Model theory makes formulas large

Anuj Dawar[1], Martin Grohe[2], Stephan Kreutzer[2], and Nicole Schweikardt[2]

[1] University of Cambridge, U.K., anuj.dawar@cl.cam.ac.uk
[2] Institut für Informatik, Humboldt Universität zu Berlin,
{grohe,kreutzer,schweika}@informatik.hu-berlin.de

Abstract. Gaifman's locality theorem states that every first-order sentence is equivalent to a local sentence. We show that there is no elementary bound on the length of the local sentence in terms of the original.

The classical Łoś-Tarski theorem states that every first-order sentence preserved under extensions is equivalent to an existential sentence. We show that there is no elementary bound on the length of the existential sentence in terms of the original. Recently, variants of the Łoś-Tarski theorem have been proved for certain classes of finite structures, among them the class of finite acyclic structures and more generally classes of structures of bounded tree width. Our lower bound also applies to these variants.

We further prove that a version of the Feferman-Vaught theorem based on a restriction by formula length necessarily entails a non-elementary blow-up in formula size.

All these results are based on a similar technique of encoding large numbers by trees of small height in such a way that small formulas can speak about these numbers. Notably, our lower bounds do not apply to restrictions of the results to structures of bounded degree. For such structures, we obtain elementary upper bounds in all cases. However, even there we can prove at least doubly exponential lower bounds.

1 Introduction

Classical results of model theory provide syntactical normal forms for various semantical properties of structures. For example, the Łoś-Tarski theorem states that every first-order definable property that is preserved under extensions of structures is actually definable by an existential first-order sentence. Gaifman's locality theorem provides a normal form for all properties definable in first-order logic. It states that each first-order definable property is definable by a local sentence, that is, a sentence where quantification is basically restricted to local neighbourhoods of elements.

Gaifman's theorem has found various applications in algorithms and complexity [9, 4, 16, 17]. In particular, there are very general algorithmic meta-theorems stating that first-order model-checking is fixed-parameter tractable for various classes of structures, such as planar graphs or graphs with excluded minors, and that first-order definable optimisation problems on such classes have polynomial time approximation schemes. These algorithms are heavily based on (an effective version of) Gaifman's theorem: First-order formulas are first translated into local formulas, and then these local formulas are algorithmically evaluated.

While it is known that the Łoś-Tarski theorem fails when restricted to all finite structures, it has recently been proved [1] that the theorem does still hold when restricted to specific "well-behaved" classes of finite structures such as acyclic structures, structures of bounded tree width, and structures of bounded degree.

In the context of algorithms, complexity, and finite model theory, questions about the efficiency of the normal forms, which are usually neglected in classical model theory, are of fundamental importance. These are the questions we address. By efficiency we mean the size of the formulas in normal form (*succinctness*) and the existence of efficient algorithms that translate formulas into their normal forms (*complexity of the translation*). We shall prove nonelementary lower bounds for the succinctness — obviously, this implies nonelementary lower bounds on the complexity of the translation. Specifically, we prove that there is no elementary function f such that every first-order sentence φ is equivalent to a local first-order sentence $\widetilde{\varphi}$ of length $||\widetilde{\varphi}|| \leq f(||\varphi||)$, not even on the class of all finite trees. Similarly, we prove that there is no elementary function f such that every first-order sentence φ that is preserved under extensions (on arbitrary structures) is equivalent to an existential first-order sentence $\widetilde{\varphi}$ of length $||\widetilde{\varphi}|| \leq f(||\varphi||)$, not even on the class of all finite trees. This provides a succinctness lower bound for both the classical Łoś-Tarski theorem and its variants for classes of finite forests and all classes of finite structures that contain all trees (but not for classes of finite structures of bounded degree).

We prove a further lower bound that is concerned with the classical Feferman-Vaught theorem. The classical 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. In particular, there is a function f such that if structures $\mathcal{A}_i$ and $\mathcal{B}_i$ (for $i = 1, 2$) satisfy the same first-order sentences of length at most $f(\ell)$, then the disjoint union of $\mathcal{A}_1$ and $\mathcal{A}_2$ satisfies the same first-order sentences of size ℓ as the disjoint union of $\mathcal{B}_1$ and $\mathcal{B}_2$. We prove a lower bound on the growth rate of f showing that f is not bounded above by an elementary function.

Technically, all our lower bound proofs rely on a suitable encoding of large natural numbers by trees of small height that can be controlled by small first-order formulas. In fact, we show — and use — that full arithmetic on a large initial segment of the positive integers can be simulated by comparably small first-order formulas that operate on the tree encodings of the numbers. It is worth mentioning that this approach can be applied in various other contexts. For example, concerning the classical decision problem, it is known that the first-order theory (and actually also the monadic second-order theory) of trees is decidable [23, 19]; and in [2] (see also [7] for related results) it has been shown that there is no *elementary* decision algorithm. A simple proof of this non-elementary lower bound can easily be obtained using the methods in the present paper (details of this can be found in the full version of this paper).

A point to note, however, is that all our non-elementary lower bounds heavily rely on the fact that the degree of the underlying structures is unbounded. In fact, when restricting attention to classes of structures of bounded degree, we can show elementary upper bounds as counterparts of the non-elementary lower bounds on classes of structures of unbounded degree. In particular, in the bounded degree case we obtain a 4-fold exponential upper bound for Gaifman's locality theorem; and we get a 5-fold exponential upper bound for the variant of the Łoś-Tarski theorem on the class of acyclic structures of bounded degree.

As far as we know, techniques similar to those applied here go back to Stockmeyer and Meyer [21]. Much later, such techniques have been employed in [10, 18, 12, 13] to prove lower bounds in parameterised complexity, respectively, on the succinctness of monadic logics. A related succinctness lower bound deserves mention. It has recently been proved by Rossman [20] that the homomorphism preservation theorem (in contrast with the Łoś-Tarski theorem) holds in the class of all finite structures. Here, it is known that there is no elementary bound on the length of the existential positive formula obtained.[3]

The rest of the paper is structured as follows. Section 2 establishes some definitions and notation and Section 3 presents the encoding of numbers by trees that is then used to prove lower bounds on the size of formulas in Gaifman normal form (Section 4) and the failure of the Feferman-Vaught theorem for formula length (Section 5). Section 6 then establishes the lower bound for the Łoś-Tarski theorem, which is based on a different encoding of numbers by trees. Finally, Section 7 contains the elementary upper bounds on classes of structures of bounded degree. Due to space limitations, many technical details of the proofs are deferred to the full version of this paper.

Acknowledgements: We would like to thank an anonymous referee for pointing us to the references [2, 7]. We gratefully acknowledge the support of the Isaac Newton Institute through its 2006 programme on Logic and Algorithms. The opportunity afforded by this programme greatly aided our collaboration.

2 Preliminaries

We use $\mathbb{R}$ to denote the set of reals and $\mathbb{N}$ to denote the set of natural numbers, i.e., the set of nonnegative integers. For natural numbers $m < n$ we write $[m, n]$ to denote the set $\{m, m{+}1, \ldots, n\}$.

We say that a function $f : \mathbb{N} \to \mathbb{R}$ is *(1-fold) exponential* if there is some polynomial p such that $f(n)$ is eventually bounded by $2^{p(n)}$. For any $k \geq 2$, a function f is called *k-fold exponential* if there is some $(k{-}1)$-fold exponential function g such that $f(n)$ is eventually bounded by $2^{g(n)}$. A function $f : \mathbb{N} \to \mathbb{R}$ is called *elementary* if it can be formed from the successor function, addition, subtraction, and multiplication using compositions, projections, bounded additions, and bounded multiplications (of the form $\sum_{z \leq y} g(\overline{x}, z)$ and $\Pi_{z \leq y} g(\overline{x}, z)$). The crucial fact for us is that a function f is bounded by an elementary function if, and only if, there exists a $k \geq 1$ such that f is bounded by a k-fold exponential function (see, e.g., [3]).

One function of particular interest for the present paper is the function *Tower* : $\mathbb{N} \to \mathbb{N}$, defined via $\mathit{Tower}(0) := 1$ and, for all $h \geq 1$, $\mathit{Tower}(h) := 2^{\mathit{Tower}(h-1)}$. I.e., $\mathit{Tower}(h)$ is a tower of 2s of height h. Note that, e.g., none of the functions $\mathit{Tower}(h)$, $\mathit{Tower}(\sqrt[4]{h})$, $\mathit{Tower}(\log h)$ is bounded by an elementary function.

A *vocabulary* is a finite set of relation symbols and constant symbols. Associated with every relation symbol R is a positive integer called the *arity* of R. In the following,

[3] This is mentioned by Rossman citing unpublished work of Gurevich and Shelah. As far as we are aware, a proof of this lower bound has not yet been published.

τ always denotes a vocabulary. A vocabulary τ is called *relational* if it does not contain any constant symbol. A τ-*structure* $\mathcal{A}$ consists of a non-empty set A, called the *universe* of $\mathcal{A}$, an element $c^{\mathcal{A}} \in A$ for each constant symbol $c \in \tau$, and a relation $R^{\mathcal{A}} \subseteq A^r$ for each r-ary relation symbol $R \in \tau$. $\mathcal{A}$ is called an *induced substructure* of a τ-structure $\mathcal{B}$ if $A \subseteq B$, $R^{\mathcal{A}} = R^{\mathcal{B}} \cap A^r$, for each relation symbol $R \in \tau$ of arity r, and $c^{\mathcal{A}} = c^{\mathcal{B}}$ for each constant symbol $c \in \tau$.

The *Gaifman graph* of a τ-structure $\mathcal{A}$ is the (undirected, loop-free) graph $\mathcal{G}_{\mathcal{A}}$ with vertex set A and an edge between two vertices $a, b \in A$ iff there exists an $R \in \tau$ and a tuple $(a_1, .., a_r) \in R^{\mathcal{A}}$ such that $a, b \in \{a_1, .., a_r\}$. The *distance* between two elements $a, b \in A$ in $\mathcal{A}$, denoted by $dist^{\mathcal{A}}(a, b)$, is defined to be the length (that is, number of edges) of the shortest path from a to b in the Gaifman graph of $\mathcal{A}$. For $r \geq 0$ and $a \in A$, the r-neighbourhood of a in $\mathcal{A}$ is the set $N_r^{\mathcal{A}}(a) = \{b \in A : dist^{\mathcal{A}}(a, b) \leq r\}$. The induced substructure of $\mathcal{A}$ with universe $N_r^{\mathcal{A}}(a)$ is denoted by $\mathcal{N}_r^{\mathcal{A}}(a)$. We omit superscripts $^{\mathcal{A}}$ if $\mathcal{A}$ is clear from the context.

We write FO(τ) to denote the class of all formulae of first-order logic over the vocabulary τ, and we write $qr(\varphi)$ to denote the *quantifier rank* of an FO(τ)-formula φ. In a natural way, we view formulas as trees (to be precise, as their *syntax trees*), where leaves correspond to the atoms of the formulas, and inner vertices correspond to Boolean connectives or quantifiers. We define the *size* (or, *length*) $||\varphi||$ of a first-order formula φ as the number of vertices in the syntax tree of φ.

Whenever we write E, it denotes a binary relation symbol. We view $\{E\}$-structures as directed graphs. For a directed graph $\mathcal{A} = (A, E^{\mathcal{A}})$ and an $a \in A$, we let A_a be the set of all vertices b such there is a path from a to b (this includes a), and we let $\mathcal{A}_a$ be the induced substructure of $\mathcal{A}$ with universe A_a. Unless we explicitly call them *undirected*, we view trees as being directed from the root to the leaves. A *forest* is a directed graph in which every vertex has indegree at most 1. Vertices of indegree 0 are called *roots* of the forest. A *tree* is a forest with exactly one root. The class of all finite forests is denoted by $\mathfrak{F}$ and the class of all finite trees by $\mathfrak{T}$. The *height* of a tree $\mathcal{T}$ is the length of the longest path in $\mathcal{T}$.

3 Encoding numbers by trees

In this section we introduce the technical machinery that is used for proving our main theorems in sections 4 and 5. We use the following encoding of natural numbers by trees, introduced in [8].

Definition 3.1 (Encoding numbers by trees). For $i, n \in \mathbb{N}$ we write bit(i, n) to denote the *i-th bit in the binary representation of* n. I.e., bit$(i, n) = 0$ if $\left\lfloor \frac{n}{2^i} \right\rfloor$ is even, and bit$(i, n) = 1$ if $\left\lfloor \frac{n}{2^i} \right\rfloor$ is odd. Inductively we define a tree $\mathcal{T}(n)$ for each natural number n as follows: $\mathcal{T}(0)$ is the one-node tree. For $n \geq 1$ the tree $\mathcal{T}(n)$ is obtained by creating a new root and attaching to it all trees $\mathcal{T}(i)$ for all i such that bit$(i, n) = 1$.

Illustrations of these trees can be found in [8]. It is straightforward to see (cf. [8, Lemma 10.20]) that for all $h, n \geq 0$, $height(\mathcal{T}(n)) \leq h \iff n < Tower(h)$. The next lemma from [8] shows that the tree encodings of numbers can be "controlled" by small first-order formulas. (In [8], the statement of the lemma is formulated for trees instead of general structures. The proof given there, however, also holds for general structures and thus leads to the following lemma.)

Lemma 3.2 ([8, Lemma 10.21]). *For every $h \geq 0$ there is an $\mathrm{FO}(E)$-formula $eq_h(x,y)$ of length $\mathcal{O}(h)$ such that for all structures $\mathcal{A} = (A, E^{\mathcal{A}})$ and $t, u \in A$ we have: If there are $m, n < Tower(h)$ such that the structures $\mathcal{A}_t$ and $\mathcal{A}_u$ are isomorphic to $\mathcal{T}(m)$ and $\mathcal{T}(n)$, resp., then* $\quad \mathcal{A} \models eq_h(t,u) \iff m = n$.

Using this, one easily obtains the following two lemmas which provide formulas of length polynomial in h that recognise tree encodings and define arithmetic on the tree encodings of numbers of size up to $Tower(h)$.

Lemma 3.3. *For every $h \geq 0$ there is a $\mathrm{FO}(E)$-formula $encoding_h(x)$ of length $\mathcal{O}(h^2)$ such that for all structures $\mathcal{A} = (T, E^{\mathcal{A}})$ and $t \in A$ we have: $\mathcal{A} \models encoding_h(t) \iff$ there is an $i \in \{0, .., Tower(h)-1\}$ such that $\mathcal{A}_t$ is isomorphic to $\mathcal{T}(i)$.*

Lemma 3.4. *For every $h \geq 0$ there are $\mathrm{FO}(E)$-formulas $bit_h(x,y)$ of size $\mathcal{O}(h)$, $less_h(x,y)$ of size $\mathcal{O}(h^2)$, $min(x)$ of constant size (not depending on h), $succ_h(x,y)$ of size $\mathcal{O}(h^3)$, and $max_h(x)$ of size $\mathcal{O}(h^4)$ such that for all structures $\mathcal{A} = (A, E^{\mathcal{A}})$ and $t, u \in A$ we have: If there are $m, n < Tower(h)$ such that the structures $\mathcal{A}_t$ and $\mathcal{A}_u$ are isomorphic to $\mathcal{T}(m)$ and $\mathcal{T}(n)$, respectively, then we have:*

(a)	$\mathcal{A} \models bit_h(t,u)$	$\iff$	$\mathrm{bit}(m,n) = 1$
(b)	$\mathcal{A} \models less_h(t,u)$	$\iff$	$m < n$
(c)	$\mathcal{A} \models min(t)$	$\iff$	*$\mathcal{A}_t$ is isomorphic to $\mathcal{T}(0)$*
(d)	$\mathcal{A} \models succ_h(t,u)$	$\iff$	$m + 1 = n$
(e)	$\mathcal{A} \models max_h(t)$	$\iff$	*$\mathcal{A}_t$ is isomorphic to $\mathcal{T}(Tower(h)-1)$.*

4 Lower bounds for the size of formulas in Gaifman normal form

The aim of this section is to prove a non-elementary succinctness gap for Gaifman's theorem. To give a precise formulation of Gaifman's theorem and our new bounds on formula length, we need to fix some (standard) notation.

For every $r \geq 0$, we let $dist_{\leq r}(x,y)$ be an $\mathrm{FO}(\tau)$-formula expressing that the distance between x and y is at most r. We often write $dist(x,y) \leq r$ instead of $dist_{\leq r}(x,y)$ and $dist(x,y) > r$ or $dist_{>r}(x,y)$ instead of $\neg dist_{\leq r}(x,y)$. An $\mathrm{FO}(\tau)$-formula $\psi(x)$ is called *r-local* if for every τ-structure $\mathcal{A}$ and every $a \in A$ we have $\mathcal{A} \models \psi(a) \iff \mathcal{N}_r^{\mathcal{A}}(a) \models \psi(a)$. A *basic local sentence* (with parameters k, r) is a sentence of the form

$$\exists x_1 \cdots \exists x_k \Big(\bigwedge_{1 \leq i < j \leq k} dist(x_i, x_j) > 2r \ \wedge \bigwedge_{1 \leq i \leq k} \psi(x_i) \Big),$$

where $\psi(x)$ is r-local.

For an $\mathrm{FO}(\tau)$-sentence φ we say that *φ is in Gaifman normal form* if φ is a Boolean combination of basic local sentences. Gaifman's well-known theorem [11] states that every first-order sentence over a relational vocabulary is equivalent to a first-order sentence in Gaifman normal form. The proof in [11] proceeds by induction on the length of the given first-order sentence φ and leads to an effective algorithm that transforms a given φ into an equivalent sentence ψ in Gaifman normal form. A closer look at Gaifman's proof shows that the size of the constructed sentence ψ may be non-elementary in the size of the original sentence φ. The main result of the present section shows that this huge increase in formula size is not just an artifact of Gaifman's proof, but that indeed there are first-order formulas φ for which the shortest equivalent formula in Gaifman normal form is non-elementarily larger than φ.

Theorem 4.1. *For every $h \geq 1$ there is an* FO(E)*-sentence φ_h of size $\mathcal{O}(h^4)$ such that every* FO(E)*-sentence in Gaifman normal form that is equivalent to φ_h on the class $\mathfrak{T}$ of finite trees has size at least Tower(h).*

Here, we show the following variant that speaks about the class of all *forests* rather than *trees*. The proof of Theorem 4.2 avoids some of the unpleasant details needed in the proof of Theorem 4.1 while still exposing the main ideas that are crucial for the proof of Theorem 4.1. The proof of Theorem 4.1 can be found in the full version of this paper.

Theorem 4.2. *For every $h \geq 1$ there is an* FO(E)*-sentence φ_h of size $\mathcal{O}(h^4)$ such that every* FO(E)*-sentence in Gaifman normal form that is equivalent to φ_h on the class $\mathfrak{F}_{\leq h}$ of finite forests of height $\leq h$ has size at least Tower(h).*

Proof. We use the tree encodings of natural numbers introduced in Section 3. For $h \geq 1$ we define the structure $\mathcal{F}_h$ to be the forest that consists of the disjoint union of all trees $\mathcal{T}(j)$ for all $j \in \{0, .., \mathit{Tower}(h)-1\}$. Furthermore, for every $i \in \{0, .., \mathit{Tower}(h)-1\}$, we let $\mathcal{F}_h^{-i}$ be the forest that consists of the disjoint union of all trees $\mathcal{T}(j)$ for all j with $j \neq i$ and $j \in \{0, .., \mathit{Tower}(h)-1\}$.

We let *root*(x) be a formula which expresses that a node x has in-degree 0, i.e., $\mathit{root}(x) := \neg \exists y\, E(y,x)$. We choose the FO$(E)$-sentence φ_h as follows: $\varphi_h :=$

$$\exists x \, \big(\mathit{root}(x) \wedge \mathit{min}(x)\big) \ \wedge \ \forall y \Big(\big(\mathit{root}(y) \wedge \neg \mathit{max}_h(y)\big) \rightarrow \exists z \big(\mathit{root}(z) \wedge \mathit{succ}_h(y,z)\big) \Big) \, .$$

Using Lemma 3.4, it is straightforward to see that $||\varphi_h|| = \mathcal{O}(h^4)$ and

$$\mathcal{F}_h \models \varphi_h \qquad \text{and,} \quad \text{for each } i < \mathit{Tower}(h), \quad \mathcal{F}_h^{-i} \not\models \varphi_h \, . \tag{1}$$

Now let ψ be an FO(E)-sentence in *Gaifman normal form* that is equivalent to φ_h on the class $\mathfrak{F}_{\leq h}$. In particular, since $\mathcal{F}_h$ as well as all the $\mathcal{F}_h^{-i}$ belong to $\mathfrak{F}_{\leq h}$, we obtain from (1) that

$$\mathcal{F}_h \models \psi \qquad \text{and,} \quad \text{for each } i < \mathit{Tower}(h), \quad \mathcal{F}_h^{-i} \not\models \psi \, . \tag{2}$$

Our aim is to show that $H := ||\psi|| \geq \mathit{Tower}(h)$. Aiming at a contradiction, let us now assume that $H < \mathit{Tower}(h)$.

Since ψ is in Gaifman normal form, it is a Boolean combination of *basic local sentences* $\chi_1, .., \chi_L$, where each χ_ℓ (for $\ell \in \{1, .., L\}$) is of the form

$$\chi_\ell \ := \ \exists x_1 \cdots \exists x_{k_\ell} \Big(\bigwedge_{1 \leq i < j \leq k_\ell} \mathit{dist}(x_i, x_j) > 2r_\ell \ \wedge \bigwedge_{1 \leq i \leq k_\ell} \psi_\ell(x_i) \Big),$$

with $k_\ell, r_\ell \geq 1$ and $\psi_\ell(x)$ a formula that is r_ℓ-local. In particular,

$$H \;:=\; ||\psi|| \;\geq\; k_1 + \cdots + k_L\,. \tag{3}$$

We can assume w.l.o.g. that there exists an $\tilde{L}$ with $0 \leq \tilde{L} \leq L$ such that

$$\text{for each } \ell \leq \tilde{L}, \quad \mathcal{F}_h \models \chi_\ell\,, \qquad \text{and} \qquad \text{for each } \ell > \tilde{L}, \quad \mathcal{F}_h \not\models \chi_\ell\,. \tag{4}$$

For all $\ell \leq \tilde{L}$ we know that $\mathcal{F}_h \models \chi_\ell$, i.e., there are nodes $t_1^{(\ell)}, \ldots, t_{k_\ell}^{(\ell)}$ in $\mathcal{F}_h$ such that the formula

$$\bigwedge_{1 \leq i < j \leq k_\ell} \mathit{dist}(x_i, x_j) > 2r_\ell \;\wedge\; \bigwedge_{1 \leq i \leq k_\ell} \psi_\ell(x_i) \tag{5}$$

is satisfied in $\mathcal{F}_h$ when interpreting each variable x_i with the node $t_i^{(\ell)}$. The set $\{\, t_i^{(\ell)} : \ell \leq \tilde{L} \text{ and } i \leq k_\ell \,\}$ consists of at most $k_1 + \cdots + k_{\tilde{L}} \leq H$ nodes (see (3)). Since we assume that $H < \mathit{Tower}(h)$, and since $\mathcal{F}_h$ consists of $\mathit{Tower}(h)$ disjoint trees, there must be at least one component $\mathcal{T}$ of $\mathcal{F}_h$ in which none of the nodes from $\{\, t_i^{(\ell)} : \ell \leq \tilde{L} \text{ and } i \leq k_\ell \,\}$ is present. Let $j \in \{0, .., \mathit{Tower}(h){-}1\}$ be such that $\mathcal{T} = \mathcal{T}(j)$.

Now, of course, the forest $\mathcal{F}_h^{-j}$, which is obtained from $\mathcal{F}_h$ by removing the component $\mathcal{T}(j)$, still contains all the nodes in $\{\, t_i^{(\ell)} : \ell \leq \tilde{L} \text{ and } i \leq k_\ell \,\}$. Considering (5), note that each formula $\psi_\ell(x_i)$ is r_ℓ-local around x_i. Thus, when interpreting x_i with the node $t_i^{(\ell)}$, the formula can only "speak" about the r_ℓ-neighbourhood of $t_i^{(\ell)}$, which is the same in $\mathcal{F}_h^{-j}$ as in $\mathcal{F}_h$. We thus obtain from (5) that $\mathcal{F}_h^{-j} \models \chi_\ell$ for each $\ell \leq \tilde{L}$.

Let us now consider the formulas χ_ℓ with $\ell > \tilde{L}$. From (4) we know that $\mathcal{F}_h \not\models \chi_\ell$, i.e., $\mathcal{F}_h \models \neg\chi_\ell$, where the formula $\neg\chi_\ell$ is of the following form:

$$\neg\exists x_1 \cdots \exists x_{k_\ell} \Big(\bigwedge_{1 \leq i < j \leq k_\ell} \mathit{dist}(x_i, x_j) > 2r_\ell \;\wedge\; \bigwedge_{1 \leq i \leq k_\ell} \psi_\ell(x_i) \Big)\,.$$

Since the formula $\psi_\ell(x_i)$ is r_ℓ-local and since $\mathcal{F}_h^{-j}$ is obtained from $\mathcal{F}_h$ by removing an entire component of $\mathcal{F}_h$, it is straightforward to see that also $\mathcal{F}_h^{-j} \models \neg\chi_\ell$. In total, we now know the following:

$$\text{for each } \ell \leq \tilde{L}, \quad \mathcal{F}_h^{-j} \models \chi_\ell\,, \qquad \text{and} \qquad \text{for each } \ell > \tilde{L}, \quad \mathcal{F}_h^{-j} \not\models \chi_\ell\,. \tag{6}$$

From (6) and (4) we obtain that $\mathcal{F}_h^{-j}$ satisfies exactly the same basic local sentences from $\{\chi_1, \ldots, \chi_L\}$ as $\mathcal{F}_h$. Since ψ is a Boolean combination of the sentences $\chi_1, .., \chi_L$, we thus have that $\mathcal{F}_h^{-j} \models \psi \iff \mathcal{F}_h \models \psi$. This, however, is a contradiction to (2). Altogether, the proof of Theorem 4.2 is complete. $\square$

To conclude this section let us mention that an easy reduction shows that Theorem 4.1 and Theorem 4.2 still hold when replacing $\mathfrak{T}$ and $\mathfrak{F}_{\leq h}$ by the classes $\mathfrak{T}^u$ and $\mathfrak{F}_{\leq h}^u$ of *undirected* trees and *undirected* forests of height at most h, respectively.

5 Failure of Feferman-Vaught theorems for formula size

The classical Feferman-Vaught theorem [6] 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. The plainest form of composition is the *disjoint union*, denoted by $\oplus$ in the following. The Feferman-Vaught theorem for disjoint union and first-order logic states that for all structures $\mathcal{A}_1, \mathcal{A}_2, \mathcal{B}_1, \mathcal{B}_2$, if the structures $\mathcal{A}_i$ and $\mathcal{B}_i$ (for $i = 1, 2$) satisfy the same first-order sentences, their disjoint unions $\mathcal{A}_1 \oplus \mathcal{A}_2$ and $\mathcal{B}_1 \oplus \mathcal{B}_2$ also satisfy the same first-order sentences. This can be stratified by the quantifier rank, that is, if $\mathcal{A}_i$ and $\mathcal{B}_i$ satisfy the same first-order sentences of quantifier rank at most q, then $\mathcal{A}_1 \oplus \mathcal{A}_2$ and $\mathcal{B}_1 \oplus \mathcal{B}_2$ also satisfy the same first-order sentences of quantifier rank at most q. This result is an immensely useful tool in analysing the expressivity of first order logic, and for deriving bounds on the quantifier rank.

To derive bounds on the formula size, it would be similarly useful to have an analogous result for formula size instead of quantifier rank. As (for a fixed, finite vocabulary) there are only finitely many first-order sentences of quantifier rank q, up to logical equivalence, we immediately get the following: There is a function f such that if the structures $\mathcal{A}_i$ and $\mathcal{B}_i$ (for $i = 1, 2$) satisfy the same first-order sentences of length at most $f(\ell)$, then $\mathcal{A}_1 \oplus \mathcal{A}_2$ and $\mathcal{B}_1 \oplus \mathcal{B}_2$ satisfy the same first-order sentences of length at most ℓ. It is not hard to derive an upper bound of $\mathit{Tower}(\mathcal{O}(\ell))$ on the function f. Maybe surprisingly, this upper bound is essentially tight:

Theorem 5.1. *There is no elementary function f such that the following holds for all trees $\mathcal{A}, \mathcal{B}, \mathcal{C} \in \mathfrak{T}$: If $\mathcal{A}$ and $\mathcal{B}$ satisfy the same first-order sentences of length at most $f(\ell)$, then $\mathcal{A} \oplus \mathcal{C}$ and $\mathcal{B} \oplus \mathcal{C}$ satisfy the same first-order sentences of length at most ℓ.*

Proof. We use the encoding and the formulas from Section 3. For every $h \geq 1$, let

$$\varphi_h \ := \ \forall x\big(\mathit{encoding}_h(x) \rightarrow \big(\mathit{max}_h(x) \vee \exists y\ \mathit{succ}_h(x, y)\big)\big).$$

Then there is a constant $c \geq 1$ such that $||\varphi_h|| \leq c \cdot h^4$ for all h.

Suppose for contradiction that f is an elementary function with the desired property. We may assume that $f(\ell) \geq \ell$ for all $\ell \geq 1$. As there are only exponentially many first-order sentences φ of a given length, there is an $h \geq 1$ such that there are less than $\mathit{Tower}(h{-}1)$ first-order sentences of length at most $f(c{\cdot}h^4)$ (up to equivalence). Let us fix such an h, and let $\ell = c{\cdot}h^4$ and $n = \mathit{Tower}(h){-}1$. For every $j \in [0, n]$, let $\mathcal{F}_j$ denote the forest consisting of the trees $\mathcal{T}(j), \ldots, \mathcal{T}(n)$, and let $\mathcal{U}_j$ be the tree obtained from $\mathcal{F}_j$ by connecting a new root with the roots of all trees in $\mathcal{F}_j$. Then there are numbers j, k such that $1 \leq j < k \leq n$, and the trees $\mathcal{U}_j$ and $\mathcal{U}_k$ satisfy the same first-order sentences of length at most $f(\ell)$. Observe that

$$\mathcal{F}_j \oplus \mathcal{T}(j{-}1) \models \varphi_h \quad \text{and} \quad \mathcal{F}_k \oplus \mathcal{T}(j{-}1) \not\models \varphi_h.$$

Now let $\mathcal{A} = \mathcal{U}_j$, $\mathcal{B} = \mathcal{U}_k$, and $\mathcal{C} = \mathcal{T}(j{-}1)$. As the new roots of $\mathcal{A}, \mathcal{B}$ are not nodes satisfying $\mathit{encoding}_h(x)$ (because $\mathcal{A}$ and $\mathcal{B}$ are isomorphic to trees $\mathcal{T}(n_\mathcal{A})$ and $\mathcal{T}(n_\mathcal{B})$ with $n_\mathcal{A}, n_\mathcal{B} \geq \mathit{Tower}(h)$), we have $\mathcal{A} \oplus \mathcal{C} \models \varphi_h$ and $\mathcal{B} \oplus \mathcal{C} \not\models \varphi_h$. Since the length of φ_h is at most ℓ and $\mathcal{A}, \mathcal{B}$ satisfy the same sentences of length at most $f(\ell)$, this is a contradiction. $\square$

6 Existential preservation on forests

A structure $\mathcal{B}$ is called an *extension* of $\mathcal{A}$ if $\mathcal{A}$ is an induced substructure of $\mathcal{B}$. Let τ be a vocabulary and let $\mathfrak{C}$ be a class of finite τ-structures that is closed under induced substructures. An FO(τ)-sentence φ is *preserved under extensions on* $\mathfrak{C}$ if the following is true for all structures $\mathcal{A}, \mathcal{B} \in \mathfrak{C}$: If $\mathcal{A} \models \varphi$ and $\mathcal{B}$ is an extension of $\mathcal{A}$, then $\mathcal{B} \models \varphi$.

The well-known *Łoś-Tarski Theorem* (see e.g. [15]) states that every first-order sentence that is preserved under extensions on the class of *all* structures (i.e., finite as well as infinite structures), is equivalent to an *existential* first-order sentence. Here, the class of *existential first-order formulas* is obtained by closing the atomic formulas and the negated atomic formulas under conjunction, disjunction, and existential quantification. While the *Łoś-Tarski theorem* fails when shifting the attention from the class of *all* structures to the class of all *finite* structures ([22, 14]), it was shown in [1] that the Łoś-Tarski theorem holds when restricted to certain "well-behaved" classes of finite structures, among them the class of all finite acyclic structures. The main result of the present section, Theorem 6.1, shows that a translation of formulas preserved under extensions into equivalent existential formulas may increase the formula size non-elementarily.

In the following, we let L and X be two unary relation symbols. An *$\{L, X\}$-labelled tree* is an $\{E, L, X\}$-structure $\mathcal{T} = (T, E^{\mathcal{T}}, L^{\mathcal{T}}, X^{\mathcal{T}})$ where $(T, E^{\mathcal{T}})$ is a tree.

Theorem 6.1. *Let τ be a vocabulary that consists of a binary relation symbol E and two unary relation symbols L and X. For every $h \geq 1$ there is a* FO(τ)*-sentence φ_h of size $2^{\mathcal{O}(h)}$ with the following properties:*

1. *φ_h is preserved under extensions on the class of all τ-structures, and*
2. *every existential* FO(τ)*-sentence ψ that is equivalent to φ_h on the class $\mathfrak{T}_{\leq h}$ of all $\{L, X\}$-labelled trees of height at most h is of size at least Tower$(h-1)$.*

Using the same approach as in the previous sections, i.e., the encoding of natural numbers by trees introduced in Section 3, it is not difficult to construct a sentence φ_h of small size which meets requirement 2. We were, however, unable to find a sentence based on this encoding which also meets requirement 1 (even when considering $\mathfrak{T}_{\leq h}$ instead of the class of all τ-structures). To prove Theorem 6.1, we therefore introduce the following encoding of numbers by $\{L, X\}$-labelled trees. The remainder of this section is devoted to the proof of Theorem 6.1.

From now on, until the end of this section, we let τ denote a vocabulary that consists of a binary relation symbol E and two unary relation symbols L and X.

Definition 6.2. For each natural number $h \geq 1$ and each $n \in \{0, 1, .., \textit{Tower}(h)-1\}$, we define the $\{L, X\}$-labelled tree $\tilde{\mathcal{T}}_h(n)$ as follows:

- $\tilde{\mathcal{T}}_1(0)$ consists of two nodes u and v such that there is an edge from u to v, and v is labelled to be a *leaf* (which is encoded by "$v \in L$") and v is labelled **0** (which is encoded by "$v \notin X$").
- $\tilde{\mathcal{T}}_1(1)$ consists of two nodes u and v such that there is an edge from u to v, and v is labelled to be a *leaf* (which is encoded by "$v \in L$") and v is labelled **1** (which is encoded by "$v \in X$").

– for $h \geq 1$ and $n \in \{0,..,\mathit{Tower}(h{+}1){-}1\} = \{0,..,2^{\mathit{Tower}(h)}{-}1\}$, the $\{L,X\}$-labelled tree $\tilde{\mathcal{T}}_{h+1}(n)$ is obtained by creating a new root, attaching to it one copy of $\tilde{\mathcal{T}}_h(i)$, for each $i \in \{0,..,\mathit{Tower}(h){-}1\}$, and labelling the root of $\tilde{\mathcal{T}}_h(i)$ with **1** if $\mathrm{bit}(i,n) = 1$, and **0** otherwise.

Note that for every fixed h, the trees $\tilde{\mathcal{T}}_h(n)$ for $n < \mathit{Tower}(h)$ all have the same shape and only vary in the labelling (w.r.t. **0** and **1**) of the children of the root. Furthermore, each path from the root of $\tilde{\mathcal{T}}_h(n)$ to a leaf has exactly length h (i.e., consists of h edges), and the nodes that are labelled L are exactly the *leaves* of $\tilde{\mathcal{T}}_h(n)$.

Unlike in the previous sections, it does not suffice to restrict attention to structures that are obtained as disjoint unions or similar, easy combinations of the trees $\tilde{\mathcal{T}}_h(n)$. Instead, we will consider a suitable notion where a node t in an arbitrary τ-structure $\mathcal{A}$ is called "h-*good*" if the substructure $\mathcal{A}_t$ is "sufficiently similar" to the tree $\tilde{\mathcal{T}}_h(n)$, for a number $n < \mathit{Tower}(h)$. The precise definition of this notion is given below. Before introducing it, however, we need the following (easy) lemma.

Lemma 6.3. *For every $h' \geq 1$ there is a universal* FO(τ)*-sentence $\mathit{forest}_{\leq h'}$ of length $\mathcal{O}(h')$ such that for every finite τ-structure $\mathcal{A} = (A, E^{\mathcal{A}}, L^{\mathcal{A}}, X^{\mathcal{A}})$ the following is true: $\mathcal{A} \models \mathit{forest}_{\leq h'} \iff (A, E^{\mathcal{A}})$ is a disjoint union of trees such that every node in $L^{\mathcal{A}}$ is a leaf, and for every root r in $\mathcal{A}$ (i.e., for every node r in A that has in-degree 0 in $E^{\mathcal{A}}$) the following is true: every path in $\mathcal{A}$ that starts in r has length at most h'.*

Definition 6.4 (h-good nodes x, and the numbers $\mathit{Rep}_h^{\mathcal{A}}(x)$ represented by them). Let $h' \geq 1$ and let $\mathcal{A}$ be a structure with $\mathcal{A} \models \mathit{forest}_{\leq h'}$. By induction on $h \in \{1,..,h'\}$ we define the following notion:

A node x of $\mathcal{A}$ is called 1-*good in* $\mathcal{A}$ iff it has at least one child y with $L^{\mathcal{A}}(y)$, and for all children y' of x in $\mathcal{A}$ the following is true: if $L^{\mathcal{A}}(y')$, then $X^{\mathcal{A}}(y') \leftrightarrow X^{\mathcal{A}}(y)$.
Every 1-good node x in $\mathcal{A}$ represents a number $\mathit{Rep}_1^{\mathcal{A}}(x) \in \{0,1\}$ as follows:

$$\mathit{Rep}_1^{\mathcal{A}}(x) = 0 \iff x \text{ has a child that belongs to } L^{\mathcal{A}} \text{ but not to } X^{\mathcal{A}}$$
$$\mathit{Rep}_1^{\mathcal{A}}(x) = 1 \iff x \text{ has a child that belongs to } L^{\mathcal{A}} \text{ and to } X^{\mathcal{A}}.$$

Let $h < h'$ be such that the notion of h-goodness as well as the numbers $\mathit{Rep}_h^{\mathcal{A}}(y)$, for all h-good nodes y in $\mathcal{A}$, are already defined. Then, a node x of $\mathcal{A}$ is called $(h{+}1)$-*good in* $\mathcal{A}$ iff the following is true: For each number $i \in \{0,..,\mathit{Tower}(h){-}1\}$ there exists a h-good child y_i of x in $\mathcal{A}$ with $\mathit{Rep}_h^{\mathcal{A}}(y_i) = i$, and for all h-good children z of x in $\mathcal{A}$ with $\mathit{Rep}_h^{\mathcal{A}}(z) = i$ the following is true: $X^{\mathcal{A}}(z) \leftrightarrow X^{\mathcal{A}}(y_i)$.
Every $(h{+}1)$-good node x in $\mathcal{A}$ represents the (uniquely defined) number

$$\mathit{Rep}_{h+1}^{\mathcal{A}}(x) = n \ \in \{0,1,..,2^{\mathit{Tower}(h)}{-}1\} = \{0,1,..,\mathit{Tower}(h{+}1){-}1\}$$

which satisfies the following: for every $i < \mathit{Tower}(h)$, $\mathrm{bit}(i,n) = 1 \iff X^{\mathcal{A}}(y_i)$.

The following notion of h-*inconsistency* can be viewed as a counterpart to the notion of h-goodness. Note, however, that h-goodness is a property of a *node* whereas h-inconsistency is a property of a whole structure.

Definition 6.5 (**h-inconsistency**)**.** Let $h' \geq 1$ and let $\mathcal{A}$ be a structure with $\mathcal{A} \models$ *forest*$_{\leq h'}$. By induction on $h \in \{1,..,h'\}$, we define the following notion:

We say that $\mathcal{A}$ is 1-*inconsistent* if there exist nodes x, y, y' such that y and y' are children of x with the following properties: y and y' both belong to $L^{\mathcal{A}}$, and we have $X^{\mathcal{A}}(y)$ and $\neg X^{\mathcal{A}}(y')$.

Let $h < h'$ be such that the notion of h-inconsistency is already defined.

We say that $\mathcal{A}$ is $(h{+}1)$-*inconsistent* if there exist nodes x, y, y' such that y and y' are children of x with the following properties: y and y' both are h-good in $\mathcal{A}$ with $\mathit{Rep}_h^{\mathcal{A}}(y) = \mathit{Rep}_h^{\mathcal{A}}(y')$, and we have $X^{\mathcal{A}}(y)$ and $\neg X^{\mathcal{A}}(y')$.

Furthermore, we say that $\mathcal{A}$ is $(\leq h)$-*inconsistent* if there exists a $\tilde{h} \in \{1,..,h\}$ such that $\mathcal{A}$ is $\tilde{h}$-inconsistent. It is straightforward (but tedious) to show the following:

Lemma 6.6. *For every $h \geq 1$ there is a* FO(τ)*-sentence φ_h of size $2^{\mathcal{O}(h)}$ such that the following is true for every τ-structure $\mathcal{A}$: $\mathcal{A} \models \varphi_h \iff \mathcal{A} \models \neg$forest$_{\leq h}$ or $\mathcal{A}$ is $(\leq h)$-inconsistent or there exists a node x that is h-good in $\mathcal{A}$.*

Furthermore, it can be shown that this sentence φ_h is preserved under extensions. This finally enables us to prove Theorem 6.1.

7 Structures of bounded degree — elementary upper bounds

All the non-elementary lower bounds in previous sections depend heavily on the fact that we consider classes of structures of unbounded degree. On classes of structures of *bounded* degree, the picture looks entirely different as we can prove elementary upper bounds as counterparts to Theorems 4.1, 5.1, and 6.1. Throughout the remainder of this section we let τ be a fixed finite relational vocabulary, and we let d be a fixed natural number. We write $\mathfrak{D}_d$ to denote the class of all τ-structures whose Gaifman graph has degree at most d. By an easy adaption of the model theoretic proof of Gaifman's theorem given in [5], one obtains the following elementary *upper* bound:

Theorem 7.1. *There is a 4-fold exponential function $g : \mathbb{N} \to \mathbb{N}$ such that for every* FO(τ)*-sentence φ there is a sentence ψ of size $\leq g(||\varphi||)$ with the following properties: ψ is a Boolean combination of basic local sentences and ψ is equivalent to φ on all structures in $\mathfrak{D}_d$.*

By similar techniques we can prove an elementary upper bound for the Feferman-Vaught theorem stratified by formula length. Furthermore, there are elementary decision algorithms for the first-order theories of classes of trees of bounded arity, in particular for the class of binary trees. Refining the methods of [1], one also obtains an elementary upper bound for the following variant of the Łoś-Tarski Theorem.

Theorem 7.2. *There is a 5-fold exponential function $f : \mathbb{N} \to \mathbb{N}$ such that any* FO(τ)*-sentence φ that is preserved under extensions on the class of acyclic structures in $\mathfrak{D}_d$ is equivalent, on this class, to an existential first-order sentence of length at most $f(||\varphi||)$.*

In all the above cases for structures of bounded degree we can also prove at least 2-fold exponential lower bounds.

References

1. A. Atserias, A. Dawar, and M. Grohe. Preservation under extensions on well-behaved finite structures. In *Proc. ICALP'05*, volume 3580 of *Spinger LNCS*, pages 1437–1449, 2005.
2. K. Compton and C. W. Henson. A uniform method for proving lower bounds on the computational complexity of logical theories. In S. Abramsky, D. M. Gabbay, and T. Maibaum, editors, *Handbook of Logic in Computer Science*, volume 5: Logic and Algebraic Methods, pages 129–216. Oxford University Press, 2000.
3. N. J. Cutland. *Computability*. Cambridge University Press, 1980.
4. A. Dawar, M. Grohe, S. Kreutzer, and N. Schweikardt. Approximation schemes for first-order definable optimisation problems. In *Proc. LICS'06*, pages 411–420, 2006.
5. H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.
6. S. Feferman and R. Vaught. The first order properties of products of algebraic systems. *Fundamenta Mathematicae*, 47:57–103, 1959.
7. J. Ferrante and C. Rackoff. *Computational Complexity of Logical Theories*, volume 718 of *Lecture Notes in Mathematics*. Springer-Verlag, 1979.
8. J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.
9. M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM*, 48:1184–1206, 2001.
10. M. Frick and M. Grohe. The complexity of first-order and monadic second-order logic revisited. *Annals of Pure and Applied Logic*, 130:3–31, 2004.
11. H. 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.
12. M. Grohe and N. Schweikardt. Comparing the succinctness of monadic query languages over finite trees. *RAIRO: Theoretical Informatics and Applications (ITA)*, 38:343–373, 2005.
13. M. Grohe and N. Schweikardt. The succinctness of first-order logic on linear orders. *Logical Methods in Computer Science*, 1(1:6):1–25, 2005.
14. Y. Gurevich. Toward logic tailored for computational complexity. In M. Richter et al., editor, *Computation and Proof Theory*, volume 1104 of *Lecture Notes in Mathematics*, pages 175–216. Springer, 1984.
15. W. Hodges. *Model Theory*. Cambridge University Press, 1993.
16. L. Libkin. On forms of locality over finite models. In *Proceedings of the 12th IEEE Symposium on Logic in Computer Science*, pages 204–215, 1997.
17. L. Libkin. Logics with counting and local properties. *Transaction on Computational Logic*, 1:33–59, 2000.
18. G. Pan and M. Vardi. Fixed-parameter hierarchies inside pspace. In *Proceedings of the 21th IEEE Symposium on Logic in Computer Science*, pages 27–36, 2006.
19. M. Rabin. Decidability of second order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969.
20. B. Rossman. Existential positive types and preservation under homomorphisisms. In *20th IEEE Symposium on Logic in Computer Science*, pages 467–476, 2005.
21. L. Stockmeyer and A. Meyer. Word problems requiring exponential time. In *Proceedings of the 5th ACM Symposium on Theory of Computing*, pages 1–9, 1973.
22. W. W. Tait. A counterexample to a conjecture of Scott and Suppes. *Journal of Symbolic Logic*, 24:15–16, 1959.
23. J. Thatcher and J. Wright. Generalised finite automata theory with an application to a decision problem of second-order logic. *Mathematical Systems Theory*, 2:57–81, 1968.