, H{-}1]$ such that $\mathrm{Bit}(i', i) = 1$.

It is easy to verify that for numbers $i, j \in [0, 2^H{-}1]$

$$\mathcal{A}_{h,i} \oplus \mathcal{A}_{h,j} \quad \models \quad \varphi_h \qquad \text{iff} \qquad i = j.$$

By Lemma 5.10, every 2-disjoint decomposition for φ_h on the class of binary forests has size at least $H = \mathrm{Tower}(h{+}3)$. Analogously to the final step of the proof of Theorem 5.3, and since φ_h has size in $O(\mathrm{Tower}(h))$, we can conclude that

$$H \quad \geqslant \quad 2^{2^{2^{\epsilon \cdot \|\varphi_h\|}}},$$

for a suitable real number $\epsilon > 0$. This completes the proof of Theorem 5.8. $\qquad\square$

For the proof of Theorem 5.9 we employ the obvious encodings of numbers by strings (cf., e.g., [FG04, HKS13]): Let $\Sigma = \{0, 1\}$. For $h \geqslant 1$ and $i \in [0, 2^{2^h}{-}1]$ let $\mathrm{bin}_{2^h}(i)$ denote the binary expansion of i of length 2^h. Strings $w \in \Sigma^+$ are represented by structures $\mathcal{B}_w \in \mathfrak{UC}$ in the usual way: the universe of $\mathcal{B}_w$ is the set of positions of the string w, the relation $S^{\mathcal{B}_w}$ is the successor relation on the positions of w, and $L_a^{\mathcal{B}_w}$ consists, for each $a \in \Sigma$, of all positions of w that carry the letter a.

We will use structures $\mathcal{B}_w$ for strings $w = \mathrm{bin}_{2^h}(i)$ with $i \in [0, 2^{2^h}{-}1]$, and we will rely on the following result of [FG04].

Lemma 5.11 (Lemma 20 in [FG04])**.** *There is a sequence of* FO(σ')*-formulas* $(eq'_h(x, y))_{h \geqslant 1}$ *of size* $\|eq'_h(x, y)\| \in O(h)$ *such that for all structures* $\mathcal{B} \in \mathfrak{UC}$ *and all nodes* a, b *of* $\mathcal{B}$ *the following holds: If* a *and* b *are the starting positions of labeled chains isomorphic to* $\mathcal{B}_{\mathrm{bin}_{2^h}}(i)$ *and* $\mathcal{B}_{\mathrm{bin}_{2^h}}(j)$*, respectively, for* $i, j \in [0, 2^{2^h}{-}1]$*, then*

$$(\mathcal{B}, a, b) \quad \models \quad eq'_h(x, y) \qquad \textit{iff} \qquad i = j. \qquad\square$$

We are now ready to prove Theorem 5.9.

Proof of Theorem 5.9. For every $h \geqslant 1$, let

$$\varphi'_h \quad := \quad \forall x \left(\mathit{root}'(x) \;\rightarrow\; \exists y \left(\mathit{root}'(y) \;\wedge\; \mathit{eq}'_h(x,y) \;\wedge\; \neg x{=}y \right) \right),$$

where $\mathit{root}'(x) := \neg \exists y \; S(y,x)$. Because eq'_h has size in $O(h)$, also φ'_h has size in $O(h)$.

Let $h \geqslant 1$ and let $H := 2^{2^h}$. For each $i \in [0, 2^H{-}1]$, let $\mathcal{A}_{h,i}$ be the disjoint union of all labeled chains $\mathcal{B}_{\mathrm{bin}_{2^h}}(i')$ for $i' \in [0, H{-}1]$ such that $\mathrm{Bit}(i', i) = 1$. It is easy to verify that for all numbers $i, j \in [0, 2^H{-}1]$ it holds that

$$\mathcal{A}_{h,i} \oplus \mathcal{A}_{h,j} \quad \models \quad \varphi'_h \qquad \text{iff} \qquad i = j.$$

Therefore, by Lemma 5.10, every 2-disjoint decomposition for φ'_h on $\mathfrak{UC}$ has size at least

$$H \quad = \quad 2^{2^h} \quad \geqslant \quad 2^{2^{\epsilon \cdot \|\varphi_h\|}},$$

for a suitable real number $\epsilon > 0$. This completes the proof of Theorem 5.9. $\quad\square$

6. Concluding remarks

In this section, we give a short summary of our main results and some directions for further research. For this, we fix a relational signature σ and a class $\mathfrak{C}_d$ of σ-structures of degree $\leqslant d$, for a $d \geqslant 3$.

Our first two main results, to which we will refer in the following with (PE) and (PH), are algorithmic versions of two preservations theorems that are restricted to the class $\mathfrak{C}_d$. Both require the class $\mathfrak{C}_d$ to be closed under induced substructures and disjoint unions.

(PE) For each sentence φ of $\mathrm{FO}{+}\mathrm{MOD}_m(\sigma)$ that is preserved under extensions on $\mathfrak{C}_d$, a $\mathfrak{C}_d$-equivalent existential $\mathrm{FO}(\sigma)$-sentence can be constructed in 5-fold exponential time.

(PH) For each sentence φ of $\mathrm{FO}{+}\mathrm{MOD}_m(\sigma)$ that is preserved under homomorphisms on $\mathfrak{C}_d$, a $\mathfrak{C}_d$-equivalent existential-positive FO-sentence can be constructed in 4-fold exponential time, provided that $\mathfrak{C}_d$ is decidable in 1-fold exponential time.

For (PE) and (PH) we have shown that a 3-fold exponential blow-up of the computed existential or existential-positive sentence is unavoidable.

Our third main result is an algorithmic version of the Feferman-Vaught theorem for disjoint sums (and, using transductions, other products of structures) of σ-structures of bounded degree.

(FV) For each formula $\varphi(\overline{x})$ of $\mathrm{FO}(\sigma_s)$, where $s \geqslant 1$, a disjoint decomposition with respect to disjoint sums and direct products of $\mathfrak{C}_d$-structures can be computed in 3-fold exponential time.

Furthermore, a matching lower bound shows that our algorithm for (FV) is basically optimal.

For most of our results it is known that there is no hope to considerably extend the class of structures to which they apply. Notably, concerning (PE) and (FV), non-elementary lower bounds on trees of unbounded degree are known from [DGKS07].

An obvious task for future research is to close the gap between the upper and lower bounds concerning (PE) and (PH). Another direction for future work is to consider other

generalised quantifiers (instead of modulo counting quantifiers) and study to what extent corresponding generalisations of our results (PE), (PH) and (FV) can be achieved.

Acknowledgement

References

[ADG08] Albert Atserias, Anuj Dawar, and Martin Grohe. Preservation under extensions on well-behaved finite structures. *SIAM J. Comput.*, 38(4):1364–1381, 2008.

[ADK06] Albert Atserias, Anuj Dawar, and Phokion G. Kolaitis. On preservation under homomorphisms and unions of conjunctive queries. *J. ACM*, 53(2):208–237, 2006.

[AG94] Miklós Ajtai and Yuri Gurevich. Datalog vs first-order logic. *J. Comput. Syst. Sci.*, 49(3):562–588, 1994.

[BK12] Benedikt Bollig and Dietrich Kuske. An optimal construction of Hanf sentences. *J. Applied Logic*, 10(2):179–186, 2012, arXiv:1105.5487 [cs.LO].

[CLRS09] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, 3rd edition, 2009.

[CM77] Ashok K. Chandra and Philip M. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *Proceedings of the 9th Annual ACM Symposium on Theory of Computing, STOC'77*, pages 77–90, 1977.

[Daw10] Anuj Dawar. Homomorphism preservation on quasi-wide classes. *J. Comput. Syst. Sci.*, 76(5):324–332, 2010, arXiv:0811.4497 [cs.LO].

[DGKS07] Anuj Dawar, Martin Grohe, Stephan Kreutzer, and Nicole Schweikardt. Model theory makes formulas large. In *Automata, Languages and Programming, 34th International Colloquium, ICALP'07*, volume 4596 of *Lecture Notes in Computer Science*, pages 913–924. Springer, 2007. Full version available as preprint NI07003-LAA, Isaac Newton Institute of Mathematical Sciences (2007).

[EF99] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 1999.

[FG04] Markus Frick and Martin Grohe. The complexity of first-order and monadic second-order logic revisited. *Ann. Pure Appl. Logic*, 130(1-3):3–31, 2004.

[FG06] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Springer, 2006.

[Gai82] Haim Gaifman. On local and non-local properties. In J. Stern, editor, *Proceedings of the Herbrand Symposium, Logic Colloquium '81*, pages 105–135. North Holland, 1982.

[GJL15] Stefan Göller, Jean Christoph Jung, and Markus Lohrey. The complexity of decomposing modal and first-order theories. *ACM Trans. Comput. Log.*, 16(1:9), 2015.

[HHS14] Frederik Harwath, Lucas Heimberg, and Nicole Schweikardt. Preservation and decomposition theorems for bounded degree structures. In *Joint Meeting of the 23rd EACSL Annual Conference on Computer Science Logic (CSL) and the 29th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS'14*, pages 49:1–49:10. ACM, 2014.

[HKS13] Lucas Heimberg, Dietrich Kuske, and Nicole Schweikardt. An optimal Gaifman normal form construction for structures of bounded degree. In *Proceedings of the 28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS'13*, pages 63–72, 2013. Full version available at `http://www2.informatik.hu-berlin.de/logik/research/HKS13.pdf`.

[HN04] Pavol Hell and Jaroslav Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.

[Hod93] Wilfrid Hodges. *Model Theory*. Cambridge University Press, 1993.

[Kre11] Stephan Kreutzer. Algorithmic meta-theorems. In *Finite and Algorithmic Model Theory*. Cambridge University Press, 2011, arXiv:0902.3616 [cs.LO]. London Mathematical Society Lecture Notes, No. 379.

[Lib04] Leonid Libkin. *Elements of Finite Model Theory*. Springer, 2004.

[Mak04]   Johann A. Makowsky. Algorithmic uses of the Feferman-Vaught theorem. *Ann. Pure Appl. Logic*,
           126(1-3):159–213, 2004.
[Nur00]   Juha Nurmonen. Counting modulo quantifiers on finite structures. *Inf. Comput.*, 160(1-2):62–87,
           2000.
[Ros08]   Benjamin Rossman. Homomorphism preservation theorems. *J. ACM*, 55(3), 2008.