

Preservation and decomposition theorems for bounded degree structures

Frederik Harwath Lucas Heimberg Nicole Schweikardt

Institut für Informatik

Goethe-Universität Frankfurt am Main

{harwath, heimberg, schweika}@cs.uni-frankfurt.de

Abstract

We provide elementary algorithms for two preservation theorems for first-order sentences with modulo m counting quantifiers ($\text{FO}+\text{MOD}_m$) on the class $\mathcal{C}_d$ of all finite structures of degree at most d : For each $\text{FO}+\text{MOD}_m$ -sentence that is preserved under extensions (homomorphisms) on $\mathcal{C}_d$, a $\mathcal{C}_d$ -equivalent existential (existential-positive) FO-sentence can be constructed in 6-fold (4-fold) exponential time. For FO-sentences, the algorithm has 5-fold (4-fold) exponential time complexity. This is complemented by lower bounds showing that for FO-sentences a 3-fold exponential blow-up of the computed existential (existential-positive) sentence is unavoidable.

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

Categories and Subject Descriptors F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—Computational Logic, Model Theory

General Terms Theory, Algorithms

Keywords computational logic, first-order logic, modulo counting quantifiers, structures of bounded degree, Hanf locality, elementary algorithms, preservation theorems, existential preservation, homomorphism preservation, Łoś-Tarski, Feferman-Vaught

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 *Łoś-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, [2, 6, 14] 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 [3] 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 [7]. A notable exception affects the Łoś-Tarski theorem for the class of acyclic finite structures of degree at most d , for which [7] 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 $\mathcal{C}$ 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 $\text{FO}+\text{MOD}_m$ of first-order logic with modulo m counting quantifiers (for these, it requires 6-fold exponential time). 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 $\mathcal{C}$ (Theorem 3.7).

Our second main result (Theorem 3.2) provides an algorithmic version of a homomorphism preservation theorem, stating that for any class $\mathcal{C}$ 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), for a given $\text{FO}+\text{MOD}_m$ -sentence that is preserved under homomorphisms on $\mathcal{C}$, 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 3.9).

We complement our preservation theorems by (non-matching) lower bounds (Theorem 3.4 and Theorem 3.3), providing a sequence of first-order sentences that are preserved under extensions (homomorphisms) for which the smallest equivalent existential (re-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CSL-LICS '14, July 14–18, 2014, Vienna, Austria.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2886-9...\$15.00.

<http://dx.doi.org/10.1145/2603088.2603130>

spectively, existential-positive) sentences are 3-fold exponentially larger. Both lower bound proofs use particular encodings of numbers by binary trees introduced in [10].

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, for disjoint unions of structures, as well as for cartesian products. Algorithmic versions of such decomposition theorems are typically of the following form (cf., [9, 12]): A given first-order sentence φ that shall be evaluated in the disjoint union or the cartesian product $\mathcal{A}$ of s structures $\mathcal{A}_1, \dots, \mathcal{A}_s$, 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. Our third main result (Theorem 4.1 and Corollary 4.6) shows that for any class $\mathcal{C}$ of structures of degree at most d , such Feferman-Vaught decompositions for disjoint unions and cartesian products can be computed in 3-fold exponential time. This is complemented by a matching lower bound (Theorem 4.4). Our lower bound, again, relies on encodings of numbers by binary trees, now along with a method of [9]. Our algorithm produces a set Δ of so-called *Hanf-formulas* and relies on a result of [4] 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 and lower bounds concerning preservation theorems. Section 4 presents our results concerning Feferman-Vaught decompositions. Due to space limitations, some technical details of our proofs are deferred to the full version of this paper.

2. Preliminaries

We write $\mathbb{Z}$ to denote the set of integers and we write $\mathbb{N}$ to denote the set of non-negative integers. Moreover, we let $\mathbb{N}_{\geq 1} := \mathbb{N} \setminus \{0\}$. For all $m, n \in \mathbb{N}$ with $m \leq n$, the expression $[m, n]$ denotes the set $\{m, \dots, n\}$.

The set of non-negative real numbers is denoted by $\mathbb{R}_{\geq 0}$. For $r > 0$, $\log(r)$ denotes the logarithm of r with respect to base 2. If f is a function from $\mathbb{N}$ to $\mathbb{R}_{\geq 0}$ then $\text{poly}(f(n))$ denotes the class of all functions $g: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ for which there exists a number $c > 0$ such that for all sufficiently large $n \in \mathbb{N}$ we have $g(n) \leq (f(n))^c$.

For an $n \in \mathbb{N}$, we write $\bar{x}$ to denote the tuple $(x_1, \dots, x_n)$. Sometimes we treat tuples as if they were sets; e.g., $a \in \bar{x}$ means $a \in \{x_1, \dots, x_n\}$. For a function $f: A \rightarrow B$ and a tuple $\bar{x} \in A^n$ we write $f(\bar{x})$ to denote the tuple $(f(x_1), \dots, f(x_n))$.

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 $\text{ar}(R)$ 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. The 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 $R^{\mathcal{A}} \subseteq A^{\text{ar}(R)}$ 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., [8, 11]); in particular, $\|\mathcal{A}\| \leq |A|^{\|\sigma\|}$.

For any relational signature σ and σ -structures $\mathcal{A}$ and $\mathcal{B}$, $\mathcal{B}$ is a *substructure* of $\mathcal{A}$ if $B \subseteq A$ and $R^{\mathcal{B}} \subseteq R^{\mathcal{A}}$ for each $R \in \sigma$. $\mathcal{B}$ is an *induced substructure* of $\mathcal{A}$ if $\mathcal{B}$ is a substructure of $\mathcal{A}$

and $R^{\mathcal{B}} := R^{\mathcal{A}} \cap B^{\text{ar}(R)}$ for each $R \in \sigma$. We 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 $\mathcal{A} \cap B$ to denote the substructure of $\mathcal{A}$ induced by $A \cap B$. Furthermore, if $A \setminus B \neq \emptyset$ then $\mathcal{A} \setminus B$ is the substructure $\mathcal{A} \cap (A \setminus B)$ of $\mathcal{A}$ obtained by deleting all elements from B .

We use the standard notation concerning first-order logic and extensions thereof, cf. [8, 11]. By $\text{qr}(\varphi)$ we denote the *quantifier rank* of φ , i.e., the maximum nesting depth of quantifiers occurring in a formula φ . By $\text{free}(\varphi)$ we denote the set of all free variables of φ . A *sentence* is a formula φ with $\text{free}(\varphi) = \emptyset$. A σ -structure $\mathcal{A}$ is called a *model* of a sentence φ if φ is satisfied in $\mathcal{A}$.

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

By $\text{FO}(\sigma)$ we denote the class of all first-order formulas of signature σ . The extension of $\text{FO}(\sigma)$ by *modulo counting quantifiers* is defined as follows: Let m be an integer such that $m \geq 2$. We write $\exists^{0 \bmod m}$ to denote the *modulo m counting quantifier*. A formula of the form $\exists^{0 \bmod m} y \psi(\bar{x}, y)$ is satisfied by a σ -structure $\mathcal{A}$ and an interpretation $\bar{a}$ of the variables $\bar{x}$ if, and only if, the number of elements $b \in A$ such that $\mathcal{A} \models \psi[\bar{a}, b]$ is a multiple of m . For a fixed number m we write $\text{FO}+\text{MOD}_m(\sigma)$ to denote the extension of $\text{FO}(\sigma)$ with the modulo m counting quantifier. Note that the quantifier rank $\text{qr}(\varphi)$ of an $\text{FO}+\text{MOD}_m$ -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 $\text{FO}+\text{MOD}_m(\sigma)$ -formula φ is its length when viewed as a word over the alphabet $\sigma \cup \{\exists^{0 \bmod m}, \exists, \forall, \neg, \wedge, \vee, \rightarrow, \leftrightarrow, (,)\} \cup \{, \} \cup \text{Var}$, 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 $(a_1, \dots, a_{\text{ar}(R)}) \in R^{\mathcal{A}}$ such that $a, b \in \{a_1, \dots, a_{\text{ar}(R)}\}$.

The *distance* $\text{dist}^{\mathcal{A}}(a, b)$ between two elements $a, b \in A$ in a σ -structure $\mathcal{A}$ is the minimal length (i.e., the number of edges) of a path from a to b in the Gaifman graph $G_{\mathcal{A}}$.

For $r \geq 0$ and $a \in A$, the *r -neighbourhood* of a in $\mathcal{A}$ is the set $N_r^{\mathcal{A}}(a) = \{b \in A : \text{dist}^{\mathcal{A}}(a, b) \leq r\}$. The *r -neighbourhood* of a set $W \subseteq A$ is the set $N_r^{\mathcal{A}}(W) := \bigcup_{a \in W} N_r^{\mathcal{A}}(a)$. If $\bar{a} = (a_1, \dots, a_n)$, we write $N_r^{\mathcal{A}}(\bar{a})$ instead of $N_r^{\mathcal{A}}(\{a_1, \dots, a_n\})$.

2.3 Bounded structures

The *degree* of a σ -structure $\mathcal{A}$ is the degree of its Gaifman graph $G_{\mathcal{A}}$. Let $\nu: \mathbb{N} \rightarrow \mathbb{N}$ be a function. A σ -structure $\mathcal{A}$ is *ν -bounded* if $|N_r^{\mathcal{A}}(a)| \leq \nu(r)$ for all $r \in \mathbb{N}$ and $a \in A$. Clearly, if $\mathcal{A}$ is ν -bounded, then it has degree $\leq \nu(1)-1$. On the other hand, if $\mathcal{A}$ has degree $\leq d$, then $\mathcal{A}$ is ν_d -bounded for $\nu_d: \mathbb{N} \rightarrow \mathbb{N}$ with $\nu_d(r) = 1 + d \cdot \sum_{i=0}^{r-1} (d-1)^i$. Thus, $\mathcal{A}$ has degree $\leq d$ 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 $\nu: \mathbb{N} \rightarrow \mathbb{N}$ 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 relational signature, let $\mathcal{A}$ be a σ -structure, let $n \geq 1$, and let $\bar{a} \in A^n$. Let $c_1, \dots, c_n$ be distinct constant symbols. For $r \geq 1$, the r -sphere around $\bar{a}$ is the $\sigma \cup \{c_1, \dots, c_n\}$ -structure $\mathcal{N}_r^{\mathcal{A}}(\bar{a}) := (\mathcal{A} \cap N_r^{\mathcal{A}}(\bar{a}), \bar{a})$, where the constant symbols $c_1, \dots, c_n$ are interpreted by the elements $a_1, \dots, a_n$.

An r -sphere with n centres is a $\sigma \cup \{c_1, \dots, c_n\}$ -structure $\tau = (B, \bar{b})$ with $\bar{b} \in B^n$ and $B = N_r^B(\bar{b})$. We say that τ is realised by $\bar{a}$ in $\mathcal{A}$ iff $\mathcal{N}_r^{\mathcal{A}}(\bar{a})$ is isomorphic to τ . By $\tau(\mathcal{A})$ we denote the set of all $\bar{a} \in A^n$ 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 $\text{sph}_\tau(\bar{x})$ of size $\text{poly}(n \nu(r))$ such that for all σ -structures $\mathcal{A}$ we have $\text{sph}_\tau(\mathcal{A}) = \tau(\mathcal{A})$.

Unless otherwise indicated, we assume an r -sphere to have only one centre. Note that, up to isomorphism, the number of possible r -spheres (with one centre), that are realisable in ν -bounded σ -structures, can easily be seen to be bounded by $2^{\text{poly}(\nu(r))}$, where the degree of the polynomial is linear in the size of σ .

2.4 Disjoint unions

For a relational signature σ , the disjoint union $\mathcal{A} = \mathcal{A}_1 \dot{\cup} \dots \dot{\cup} \mathcal{A}_s$ of σ -structures $\mathcal{A}_1, \dots, \mathcal{A}_s$ (for $s \in \mathbb{N}_{\geq 1}$) is defined (up to isomorphism) as follows: Let A be a set of size $|A_1| + \dots + |A_s|$ and let, for each $i \in [1, s]$, $f_i: A_i \rightarrow A$ be an injective function such that $f_1(A_1), \dots, f_s(A_s)$ is a partition of A . Now, $\mathcal{A}$ is the σ -structure with universe A and $R^{\mathcal{A}} = \bigcup_{i=1}^s \{f_i(\bar{a}) : \bar{a} \in R^{A_i}\}$, for each $R \in \sigma$. The mapping of $\mathcal{A}$ is the function $\pi: A \rightarrow (A_1 \dot{\cup} \dots \dot{\cup} A_s)$ with $\pi(a) = b$ iff there is an $i \in [1, s]$ such that $f_i(b) = a$.

If the universes of $\mathcal{A}_1, \dots, \mathcal{A}_s$ are pairwise disjoint, we let $A := A_1 \dot{\cup} \dots \dot{\cup} A_s$, and hence 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

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 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 $\exists x_1 \dots \exists x_n \varphi$, where φ is quantifier-free. It is straightforward to see that every existential FO(σ)-sentence is preserved under extensions on the class of all σ -structures.

A homomorphism of σ -structures $\mathcal{A}$ and $\mathcal{B}$ is a mapping $h: A \rightarrow B$ such that for each relation $R \in \sigma$ with $r := \text{ar}(R)$ and all $(a_1, \dots, a_r) \in A^r$, if $(a_1, \dots, a_r) \in R^{\mathcal{A}}$ then $(h(a_1), \dots, h(a_r)) \in R^{\mathcal{B}}$. 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 $\neg, \rightarrow, \leftrightarrow$. 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 the class of all σ -structures.

Throughout this section, let σ be a fixed finite relational signature and let $\nu: \mathbb{N} \rightarrow \mathbb{N}$ be a fixed time-constructible strictly increasing function that is at most 1-fold exponential.

In this section, we explore the complexity of constructing existential (existential-positive) FO(σ)-sentences for FO+MOD $_m$ (σ)-sentences that are preserved under extensions (homomorphisms) on a class of ν -bounded σ -structures that is closed under disjoint

unions and closed under induced substructures (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 $d \in \mathbb{N}$, 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.

3.1 Summary of this section's main results

Upper bounds 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.

	extensions		homomorphisms
	FO+MOD $_m$	FO	FO+MOD $_m$
exponential ν	6-exp	5-exp	4-exp
polynomial ν	4-exp	3-exp	3-exp

Table 1. The time complexity of our algorithms. The class $\mathfrak{C}_d$ of all structures of degree $\leq d$ for $d \geq 3$ is a case of “exponential ν ”.

The precise statement of this section's first main result reads as follows; a proof is given in Subsection 3.2 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 which, given an input FO+MOD $_m$ (σ)-sentence φ of quantifier rank $q \geq 0$, for an $m \geq 2$, constructs in time*

$$\|\varphi\| \cdot 2^{m \cdot 2^{\nu(2^{\text{poly}(\nu(3^q))})}} \quad (1)$$

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

If φ does not contain any modulo counting quantifier (i.e., φ is an FO(σ)-sentence), the algorithm uses only time

$$\|\varphi\| \cdot 2^{\nu(2^{\text{poly}(\nu(3^q))})}. \quad (2)$$

Furthermore, the degree of the polynomials in the expressions (1) and (2) is linear in the size of the signature σ .

Consequently, if the function ν is exponential and φ is an FO+MOD $_m$ (σ)-sentence, then the algorithm uses 6-fold exponential time — e.g, if $\nu = \nu_d$ for $d \geq 3$, it uses time

$$2^{m \cdot 2^{d^2 2^{O(\|\varphi\|)}}}, \quad \text{and 5-fold exponential time } 2^{d^2 2^{d^2 2^{O(\|\varphi\|)}}},$$

for each FO(σ)-sentence φ . If ν is polynomial and φ is an FO+MOD $_m$ (σ)-sentence, then the construction requires only 4-fold exponential time

$$2^{m \cdot 2^{2^2 2^{O(\|\varphi\|)}}}, \quad \text{and 3-fold exponential time } 2^{2^2 2^{O(\|\varphi\|)}}$$

for each FO(σ)-sentence φ .

This section's second main result reads as follows; a proof is given in Subsection 3.3 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 $t: \mathbb{N} \rightarrow \mathbb{N}$. There is an algorithm which, given an input FO+MOD $_m$ (σ)-sentence φ of quantifier rank $q \geq 0$, for an $m \geq 2$, constructs in time*

$$2^{O(\|\varphi\|) \cdot 2^{\|\sigma\| \cdot \nu(3^{O(q)})}} \cdot t(2^{\|\sigma\| \cdot \nu(3^{O(q)})}) \quad (3)$$

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

For example, let $\mathfrak{C}_\nu$ be the class of all ν -bounded finite σ -structures. If ν is exponential and φ is an $\text{FO}+\text{MOD}_m(\sigma)$ -sentence, then the construction can be carried out in 4-fold exponential time — e.g. if $\nu = \nu_d$ for $d \geq 3$ (i.e., $\mathfrak{C}_\nu$ is the class of all finite structures of degree at most d), in time

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

If ν is polynomial and φ is an $\text{FO}+\text{MOD}_m(\sigma)$ -sentence, then the construction requires only 3-fold exponential time

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

Lower bounds Our upper bounds 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 3.3. 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 an $\epsilon > 0$ and a sequence $(\varphi_h)_{h \geq 1}$ of $\text{FO}(\sigma)$ -sentences of increasing size such that for each $h \geq 1$ 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 3.4. Let $\sigma' := \{S_0, S_1\} \cup \{V_M : M \subseteq \{0, 1\}\}$, where S_0, S_1 are binary 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 an $\epsilon > 0$ and a sequence $(\varphi'_h)_{h \geq 1}$ of $\text{FO}(\sigma')$ -sentences of increasing size such that for each $h \geq 1$ 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 the previous theorems 3.3 and 3.4, 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.

The key idea for the proofs of Theorem 3.3 and Theorem 3.4, which will be included in the full version of this paper, is an encoding of numbers by binary trees, introduced in [10]. The main challenge here is to find sequences of sentences that not only have large minimal models but are also preserved under extensions (homomorphisms). Towards this end, the auxiliary unary relation symbols in σ and σ' are introduced to interpret binary tree encodings in ordered binary forests.

3.2 Preservation under extensions: Proof of Theorem 3.1

This subsection is devoted to the proof of Theorem 3.1. Consider a class $\mathfrak{C}$ of σ -structures and an $\text{FO}+\text{MOD}_m(\sigma)$ -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 .

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

In both cases, the following generalisation of Hanf's theorem (see e.g. [8, 11]) by Nurmonen [13] to sentences with modulo counting quantifiers is used. 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 $\text{FO}+\text{MOD}_m(\sigma)$ -sentences of quantifier rank at most q .

Theorem 3.5 (Theorem 3.4 in [13]). Let $\mathcal{A}, \mathcal{B}$ be σ -structures. Let $m \geq 2$ and $q \geq 0$. Suppose that for some $e \geq 0$ 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

$$|\tau(\mathcal{A})| = |\tau(\mathcal{B})| \quad \text{or} \quad |\tau(\mathcal{A})| \geq t \text{ and } |\tau(\mathcal{B})| \geq t,$$

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

A set of elements B in a σ -structure $\mathcal{A}$ is r -scattered, for an $r \geq 0$, 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 3.6. Let $\mathcal{A}$ be a ν -bounded σ -structure and let $m, r \geq 0$. If $|A| > (m-1) \cdot \nu(2r)$ then there exists an r -scattered subset of A of cardinality m .

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

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

The main combinatorial contribution of this subsection is provided by the following theorem:

Theorem 3.7. Let $\mathfrak{C}_\nu$ be a class of ν -bounded σ -structures that is closed under disjoint unions and induced substructures. Let $m \geq 2$ and let φ be an $\text{FO}+\text{MOD}_m(\sigma)$ -sentence of quantifier rank $q \geq 0$ that is preserved under extensions on $\mathfrak{C}_\nu$. There is a number

$$N(\nu, m, q) \in m \cdot 2^{\nu(2^{\text{poly}(\nu(3^q))})}$$

such that every $\mathfrak{C}_\nu$ -minimal model of φ has size $\leq N(\nu, m, q)$.

Proof. Let $r := 3^q$ and let s be the number of non-isomorphic r -spheres (with one centre) in structures in $\mathfrak{C}_\nu$. Let $R := 2sr$ and let S be the number of non-isomorphic R -spheres (with one centre) in structures in $\mathfrak{C}_\nu$. Finally, let $t := q \cdot (\nu(r)+1) + 1$ be the threshold from Theorem 3.5 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 $|A| > (2Stm-1) \cdot \nu(2R)$. By Lemma 3.6, 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.




Figure 1. The first six elements of the sequence $(C_i)_{i \geq 0}$. Note that for all even i , C_i is a disjoint extension of $\mathcal{A}$ with an induced substructure of $\mathcal{A}$ while for each odd i , C_i is a proper induced substructure of $\mathcal{A}$.

An r -sphere τ is *frequent* in a structure $\mathcal{A}$ if $|\tau(\mathcal{A})| \geq t$. Otherwise, it is *rare*. 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 $2tm \geq t$ times in $\mathcal{A}$.

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

$$\mathcal{N}_R^{\mathcal{A}}(a) \cong \mathcal{N}_R^{\mathcal{A}}(b) \quad \text{for all } a \in X \text{ and } b \in X' \setminus X \quad (4)$$

the following holds for each r -sphere τ : If $|\tau(\mathcal{A}) \cap N_{R-r}^{\mathcal{A}}(X)| \geq t$, also $|\tau(\mathcal{A} \setminus N_{R-r}^{\mathcal{A}}(X))| \geq t$. That is, 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 $(C_i)_{i \geq 0}$ and $(\mathcal{D}_{2i})_{i \geq 0}$ of structures in $\mathfrak{C}_\nu$. Let $C_0 := \mathcal{A}$ and $C_1 := \mathcal{A} \setminus X$. For each even $i > 1$, let $\mathcal{D}_i := C_{i-1} \cap N_{2(i-1)r}^{\mathcal{A}}(X)$. Note that for odd i , $\mathcal{D}_i$ is neither defined nor required. For each $i \geq 2$, let

$$C_i := \begin{cases} \mathcal{A} \dot{\cup} \mathcal{D}'_i & \text{if } i \text{ is even,} \\ (\mathcal{A} \setminus N_{2(i-1)r}^{\mathcal{A}}(X)) \dot{\cup} \mathcal{D}_{i-1} & \text{if } i \text{ is odd,} \end{cases}$$

where $\mathcal{D}'_i$ is a structure isomorphic to $\mathcal{D}_i$ whose universe is disjoint with $\mathcal{A}$. For all even i , C_i is a disjoint extension of $\mathcal{A}$, and for all odd i , 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 $1 \leq i < s$, there is a number $k \in \mathbb{Z}$ such that $|\tau(\mathcal{A} \setminus N_{2(i-1)r}^{\mathcal{A}}(X))| = |\tau(\mathcal{A})| + ktm$ and therefore $|\tau(\mathcal{A} \setminus N_{2(i-1)r}^{\mathcal{A}}(X))| \equiv |\tau(\mathcal{A})| \pmod{m}$. Furthermore, since $|X| \equiv 0 \pmod{m}$, it is straightforward to see that $|\tau(\mathcal{D}'_i)| = |\tau(\mathcal{D}_i)| \equiv 0 \pmod{m}$ for all even $i \leq s$. This immediately proves the following:

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

Below, we will prove the following:

Claim 2. The following holds for all $i < s$ and every r -sphere τ :

- (a) If τ is frequent in C_i , it is also frequent in C_{i+1} .
- (b) If τ is rare in C_i and rare in C_{i+1} , then $|\tau(C_i)| = |\tau(C_{i+1})|$.

Whereas every r -sphere that is frequent in C_i is also frequent in C_{i+1} , the opposite is not true: There may be r -spheres that are rare in C_i but that occur frequently in C_{i+1} .

However, since there are only s non-isomorphic r -spheres in structures from $\mathfrak{C}_\nu$, and C_0 already contains frequent r -spheres, Claim 2 (a) implies that there has to be an $i < s$ such that all frequent r -spheres of C_{i+1} are frequent already in C_i . Thus, for this particular i we know that any r -sphere is either frequent in C_{i+1} and in C_i or it is rare in C_{i+1} and in C_i .

Hence, with Claim 2 (b) it follows that for every r -sphere τ , either $|\tau(C_i)| = |\tau(C_{i+1})|$ or τ is frequent in C_i and C_{i+1} .

Together with Claim 1, we can conclude from Theorem 3.5 that $C_i \equiv_m^{q} C_{i+1}$. Therefore, by using Claim 2, the proof of Theorem 3.7 can be completed as follows:

In case that i is even, we let $\mathcal{B} := C_i$ and $\mathcal{A}' := C_{i+1}$; and in case that i is odd, we let $\mathcal{B} := C_{i+1}$ and $\mathcal{A}' := C_i$. 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, $|A| \leq (2Stm-1) \cdot \nu(2R)$. Recall from Section 2.3 that $S \in 2^{\text{poly}(\nu(R))}$. Using estimates for R and t , it follows that

$$N(\nu, m, q) := (2Stm-1) \cdot \nu(2R) \in m \cdot 2^{\nu(2^{\text{poly}(\nu(3q))})}.$$

All that remains to be done to finish the proof of Theorem 3.7 is to prove Claim 2.

Proof of Claim 2: Observe that, for all $i, j \leq s$,

$$C_i \cap (\mathcal{A} \setminus N_{R-2r}^{\mathcal{A}}(X)) \cong C_j \cap (\mathcal{A} \setminus N_{R-2r}^{\mathcal{A}}(X)). \quad (5)$$

Let $i < s$. For the proof of Claim 2 (b) let τ be an r -sphere that is rare in C_i and C_{i+1} . Since X is an R -scattered set of size $\geq t$, the rareness of τ implies that $\tau(C_i)$ and $\tau(C_{i+1})$ are subsets of $\mathcal{A} \setminus N_{R-r}^{\mathcal{A}}(X)$. Hence, (5) implies that $|\tau(C_i)| = |\tau(C_{i+1})|$. This proves Claim 2 (b).

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

Even i : Recall that $C_0 = \mathcal{A}$ and for each even $i \geq 2$, $C_i = \mathcal{A} \dot{\cup} \mathcal{D}'_i$. Let τ be an r -sphere that is frequent in C_i .

If τ is realised in C_i by an element in $N_{R-r}^{\mathcal{A}}(X)$ then, since X is R -scattered and $|X| \geq t$, $|\tau(\mathcal{A}) \cap N_{R-r}^{\mathcal{A}}(X)| \geq t$. By observation (4) we obtain that also $|\tau(\mathcal{A}) \cap N_{R-r}^{\mathcal{A}}(X' \setminus X)| \geq t$. Since $X \subseteq X'$, X' is R -scattered, and $2ir < R$, we obtain that $|\tau(\mathcal{A} \setminus N_{2ir}^{\mathcal{A}}(X))| \geq t$. Thus, τ is also frequent in C_{i+1} .

If τ is realised in C_i by an element in $\mathcal{D}'_i$ (for $i \geq 2$), then, since $\mathcal{D}'_i \cong \mathcal{D}_i$, τ is realised in C_{i+1} also by an element in $\mathcal{D}_i \subseteq N_{R-r}^{\mathcal{A}}(X)$. Thus, τ is also frequent in C_{i+1} .

Otherwise, we know that $|\tau(\mathcal{A}) \setminus N_{R-r}^{\mathcal{A}}(X)| \geq t$. Then it follows from (5) that τ is also frequent in C_{i+1} .

Odd i : Recall that $C_1 = \mathcal{A} \setminus X$ and that, for each odd $i \geq 3$, $C_i = (\mathcal{A} \setminus N_{2(i-1)r}^{\mathcal{A}}(X)) \dot{\cup} \mathcal{D}_{i-1}$. Let τ be an r -sphere that is frequent in C_i .

If τ is realised in C_i by an element in $N_{2(i-1)r+r}^{\mathcal{A}}(X)$ then, since X is R -scattered and $|X| \geq t$, $|\tau(C_i) \cap N_{2(i-1)r+r}^{\mathcal{A}}(X)| \geq t$. From $\mathcal{D}_{i+1} = C_i \cap N_{2(i-1)r+2r}^{\mathcal{A}}(X)$ we obtain that $|\tau(\mathcal{D}_{i+1})| \geq t$. Since $\mathcal{D}'_{i+1} \cong \mathcal{D}_{i+1}$, we have $|\tau(\mathcal{D}'_{i+1})| \geq t$. Consequently, since $C_{i+1} = \mathcal{A} \dot{\cup} \mathcal{D}'_{i+1}$, also $|\tau(C_{i+1})| \geq t$. Thus, τ is frequent in C_{i+1} .

Otherwise, we know that $|\tau(C_i) \cap (\mathcal{A} \setminus N_{2(i-1)r+r}^{\mathcal{A}}(X))| \geq t$. Since for every element $a \in \mathcal{A} \setminus N_{2(i-1)r+r}^{\mathcal{A}}(X)$ the r -spheres of

a in $\mathcal{A} \setminus N_{2(i-1)r}^{\mathcal{A}}(X)$ and $\mathcal{A}$ are isomorphic, 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 3.7. $\square$

The following lemma is a slight generalisation of Lemma 8.4 in [7] 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 $N \geq 0$ and an $\text{FO}+\text{MOD}_m(\sigma)$ -sentence φ of quantifier rank $q \geq 0$, for an $m \geq 2$, constructs in time $\|\varphi\| \cdot 2^{O(N \cdot q)}$ an existential $\text{FO}(\sigma)$ -sentence ψ such that the following holds: If φ is preserved under extensions on $\mathfrak{C}$ and every $\mathfrak{C}$ -minimal model of φ has at most N elements, then ψ is $\mathfrak{C}$ -equivalent to φ .*

Furthermore, if φ does not contain any modulo counting quantifier, the construction only takes time $O(\|\varphi\| \cdot N^q)$.

Proof. Let $m \geq 2$ and let φ be an $\text{FO}+\text{MOD}_m(\sigma)$ -sentence that is preserved under extensions on $\mathfrak{C}$ and assume that, for a number $N \geq 0$, each $\mathfrak{C}$ -minimal model of φ has at most N elements.

Let $x_1, \dots, x_k$ be, for a $k \geq 1$, an enumeration of all variables that appear in φ and suppose, without loss of generality, that every variable in φ is only bound in one place. Let $y_1, \dots, y_N$ be a collection of N new variables distinct from $x_1, \dots, x_k$.

We construct from the sentence φ a new quantifier-free formula with free variables $y_1, \dots, y_N$. To this end, for each formula ψ with free variables among $x_1, \dots, x_k$ and each function $s: [1, k] \rightarrow [1, N]$, a formula $(\psi)_s$ is defined inductively as follows:

- If ψ is atomic then $(\psi)_s$ is the formula obtained by replacing all occurrences of variables x_i in ψ by $y_{s(i)}$.
- If $\psi = \neg\psi'$ then $(\psi)_s := \neg(\psi')_s$. For each Boolean connective $\star \in \{\wedge, \vee, \rightarrow, \leftrightarrow\}$, $(\psi' \star \psi'')_s := (\psi')_s \star (\psi'')_s$.
- If $\psi = \exists x_i \psi'$ then $(\psi)_s := \bigvee_{j=1}^N (\psi')_{s[i \rightarrow j]}$ where $s[i \rightarrow j]$ is the function s' that agrees with s on all values except i and $s'(i) = j$.
- If $\psi = \forall x_i \psi'$ then $(\psi)_s := \bigwedge_{j=1}^N (\psi')_{s[i \rightarrow j]}$.
- If $\psi = \exists^{0 \bmod m} x_i \psi'$ then $(\psi)_s :=$

$$\bigvee_{\substack{M \subseteq [1, N], \\ |M| \equiv 0 \bmod m}} \left(\bigwedge_{\substack{j, k \in M, \\ j \neq k}} \neg y_j = y_k \wedge \bigwedge_{j=1}^N \gamma_j^M \right),$$

where, for each $j \in M$, $\gamma_j^M := (\psi')_{s[i \rightarrow j]}$ and, for each $j \in [1, N] \setminus M$, $\gamma_j^M := \neg(\psi')_{s[i \rightarrow j]}$.

Now, let φ' be the sentence $\exists y_1 \dots \exists y_N (\varphi)_s$ where, since φ does not have any free variable, s can be an arbitrary function from $[1, k]$ to $[1, N]$. It is easily seen that a σ -structure $\mathcal{A} \in \mathfrak{C}$ satisfies φ' iff $\mathcal{A}$ contains an induced substructure $\mathcal{A}'$ with $\leq N$ elements which satisfies φ . It is not difficult to show that $\mathcal{A} \models \varphi'$ iff $\mathcal{A} \models \varphi$:

For the direction “ $\implies$ ” assume that $\mathcal{A} \models \varphi'$. Thus, $\mathcal{A}$ contains an induced substructure $\mathcal{A}'$ with $\mathcal{A}' \models \varphi$. Since $\mathcal{A} \in \mathfrak{C}$ and $\mathfrak{C}$ is closed under induced substructures, also $\mathcal{A}' \in \mathfrak{C}$. As $\mathcal{A}' \models \varphi$, and φ is preserved under extensions on $\mathfrak{C}$, we obtain that also $\mathcal{A} \models \varphi$.

For the direction “ $\impliedby$ ” assume that $\mathcal{A} \models \varphi$. Let $\mathcal{A}'$ be a minimal induced substructure of $\mathcal{A}$ 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, $\mathcal{A}'$ has at most N elements. By construction of φ' we obtain that $\mathcal{A} \models \varphi'$.

For a bound on the size of φ' , 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 N , for first-order quantifiers, and by a factor in $2^{O(N)}$, for modulo counting quantifiers. It follows that $\|\varphi'\| \in \|\varphi\| \cdot 2^{O(N \cdot q)}$, where q is the quantifier rank of φ . If φ does not contain any modulo counting

quantifier, $\|\varphi'\| \in O(\|\varphi\| \cdot N^q)$. It is easy to see that the inductive translation can be also carried out in time $O(\|\varphi'\|)$. $\square$

Theorem 3.1 is now obtained by a straightforward combination of Lemma 3.8 and Theorem 3.7.

Proof of Theorem 3.1. Let $m \geq 2$ and let φ be an $\text{FO}+\text{MOD}_m(\sigma)$ -sentence of quantifier rank $q \geq 0$ that is preserved under extensions on $\mathfrak{C}_\nu$. Let $N = N(\nu, m, q)$ be the upper bound on the size of $\mathfrak{C}_\nu$ -minimal models of φ , obtained from Theorem 3.7.

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

$$\|\varphi\| \cdot 2^{O(N \cdot q)} \subseteq \|\varphi\| \cdot 2^{m \cdot 2^{\nu(2^{\text{poly}(\nu(3^q))})}},$$

and, for each $\text{FO}(\sigma)$ -sentence φ , in time

$$O(\|\varphi\| \cdot N^q) \subseteq \|\varphi\| \cdot 2^{\nu(2^{\text{poly}(\nu(3^q))})}.$$

This concludes the proof of Theorem 3.1. $\square$

3.3 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 3.9. *Let $\mathfrak{C}_\nu$ be a class of ν -bounded σ -structures that is closed under disjoint unions and induced substructures. Let $m \geq 2$ and let φ be an $\text{FO}+\text{MOD}_m(\sigma)$ -sentence of quantifier rank $q \geq 0$ that is preserved under homomorphisms on $\mathfrak{C}_\nu$. There is a number*

$$N(\nu, q) \in 2^{\nu(3^{O(q)})}$$

such that every $\mathfrak{C}_\nu$ -minimal model of φ has size $\leq N(\nu, q)$.

Proof. The proof is similar to the proof of Lemma 7.1 in [1]. However, it does not rely on Gaifman’s theorem but uses Nurmonen’s generalisation of Hanf’s theorem, stated in Theorem 3.5. Towards applying Theorem 3.5, we let $r := 3^q$, let $t := q \cdot (\nu(r)+1) + 1$, and let s be the number of non-isomorphic $2r$ -spheres (with one centre) in structures in $\mathfrak{C}_\nu$.

Let $\mathcal{A}$ be a $\mathfrak{C}_\nu$ -minimal model of φ . Towards a contradiction, assume that $|\mathcal{A}| > s \cdot \nu(4r)$. By Lemma 3.6, $\mathcal{A}$ 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 $a_1, a_2 \in \mathcal{A}$ with disjoint and isomorphic $2r$ -neighbourhoods.

Let $\mathcal{A}' := \mathcal{A} \setminus \{a_1\}$. Clearly, the r -spheres of elements in $\mathcal{A} \setminus N_r^{\mathcal{A}}(a_1)$ 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 3.5, and therefore we have $\mathcal{B} \equiv_m^q \mathcal{B}'$. 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, $|\mathcal{A}| \leq s \cdot \nu(4r)$. Recalling from Section 2.3 that $s \in 2^{\text{poly}(\nu(2r))}$, one obtains that

$$N(\nu, q) := s \cdot \nu(4r) \in 2^{\nu(3^{O(q)})}.$$

This concludes the proof of Theorem 3.9. $\square$

The following lemma provides a construction of existential-positive FO(σ)-sentences for FO+MOD $_m(\sigma)$ -sentences that are preserved under homomorphisms. The proof is an algorithmic version of the proof of Theorem 3.1 in [2].

Lemma 3.10. *Let $\mathfrak{C}$ be a class of σ -structures that is decidable in time $t(n)$ for some function $t: \mathbb{N} \rightarrow \mathbb{N}$. There is an algorithm that, on input of a number $N \geq 0$ and an FO+MOD $_m(\sigma)$ -sentence φ of quantifier rank $q \geq 0$, for an $m \geq 2$, constructs in time*

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

an existential-positive FO(σ)-sentence ψ of size $2^{O(N^{\|\sigma\|})}$, 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., $\gamma_{\mathcal{A}} = \exists x_1 \dots \exists x_{|\mathcal{A}|} \theta_{\mathcal{A}}$ where $x_1, \dots, x_{|\mathcal{A}|}$ are variables representing the elements $a_1, \dots, a_{|\mathcal{A}|}$ of $\mathcal{A}$'s universe A , and $\theta_{\mathcal{A}}$ is the conjunction of all atoms of the form $R(x_{i_1}, \dots, x_{i_r})$ where $R \in \sigma$, $r = \text{ar}(R)$, $i_1, \dots, i_r \in \{1, \dots, |\mathcal{A}|\}$, and $(a_{i_1}, \dots, a_{i_r}) \in R^{\mathcal{A}}$. The well-known Chandra-Merlin theorem [5] 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 $O(\|\mathcal{A}\|) = O(|\mathcal{A}|^{\|\sigma\|})$.

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

1. computes the set M that consists all models of φ with universe $\{1, \dots, n\}$, for $n \leq N$, that belong to $\mathfrak{C}$,
2. if $M = \emptyset$, it outputs the formula $\psi := \text{false}$,
3. if $M \neq \emptyset$, it outputs the formula $\psi := \bigvee_{\mathcal{A} \in M} \gamma_{\mathcal{A}}$.

Clearly, the size of ψ is in

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

where $2^{O(N^{\|\sigma\|})}$ is an upper bound on the number of all σ -structures with universe $\{1, \dots, n\}$ for $n \leq N$. Obviously, ψ is an existential-positive FO(σ)-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 $A = \{1, \dots, n\}$ and $n \leq N$, 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 $t(\|\mathcal{A}\|) \leq t(N^{\|\sigma\|})$. Using the naive model checking algorithm for FO+MOD $_m$, the question whether $\mathcal{A} \models \varphi$ can be answered within time $\|\mathcal{A}\|^{O(\|\varphi\|)} \subseteq N^{\|\sigma\| \cdot O(\|\varphi\|)}$. Since there are at most $2^{O(N^{\|\sigma\|})}$ σ -structures with universe $\{1, \dots, n\}$ and $n \leq N$, the entire computation of M takes time at most

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

$\square$

Theorem 3.2 is now obtained by a straightforward combination of Lemma 3.10 and Theorem 3.9 (in the analogous way as Theorem 3.1 was obtained by combining Lemma 3.8 with Theorem 3.7).

4. Feferman-Vaught decompositions

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

Let $(P_i)_{i \geq 1}$ be a sequence of unary relation symbols which are not already contained in σ . For every $s \geq 1$, σ_s denotes the signature $\sigma \cup \{P_1, \dots, P_s\}$.

Recall that a disjoint union $\mathcal{B} = \mathcal{A}_1 \dot{\cup} \dots \dot{\cup} \mathcal{A}_s$ of σ -structures $\mathcal{A}_1, \dots, \mathcal{A}_s$ involves injective functions $f_i: A_i \rightarrow B$ for each $i \in [1, s]$ such that $f_1(A_1), \dots, f_s(A_s)$ is a partition of B . The disjoint sum $\mathcal{A}_1 \oplus \dots \oplus \mathcal{A}_s$ of $\mathcal{A}_1, \dots, \mathcal{A}_s$ is a σ_s -structure $\mathcal{A}$ that expands the disjoint union of $\mathcal{A}_1, \dots, \mathcal{A}_s$ by the unary relations $P_i^{\mathcal{A}} := f_i(A_i)$, for all $i \in [1, s]$. Clearly, $P_1^{\mathcal{A}}, \dots, P_s^{\mathcal{A}}$ 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 $\mathcal{A}$.

This section provides an algorithm that computes, for every $s \geq 1$ and each FO(σ_s)-formula φ , a *Feferman-Vaught decomposition*, i.e., a decomposition into a Boolean combination of FO(σ)-formulas that is equivalent to φ on each disjoint sum of s structures in $\mathfrak{C}_\nu$. 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.

We also 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 stating 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* [12]. They give conditions for the validity of an FO(σ_s)-formula in a disjoint sum of structures in terms of a Boolean combination of FO(σ)-formulas that speak about the component structures of the disjoint sum.

Let $s \geq 1$ and let $\bar{x}$ be a tuple of $n \geq 0$ variables. For each $i \in [1, s]$, let Δ_i be a finite set of FO(σ)-formulas δ with $\text{free}(\delta) \subseteq \bar{x}$ and let β be a propositional formula with variables from the set $X_D := \{X_{i,\delta} : i \in [1, s], \delta \in \Delta_i\}$. The tuple

$D = (\Delta_1, \dots, \Delta_s, \beta)$ is an s -reduction sequence over $\bar{x}$ (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 $\mathcal{A}_1, \dots, \mathcal{A}_s$ be σ -structures and let $\bar{a}$ be a tuple $(a_1, \dots, a_n)$ of $(\bigcup_{i=1}^s \mathcal{A}_i)^n$. We say that $(\mathcal{A}_1, \dots, \mathcal{A}_s, \bar{a})$ is a *model of the reduction sequence* D , in symbols: $(\mathcal{A}_1, \dots, \mathcal{A}_s, \bar{a}) \models D$, iff $\mu \models \beta$, where $\mu: X_D \rightarrow \{0, 1\}$ is the truth assignment such that for each $i \in [1, s]$ and $\delta \in \Delta_i$,

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

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

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

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

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 $\text{FO}(\sigma_s)$ -formula $\varphi(\bar{x})$ is a Boolean combination of $\text{FO}(\sigma)$ -formulas from sets Δ_i , $i \in [1, s]$. This Boolean combination is equivalent to φ on every s -disjoint sum $\mathcal{A}_1 \oplus \dots \oplus \mathcal{A}_s$ and, for each $i \in [1, s]$, every $\text{FO}(\sigma)$ -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 provides a 3-fold exponential algorithm that computes a disjoint decomposition $D = (\Delta_1, \dots, \Delta_s, \beta)$ for an input $\text{FO}(\sigma_s)$ -formula $\varphi(\bar{x})$ on $\mathfrak{C}_\nu$. The formulas in Δ_i are of a particular form called *Hanf-formulas*, defined as follows.

A *Hanf-formula* with $n \geq 0$ free variables $\bar{x}$ is a formula of the form $\exists^{\geq k} y \text{ sph}_\tau(\bar{x}, y)$ where τ is the isomorphism type of a finite r -sphere (for an $r \in \mathbb{N}$) with $n+1$ centres.

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

$$\exists^{\geq k} y \varphi(\bar{x}, y)$$

as a shorthand for the formula

$$\exists y_1 \dots \exists y_k \left(\bigwedge_{1 \leq i < j \leq k} \neg y_i = y_j \wedge \forall y \left(\bigvee_{1 \leq i \leq k} y = y_i \rightarrow \varphi(\bar{x}, y) \right) \right).$$

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

We are now ready to give a precise statement of this section's main result:

Theorem 4.1. *Let $s \geq 1$. There is an algorithm which, given an input $\text{FO}(\sigma_s)$ -formula $\varphi(\bar{x})$ of quantifier rank $q \geq 0$, constructs in time*

$$s \cdot 2^{\text{poly}(\|\varphi\| \cdot \nu(4^q))} \quad (6)$$

a disjoint decomposition $(\Delta_1, \dots, \Delta_s, \beta)$ for $\varphi(\bar{x})$ on $\mathfrak{C}_\nu$.

Furthermore, the sets $\Delta_1, \dots, \Delta_s$ contain only *Hanf-formulas* over σ , and the degree of the polynomial in the expression (6) is linear in $\|\sigma\| + s$.

Thus, if the function ν is exponential, then φ can be decomposed in 3-fold exponential time — e.g., if $\nu = \nu_d$ for $d \geq 3$, then φ can be decomposed in time

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

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

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

The remainder of this subsection is devoted to the proof of Theorem 4.1.

The first step of the algorithm employs an algorithm by Bollig and Kuske [4] to transform an $\text{FO}(\sigma)$ -formula φ into a $\mathfrak{C}_\nu$ -equivalent formula φ^H in *Hanf normal form*, i.e., a Boolean combination of Hanf-formulas. While [4] provided a 3-fold exponential algorithm that transforms a given $\text{FO}(\sigma)$ -formula ψ into a formula in Hanf normal form that is equivalent to ψ on all σ -structures of degree at most d (for $d \geq 1$), in the slightly more general setting of ν -bounded structures, their proof yields the following (a proof can be found in the full version of [10]).

Theorem 4.2 ([4, 10]). *There is an algorithm which transforms an input $\text{FO}(\sigma)$ -formula φ of quantifier rank $q \geq 0$ in time*

$$2^{\text{poly}(\|\varphi\| \cdot \nu(4^q))} \quad (7)$$

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 $\exists^{\geq k} y \text{ sph}_\tau(\bar{x}, y)$, where $k \leq \|\varphi\| \cdot (q+1) \cdot \nu(4^q)$, τ is a ν -bounded r -sphere of radius $r \leq 4^q$, and $|\bar{x}| \leq |\text{free}(\varphi)|$.

Furthermore, the degree of the polynomial in the expression (7) is linear in the size of the signature σ .

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 $s \geq 1$. There is an algorithm which, given an input Hanf-formula $\psi(\bar{x}) := \exists^{\geq k} x_{n+1} \text{ sph}_\tau(\bar{x}, x_{n+1})$ with $n \geq 0$ free variables, where $k \geq 1$ and, for an $r \geq 0$, τ is an r -sphere with $n+1$ centres over the signature σ_s , constructs in time*

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

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

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

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

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 $P_1^{\mathcal{T}}, \dots, P_s^{\mathcal{T}}$, 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 $\mathcal{T}$ is an induced substructure of an s -disjoint sum or not.

Case 1: $\mathcal{T}$ is not an induced substructure of an s -disjoint sum. Then, there is no s -disjoint sum $\mathcal{A}$ and $\bar{a} \in \mathcal{A}^{n+1}$ such that $\mathcal{N}_r^{\mathcal{A}}(\bar{a}) \cong \tau$. 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 $\psi(\bar{x})$. 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 $N_r^T(c_i)$ belong to P_j^T . For each $i \in [1, s]$, let $\bar{c}_i$ be the subsequence of $\bar{c}$ induced by all $c_j \in \bar{c} \cap P_i^T$, and let $\bar{x}_i$ be the subsequence of $\bar{x}$ induced by all x_j such that $c_j \in \bar{c} \cap P_i^T$. W.l.o.g., $c_{n+1} \in P_s^T$.

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

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

$$\bigwedge_{i \in [1, s], \bar{c} \cap P_i^T \neq \emptyset} X_{i, \delta_i}.$$

It remains to define the formulas $\delta_i(\bar{x}_i)$. Let $\mathcal{T}_i := \mathcal{N}_r^T(\bar{c}_i)$ and let τ_i be the r -sphere $(\mathcal{T}_i, \bar{c}_i)$. Let $\mathcal{T}'_i := (\mathcal{T}_i)_{|\sigma}$ be the σ -reduct of $\mathcal{T}_i$, and let τ'_i be the r -sphere $(\mathcal{T}'_i, \bar{c}_i)$.

Recall that we assume w.l.o.g. that $c_{n+1} \in P_s^T$. For $i=s$ we let

$$\delta_s(\bar{x}_s) := \exists^{\geq k} x_{n+1} \text{ sph}_{\tau'_s}(\bar{x}_s, x_{n+1}).$$

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

$$\delta_i(\bar{x}_i) := \exists^{\geq 1} x_{n+1} \text{ sph}_{(\mathcal{T}'_i, \bar{c}_i c)}(\bar{x}_i, x_{n+1}),$$

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

Claim 1. $(\Delta_1, \dots, \Delta_s, \beta)$ is a disjoint decomposition for $\psi(\bar{x})$ on the class of all σ -structures.

Proof of Claim 1: Let $\mathcal{A} = \mathcal{A}_1 \oplus \dots \oplus \mathcal{A}_s$ be the disjoint sum of σ -structures $\mathcal{A}_1, \dots, \mathcal{A}_s$, let π be the mapping of $\mathcal{A}$, and let $\bar{a} = (a_1, \dots, a_n) \in A^n$. For every $i \in [1, s]$, let $\bar{a}_i$ be the subsequence of $\bar{a}$ induced by all $a_j \in \bar{a} \cap P_i^A$. By definition of a disjoint decomposition, we have to show that

$$(\mathcal{A}, \bar{a}) \models \exists^{\geq k} x_{n+1} \text{ sph}_{(\mathcal{T}, \bar{c})}(\bar{x}, x_{n+1}) \quad (= \psi(\bar{x}))$$

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

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

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

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

“ $\impliedby$ ”: Suppose that (a) and (b) hold for all $i \in [1, s]$ with $\bar{c} \cap P_i^T \neq \emptyset$. Then, for each $i \in [1, s-1]$, (b) and the definition of δ_i imply that $\mathcal{N}_r^{A_i}(\pi(\bar{a}_i)) \cong \tau'_i$.

Furthermore, by (b) and the definition of δ_s , there are at least k elements $a \in A_s$ such that $\mathcal{N}_r^{A_s}(\pi(\bar{a}_s a)) \cong \tau'_s$.

Recall that $\mathcal{A}$ is the disjoint sum of $\mathcal{A}_1, \dots, \mathcal{A}_s$. Therefore, the r -neighbourhoods of all tuples $\bar{a}_i$ and $\bar{a}_j$, 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 $\bar{a}$ and each element a as chosen above, $\mathcal{A} \models \text{sph}_\tau[\bar{a}, a]$, and consequently $(\mathcal{A}, \bar{a})$ satisfies $\psi(\bar{x})$.

This completes the proof of Claim 1.

Claim 2. The disjoint decomposition $(\Delta_1, \dots, \Delta_s, \beta)$ can be computed in time $O(s + \|\psi\|)$.

Proof of Claim 2: Let $\psi(\bar{x}) := \exists^{\geq k} x_{n+1} \text{ sph}_\tau(\bar{x}, x_{n+1})$ be a Hanf-formula of signature σ_s , where $k \geq 1$, $n \geq 0$, $r \geq 0$, and $\tau = (\mathcal{T}, \bar{c})$ is an r -sphere with $n+1$ centres. On input of $\psi(\bar{x})$ 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 $\delta_i(\bar{x}_i)$.

For Step (1), note that the structure $\mathcal{T}$ can easily be reconstructed from the formula sph_τ in time $O(\|\text{sph}_\tau\|)$. 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 $P_1^T, \dots, P_s^T$. 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 $P_1^T, \dots, P_s^T$ occur in a tuple of these relations. Both passes take time $O(\|\mathcal{T}\|) \subseteq O(\|\text{sph}_\tau\|)$.

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

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

Altogether, the algorithm takes time $O(s + \|\psi\|)$. This completes the proof of Claim 2 and Lemma 4.3. $\square$

By using Theorem 4.2 and Lemma 4.3, we can now prove Theorem 4.1.

Proof of Theorem 4.1. Given an input FO(σ_s)-formula $\varphi(\bar{x})$ of quantifier rank $q \geq 0$ and with $n \geq 0$ free variables, our algorithm proceeds as follows:

- (1) We employ Theorem 4.2 to transform $\varphi(\bar{x})$ into a $\mathfrak{C}_\nu$ -equivalent formula $\varphi^H(\bar{x})$ in Hanf normal form.

This takes time $2^{\text{poly}(\|\varphi\| \cdot \nu(4^q))}$. The Hanf-formulas occurring in this formula are of the form $\exists^{\geq k} x_{n+1} \text{ sph}_\tau(\bar{x}, x_{n+1})$, where $k \leq \|\varphi\| \cdot (q+1) \cdot \nu(4^q)$, and τ is a ν -bounded r -sphere of radius $r \leq 4^q$ with at most $n+1$ centres. Since $q+2 \leq \|\varphi\|$, we have

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

- (2) For a suitable $L \geq 1$, let $\psi_1(\bar{x}), \dots, \psi_L(\bar{x})$ be the Hanf-formulas occurring in $\varphi^H(\bar{x})$.

For each $\ell \in [1, L]$, we employ Lemma 4.3 to compute a disjoint decomposition $(\Delta_{1,\ell}, \dots, \Delta_{s,\ell}, \beta_\ell)$ for the the Hanf-formula $\psi_\ell(\bar{x})$. 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 $\exists^{\geq k} x_{n+1} \text{ sph}_\tau(\bar{x}, x_{n+1})$ of $\varphi^H(\bar{x})$, this takes time

$$O(s + (k^2 + \|\text{sph}_\tau\|)) = O(s) + \text{poly}(\|\varphi\| \cdot \nu(4^q)).$$

Since $L < \|\varphi^H\|$ and $\|\varphi^H\| \in 2^{\text{poly}(\|\varphi\| \cdot \nu(4^q))}$, this step takes time

$$L \cdot (O(s) + \text{poly}(\|\varphi\| \cdot \nu(4^q))) \subseteq s \cdot 2^{\text{poly}(\|\varphi\| \cdot \nu(4^q))}.$$

- (3) We output the disjoint decomposition $D := (\Delta_1, \dots, \Delta_s, \beta)$, where $\Delta_i := \bigcup_{\ell=1}^L \Delta_{i,\ell}$, for each $i \in [1, s]$, and the propositional formula β is obtained from $\varphi^H(\bar{x})$ by replacing each Hanf-formula $\psi_\ell(\bar{x})$, for every $\ell \in [1, L]$, by the propositional formula β_ℓ .

Again, this takes time $s \cdot 2^{\text{poly}(\|\varphi\| \cdot \nu(4^q))}$,

which also bounds the total running time of the algorithm. It is straightforward to verify that D is indeed a disjoint decomposition for $\varphi(\bar{x})$ on $\mathfrak{C}_\nu$, that is, for all structures $\mathcal{A}_1, \dots, \mathcal{A}_s \in \mathfrak{C}_\nu$, their disjoint sum $\mathcal{A} = \mathcal{A}_1 \oplus \dots \oplus \mathcal{A}_s$ with mapping π , and all $\bar{a} \in A^n$, $(\mathcal{A}, \bar{a}) \models \varphi(\bar{x})$ iff $(\mathcal{A}_1, \dots, \mathcal{A}_s, \pi(\bar{a})) \models D$. This completes the proof of Theorem 4.1. $\square$

4.3 A matching lower bound

The following two lower bounds show that, for structures of degree 3 or 2, the algorithm of Theorem 4.1 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 4.4. *There is an $\epsilon > 0$ and a sequence of $\text{FO}(E)$ -sentences $(\varphi_h)_{h \geq 1}$ of increasing size such that, for every $h \geq 1$, every 2-disjoint decomposition for φ_h on finite binary forests has size at least $2^{2^{\epsilon \cdot \|\varphi_h\|}}$.*

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 4.5. *There is an $\epsilon > 0$ and a sequence of $\text{FO}(\sigma')$ -sentences $(\varphi'_h)_{h \geq 1}$ of increasing size such that, for every $h \geq 1$, 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 23 in (the full version of) [9]. The proofs, which will be included in the full version of this paper, employ encodings of numbers by binary trees, introduced in [10], and by labeled chains.

4.4 Generalised decompositions

By using the well-known notion of *transductions* (also known as *first-order interpretations*, cf. [8]), it is not difficult to transfer the algorithm from Theorem 4.1 to an algorithm producing more general decompositions than just *disjoint* decompositions. This way, we obtain the following result on *cartesian products* (also known as *tensor products*) rather than disjoint sums.

Let $s \geq 1$. For σ -structures $\mathcal{A}_1, \dots, \mathcal{A}_s$, the *cartesian product* $\mathcal{A}_1 \otimes \dots \otimes \mathcal{A}_s$ is the σ -structure $\mathcal{A} = (A, (R^{\mathcal{A}})_{R \in \sigma})$, where the universe A is the set $A_1 \times \dots \times A_s$, and for each relation symbol $R \in \sigma$ of arity $r := \text{ar}(R)$, $R^{\mathcal{A}}$ is the set of all tuples $((a_{1,1}, \dots, a_{1,s}), \dots, (a_{n,1}, \dots, a_{n,s}))$ in A^r such that, for each $i \in [1, s]$, $(a_{1,i}, \dots, a_{n,i}) \in R^{\mathcal{A}_i}$.

Corollary 4.6. *Let $s \geq 1$. There is an algorithm which, given an input $\text{FO}(\sigma)$ -formula $\varphi(\bar{x})$ of quantifier rank $q \geq 0$ and with $n \geq 0$*

free variables, constructs in time

$$2^{\text{poly}(s \cdot \|\varphi\| \cdot \nu(4^{s \cdot q}))} \quad (8)$$

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

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

Furthermore, the sets $\Delta_1, \dots, \Delta_s$ consist of Hanf-formulas over σ , and the degree of the polynomial in the expression (8) is linear in $\|\sigma\| + s$.

A proof of Corollary 4.6 will be included in the full version of this paper.

Acknowledgments

We would like to thank the anonymous referees for their valuable comments. Furthermore, the second author would like to thank Dietrich Kuske for inspiring discussions on Feferman-Vaught decompositions.

References

- [1] M. Ajtai and Y. Gurevich. Datalog vs. first-order logic. In *Proc. FOCS'89*, pages 142–147, 1989.
- [2] A. Atserias, A. Dawar, and P. G. Kolaitis. On preservation under homomorphisms and unions of conjunctive queries. *J. ACM*, 53(2): 208–237, 2006.
- [3] A. Atserias, A. Dawar, and M. Grohe. Preservation under extensions on well-behaved finite structures. *SIAM J. Comput.*, 38(4):1364–1381, 2008.
- [4] B. Bollig and D. Kuske. An optimal construction of Hanf sentences. *J. Applied Logic*, 10(2):179–186, 2012.
- [5] A. K. Chandra and P. M. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *Proc. STOC'77*, pages 77–90, 1977.
- [6] A. Dawar. Homomorphism preservation on quasi-wide classes. *J. Comput. Syst. Sci.*, 76(5):324–332, 2010.
- [7] A. Dawar, M. Grohe, S. Kreutzer, and N. Schweikardt. Model Theory Makes Formulas Large. In *Proc. ICALP'07*, pages 1076–1088, 2007. Full version available as preprint NI07003-LAA, Isaac Newton Institute of Mathematical Sciences (2007).
- [8] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 1999.
- [9] S. Göller, J. C. Jung, and M. Lohrey. The complexity of decomposing modal and first-order theories. In *Proc. LICS'12*, pages 325–334, 2012. Full version available at <http://www.eti.uni-siegen.de/ti/veroeffentlichungen/12-modal.pdf>.
- [10] L. Heimberg, D. Kuske, and N. Schweikardt. An optimal Gaifman normal form construction for structures of bounded degree. In *Proc. LICS'13*, pages 63–72, 2013. Full version available at <http://www.tks.cs.uni-frankfurt.de/research/HKS13-Gaifman-FullVersion.pdf>.
- [11] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [12] J. Makowsky. Algorithmic uses of the Feferman-Vaught theorem. *Ann. Pure Appl. Logic*, 126(1-3):159–213, 2004. ISSN 0168-0072.
- [13] J. Nurmonen. Counting modulo quantifiers on finite structures. *Information and Computation*, 160(1-2):62–87, 2000.
- [14] B. Rossman. Homomorphism preservation theorems. *J. ACM*, 55(3), 2008.