Model theory makes formulas large

Anuj Dawar* Martin Grohe** Stephan Kreutzer** Nicole Schweikardt**

* University of Cambridge, U.K., Email: anuj.dawar@ cl.cam.ac.uk
** Humboldt-Universität zu Berlin, Germany, Email: {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. Gaifman's theorem is an essential ingredient in several algorithmic meta theorems for first order logic. Our result has direct implications for the running time of the algorithms.

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 trees and more generally classes of structures of bounded tree width. Our lower bound also applies to these variants.

The first-order theory of trees is decidable. We prove that there is no elementary decision algorithm.

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 [10, 6, 17, 18]. In particular, there are very general algorithmic meta-theorems stating that deciding first-order properties of various classes of structures, such as planar graphs or graphs with excluded minors, is fixed-parameter tractable, and that first-order definable optimization 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 trees, structures of bounded tree width, and structures of bounded degree. These results are part of recent efforts in finite model-theory towards developing a model theory for "well-behaved" classes of finite structures [1, 2, 3].

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 question for 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 the classes of finite trees and all classes of finite structures that contain all trees (but not for classes of finite structures of bounded degree).

We prove two further, related lower bounds. The first is concerned with the classical decision problem. It is known that the first-order theory (and actually also the monadic second-order theory) of trees is decidable [27, 22]. We prove that there is no elementary decision algorithm. Finally, we prove that a version of the Feferman-Vaught theorem based on a restriction of formulas by formula length necessarily entails a non-elementary blow-up in formula size.

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. Let us emphasize, however, 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 [24]. Much later, such techniques have been employed in [11, 20, 13, 14] to prove lower bounds in parameterized complexity, respectively, on the succinctness of monadic logics. Let us also mention a related succinctness lower bound: It has recently been proved by Rossman [23] 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.[1]

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), the lower bounds on the complexity of deciding the first-order theory of trees (Section 5) and also the failure of the Feferman-Vaught theorem for formula length (Section 6). Section 7 then establishes the lower bound for the Łoś-Tarski theorem, which is based on a different encoding of numbers by trees. Finally, Section 8 contains the elementary upper bounds on classes of structures of bounded degree.

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., [5]).

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

[1] Rossman mentions this in his paper, referring to unpublished work of Gurevich and Shelah. As far as we know, a proof of this lower bound result has not been published yet.

symbol R is a positive integer called the *arity* of R. In the following, τ always denotes 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-neighborhood 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 (precisely, 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 of φ's syntax tree.

E always 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 leafs. 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, 5, and 6. We use the following encoding of natural numbers by trees, introduced in [9].

Definition 3.1 (Encoding numbers by trees). For natural numbers i, 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 $\lfloor \frac{n}{2^i} \rfloor$ is even, and bit$(i, n) = 1$ if $\lfloor \frac{n}{2^i} \rfloor$ 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$(i, n) = 1$.

These trees are illustrated in Figures 1 and 2. Figure 1 shows the trees $\mathcal{T}(0)$ up to $\mathcal{T}(10)$, and Figure 2 shows the tree $\mathcal{T}(2^{2^{10}})$.

It is straightforward to see (cf. [9, Lemma 10.20])[2] that

$$\text{for all } h, n \geq 0, \quad height(\mathcal{T}(n)) \leq h \iff n < Tower(h)\,. \tag{1}$$

The next lemma from [9] shows that the tree encodings of numbers can be "controlled" by small first-order formulas. (In [9], 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 ([9, Lemma 10.21]). *For every $h \geq 0$ there is a* 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* $\mathcal{A} \models eq_h(t, u) \iff m = n\,.$

[2]Note that our function $Tower(h)$ is slightly different than [9]'s function $tow(h)$; precisely, we have $tow(h) = Tower(h-1)$. Thus, equation (1) looks slightly different than in [9, Lemma 10.20].

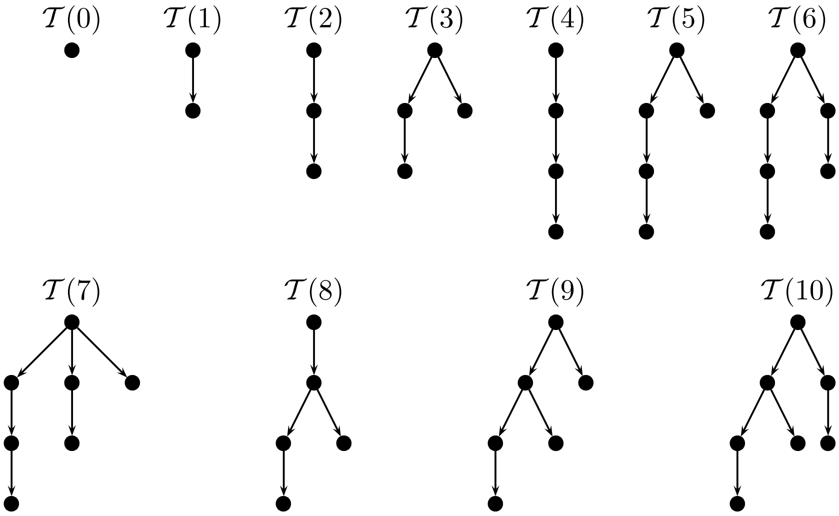

Figure 1: The trees $\mathcal{T}(0)$ up to $\mathcal{T}(10)$

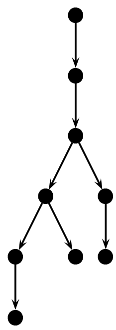

Figure 2: The tree $\mathcal{T}(2^{2^{10}})$

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

Proof. Since *Tower*$(0) = 1$, the formula $\mathit{encoding}_0$ has to express that $\mathcal{T}_t$ is isomorphic to $\mathcal{T}(0)$, i.e., to the one-node tree. We can thus choose $\mathit{encoding}_0(x) := \neg\exists y\, E(x, y)$.

For $h \geq 1$ we choose

$$\mathit{encoding}_h(x) \;:=$$
$$\forall y\big(E(x,y) \rightarrow \mathit{encoding}_{h-1}(y)\big) \;\wedge\; \forall y\, \forall y' \Big(\big(E(x,y) \wedge E(x,y') \wedge \neg y{=}y'\big) \rightarrow \neg \mathit{eq}_{h-1}(y,y')\Big)$$

It is straightforward to see that this formula has the intended meaning. Furthermore, considering the length of the formula $\mathit{encoding}_h$, there is a $c > 0$ such that

$$\begin{aligned} ||\mathit{encoding}_h|| &\leq ||\mathit{encoding}_{h-1}|| + ||\mathit{eq}_{h-1}|| + c \\ &\leq ||\mathit{encoding}_{h-2}|| + ||\mathit{eq}_{h-2}|| + ||\mathit{eq}_{h-1}|| + 2{\cdot}c \\ &\leq \cdots \\ &\leq ||\mathit{encoding}_0|| + \textstyle\sum_{h'=0}^{h-1} ||\mathit{eq}_{h'}|| + c{\cdot}h \quad = \; \mathcal{O}(h^2) \end{aligned}$$

(recall that due to Lemma 3.2, $||\mathit{eq}_{h'}|| = \mathcal{O}(h'))$. □

Lemma 3.4. *For every $h \geq 0$ there are* 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) $\mathcal{A} \models bit_h(t,u) \iff \text{bit}(m,n) = 1$.
(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* $\mathcal{T}(Tower(h){-}1)$.

Proof. (a): Choose $bit_h(x,y) := \exists y' \big(E(y,y') \wedge eq_h(y',x)\big)$.

(b): We define the formulas $less_h(x,y)$ by induction on h.

Since $Tower(0) = 1$, we know that $m, n < Tower(0) \iff m = n = 0$. Thus, $less_0(x,y)$ can be chosen as a formula that is *never* satisfied, e.g. $less_0(x,y) := \exists x\, \neg x{=}x$.
For $h \geq 1$ we choose

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

Along Definition 3.1 it is straightforward to see that the formula $less_h$ expresses that there exists an i (corresponding to the variable y' in $less_h$) such that $\text{bit}(i,n) = 1$, $\text{bit}(i,m) = 0$, and for each j (corresponding to the variable x'' in $less_h$) with $j > i$ and $\text{bit}(j,m) = 1$ we have $\text{bit}(j,n) = 1$. Thus, the formula expresses that $m < n$. Furthermore, considering the length of the formula $less_h$, there is a $c > 0$ such that

$$\begin{aligned} ||less_h|| &\leq ||less_{h-1}|| + 2{\cdot}||eq_{h-1}|| + c \\ &\leq ||less_{h-2}|| + 2{\cdot}||eq_{h-2}|| + 2{\cdot}||eq_{h-1}|| + 2{\cdot}c \\ &\leq \cdots \\ &\leq ||less_0|| + 2 \cdot \textstyle\sum_{h'=0}^{h-1} ||eq_{h'}|| + c{\cdot}h \quad = \mathcal{O}(h^2) \end{aligned}$$

(recall that due to Lemma 3.2, $||eq_{h'}|| = \mathcal{O}(h')$).

(c): Since $\mathcal{T}(0)$ is the one-node tree, we can choose $min(x) := \neg\exists y\, E(x,y)$.

(d): We define $succ_h$ by induction on h and make use of the formulas eq_{h-1} and $less_{h-1}$. Of course, for $h = 0$, $succ_0$ can be chosen to be a formula that is *never* satisfied.

For $h \geq 1$, the first two lines of the following formula $succ_h$ express that there is a number i (corresponding to the variable y') such that i is the smallest number with $\text{bit}(i,n) = 1$, and that for this particular i we have $\text{bit}(i,m) = 0$. Lines 3 and 4 express for each $j > i$ that $\text{bit}(j,n) = 1 \Longrightarrow \text{bit}(j,m) = 1$ and, vice versa, $\text{bit}(j,m) = 1 \Longrightarrow \text{bit}(j,n) = 1$. The last two lines expresses the following: If $i \neq 0$ then $\text{bit}(0,m) = 1$ and for each $j < i$ with $\text{bit}(j,m) = 1$ we have $\big(j{+}1 = i$ or $\text{bit}(j{+}1,m) = 1\big)$.

Altogether, the formula $succ_h$ hence expresses that $m + 1 = n$.

$$\begin{aligned} & succ_h(x,y) \;:=\; \exists y' \Big(\\ & \quad E(y,y') \wedge \forall y'' \Big((E(y,y'') \wedge \neg y''{=}y') \rightarrow less_{h-1}(y',y'')\Big) \wedge \forall x' \Big(E(x,x') \rightarrow \neg eq_{h-1}(x',y')\Big) \wedge \\ & \quad \forall y'' \Big(\big(E(y,y'') \wedge less_{h-1}(y',y'')\big) \rightarrow \exists x'' \big(E(x,x'') \wedge eq_{h-1}(x'',y'')\big)\Big) \wedge \\ & \quad \forall x'' \Big(\big(E(x,x'') \wedge less_{h-1}(y',x'')\big) \rightarrow \exists y'' \big(E(y,y'') \wedge eq_{h-1}(y'',x'')\big)\Big) \wedge \\ & \quad \neg min(y') \rightarrow \Big(\exists x' \big(E(x,x') \wedge min(x')\big) \wedge \\ & \qquad\qquad \forall x' \Big(\big(E(x,x') \wedge less_{h-1}(x',y')\big) \rightarrow \big(\exists z (succ_{h-1}(x',z) \wedge (z{=}y' \vee E(x,z)))\big)\Big)\Big) . \end{aligned}$$

Since $||less_{h-1}|| = \mathcal{O}((h{-}1)^2)$ and $||eq_{h-1}|| = \mathcal{O}(h{-}1)$, we obtain (in a similar way as in the proof of Lemma 3.3) that $||succ_h|| = \mathcal{O}(\sum_{h' \leq h} h'^2) = \mathcal{O}(h^3)$.

(e): Since $\mathit{Tower}(0) = 1$, the formula max_0 has to express that $\mathcal{A}_t$ is isomorphic to $\mathcal{T}(0)$, i.e., to the one-node tree. We can thus choose $\mathit{max}_0(x) := \neg\exists y\, E(x,y)$.

For $h \geq 1$ we choose

$$\mathit{max}_h(x) \;:=\; \mathit{encoding}_h(x) \;\wedge\; \mathit{max}'_h(x)\,,$$

where the formula $\mathit{max}'_h(x)$ is defined by induction on h as follows: For $h = 0$ we let $\mathit{max}'_0(x) := \mathit{max}_0(x) := \neg\exists y\, E(x,y)$. For $h \geq 1$ note that $\mathit{Tower}(h){-}1 = 2^{\mathit{Tower}(h-1)}{-}1$. Thus, the tree $\mathcal{T}' := \mathcal{T}(\mathit{Tower}(h){-}1)$ consists of a root node which, for each i with $0 \leq i < \mathit{Tower}(h{-}1)$, has a child t_i such that $\mathcal{T}'_{t_i}$ is isomorphic to $\mathcal{T}(i)$. In particular, the root of $\mathcal{T}'$ has a child t_0 and a child t_{max} such that $\mathcal{T}'_{t_0}$ is isomorphic to $\mathcal{T}(0)$ and $\mathcal{T}'_{t_{max}}$ is isomorphic to $\mathcal{T}(\mathit{Tower}(h{-}1){-}1)$. We choose

$$\begin{aligned}\mathit{max}'_h(x) \;:=\;& \exists y\,\big(E(x,y)\wedge \mathit{min}(y)\big) \;\wedge\\ & \forall y\,\Big(E(x,y) \;\rightarrow\; \big(\,\mathit{max}'_{h-1}(y) \;\vee\; \exists z\,\big(E(x,z)\wedge \mathit{succ}_{h-1}(y,z)\big)\,\big)\Big)\,.\end{aligned}$$

It is straightforward to see that the formula $\mathit{max}_h(x)$ expresses the desired property. Furthermore, since $||\mathit{succ}_{h-1}|| = \mathcal{O}((h{-}1)^3)$ and $||\mathit{eq}_{h-1}|| = \mathcal{O}(h{-}1)$, we obtain (in a similar way as in the proof of Lemma 3.3) that $||\mathit{max}'_h|| = \mathcal{O}(\sum_{h'\leq h} h'^3) = \mathcal{O}(h^4)$. Thus, also $||\mathit{max}_h|| = \mathcal{O}(h^4)$.

Finally, the proof of Lemma 3.4 is complete.

□

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 $\mathit{dist}_{\leq r}(x,y)$ be an FO(τ)-formula expressing that the distance between x and y is at most r. We often write $\mathit{dist}(x,y) \leq r$ instead of $\mathit{dist}_{\leq r}(x,y)$ and $\mathit{dist}(x,y) > r$ or $\mathit{dist}_{>r}(x,y)$ instead of $\neg\mathit{dist}_{\leq r}(x,y)$. An 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

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

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

For an FO(τ)-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 from [12] states that every first-order sentence over a relational vocabulary is equivalent to a first-order sentence in Gaifman normal form.

Theorem 4.1 (Gaifman [12]). *Every first-order sentence over a relational vocabulary is equivalent to a first-order sentence in Gaifman normal form.*

The proof in [12] 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 present section's main result shows that this huge increase of 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 4.2. *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 $\mathfrak{T}$ of finite trees has size at least $\mathit{Tower}(h)$.*

Before proving Theorem 4.2, we first show the following variant that speaks about the class of all *forests* rather than *trees*. The proof of Theorem 4.3 will avoid some of the unpleasant details needed in the proof of Theorem 4.2 while still exposing the main ideas that are crucial for the proof of Theorem 4.2.

Theorem 4.3. *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

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

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 (2) that

$$\mathcal{F}_h \models \psi \qquad \text{and,} \quad \text{for each } i < \mathit{Tower}(h), \quad \mathcal{F}_h^{-i} \not\models \psi\,. \tag{3}$$

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,

$$k_1 + \cdots + k_L \;\leq\; ||\psi|| \;=:\; H\,. \tag{4}$$

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

$$\text{for each } \ell \leq \tilde{L}, \quad \mathcal{F}_h \models \chi_\ell\,, \qquad \text{and} \qquad \text{for each } \ell > \tilde{L}, \quad \mathcal{F}_h \not\models \chi_\ell\,. \tag{5}$$

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

$$\bigwedge_{1\leq i<j\leq k_\ell} \mathit{dist}(x_i,x_j) > 2r_\ell \;\wedge \bigwedge_{1\leq i\leq k_\ell} \psi_\ell(x_i) \tag{6}$$

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 (4)). Since we assume that $H < \mathit{Tower}(h)$, and since $\mathcal{F}_h$ consists of *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 (6), 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-neighborhood of $t_i^{(\ell)}$, which is the same in $\mathcal{F}_h^{-j}$ as in $\mathcal{F}_h$. We thus obtain from (6) 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 (5) 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:

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

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:

$$\text{for each } \ell \le \tilde{L}, \quad \mathcal{F}_h^{-j} \models \chi_\ell\,, \qquad \text{and} \qquad \text{for each } \ell > \tilde{L}, \quad \mathcal{F}_h^{-j} \not\models \chi_\ell\,. \tag{7}$$

From (7) and (5) 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 (3). Altogether, the proof of Theorem 4.3 is complete. □

We are now ready for the proof of Theorem 4.2.

Proof of Theorem 4.2:
We use a modification of the proof of Theorem 4.3, where instead of a single forest $\mathcal{F}_h$ we consider a series of trees $\mathcal{F}_{h,R}$ for all $R > 2h$. Precisely, for each $R > 2h$ the structure $\mathcal{F}_{h,R}$ is defined as follows:

- $\mathcal{F}_{h,R}$ contains a disjoint copy $\mathcal{T}_j$ of the tree $\mathcal{T}(j)$, for each $j \in \{0, .., \mathit{Tower}(h){-}1\}$.
- Additionally, there is a path $a_0, a_1, .., a_{R\cdot(\mathit{Tower}(h)+1)}$ of $1 + R\cdot(\mathit{Tower}(h){+}1)$ distinct nodes that do not belong to any of the trees $\mathcal{T}_j$.
- For each number $j \in \{0, .., \mathit{Tower}(h){-}1\}$ there is an edge from node $a_{R\cdot(j+1)}$ to the root of $\mathcal{T}_j$.

Note that the resulting structure $\mathcal{F}_{h,R}$ is indeed a *tree*.

For each $i \in \{0, .., \mathit{Tower}(h){-}1\}$ we let $\mathcal{F}_{h,R}^{-i}$ be the tree obtained from $\mathcal{F}_{h,R}$ by deleting the entire subtree $\mathcal{T}_i$.

Instead of the formula $\mathit{root}(x)$ from the proof of Theorem 4.3, we now choose $\mathit{root}_h(x)$ to be a formula which expresses the following:

1. every directed path starting in x has length at most h, and
2. there is a node y for which we have $E(y, x)$ and for which there exists a directed path of length at least $2h$ that starts in y.

It should be clear that this can be formalised by a FO(E)-formula $\mathit{root}_h(x)$ of size $\mathcal{O}(h)$. Furthermore, if $\mathcal{A}$ is the tree $\mathcal{F}_{h,R}$ or one of the trees $\mathcal{F}_{h,R}^{-i}$, and t is a node in $\mathcal{A}$, then

$$\mathcal{A} \models \mathit{root}_h(t) \iff t \text{ is the root of one of the trees } \mathcal{T}_j\,.$$

Now, the sentence φ_h is chosen in the same way as in the proof of Theorem 4.3, where instead of the formula $\mathit{root}(x)$ now the formula $\mathit{root}_h(x)$ is used. One then obtains that $||\varphi_h|| = \mathcal{O}(h^4)$ and, for all $R > 2h$,

$$\mathcal{F}_{h,R} \models \varphi_h \qquad \text{and,} \quad \text{for each } i < \mathit{Tower}(h), \quad \mathcal{F}_{h,R}^{-i} \not\models \varphi_h\,. \tag{8}$$

Now let ψ be a FO(E)-sentence in *Gaifman normal form* that is equivalent to φ_h on the class $\mathfrak{T}$ of finite trees. In particular, since $\mathcal{F}_{h,R}$ as well as all the $\mathcal{F}_{h,R}^{-i}$ belong to $\mathfrak{T}$, we obtain from (8) that

$$\mathcal{F}_{h,R} \models \psi \qquad \text{and,} \quad \text{for each } i < \mathit{Tower}(h), \quad \mathcal{F}_{h,R}^{-i} \not\models \psi\,. \tag{9}$$

Our aim is to show that $H := ||\psi|| \ge \mathit{Tower}(h)$. 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* $\chi_1, .., \chi_L$. We choose the same notation concerning these sentences as in the proof of Theorem 4.3, i.e., k_ℓ, r_ℓ denote the according parameters of χ_ℓ. In particular, we again have

$$k_1 + \cdots + k_L \leq ||\psi|| =: H . \tag{10}$$

Letting $r := \max\{r_1, .., r_L\}$, we now choose a particular number R as follows:

$$R := 1 + \max\{2h, 4r\} . \tag{11}$$

I.e., from now on R is a fixed number which is larger than $2h$ and larger than four times the "locality radius" of each of the formulas $\chi_1, .., \chi_L$. Using this fixed R, we can assume w.l.o.g. that there exists a $\tilde{L}$ with $0 \leq \tilde{L} \leq L$ such that

$$\text{for each } \ell \leq \tilde{L}, \quad \mathcal{F}_{h,R} \models \chi_\ell , \qquad \text{and} \qquad \text{for each } \ell > \tilde{L}, \quad \mathcal{F}_{h,R} \not\models \chi_\ell . \tag{12}$$

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

$$\bigwedge_{1 \leq i < j \leq k_\ell} \mathit{dist}(x_i, x_j) > 2r_\ell \ \wedge \bigwedge_{1 \leq i \leq k_\ell} \psi_\ell(x_i) \tag{13}$$

is satisfied in $\mathcal{F}_{h,R}$ 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 a total number of at most $k_1 + \cdots + k_{\tilde{L}} \leq H$ nodes (see (10)). Since we assume that $H < \mathit{Tower}(h)$, and since $\mathcal{F}_{h,R}$ contains all the trees $\mathcal{T}_j$, for all $j \in \{0, .., \mathit{Tower}(h){-}1\}$, and they are of pairwise distance $\geq R$ in $\mathcal{F}_{h,R}$, there must be at least one $j \in \{0, .., \mathit{Tower}(h){-}1\}$ such that none of the nodes from $\{ \, t_i^{(\ell)} \ : \ \ell \leq \tilde{L} \text{ and } i \leq k_\ell \, \}$ is contained in the $\frac{R}{2}$-neighbourhood of $\mathcal{T}_j$.

Now, of course, the tree $\mathcal{F}_{h,R}^{-j}$, which is obtained from $\mathcal{F}_{h,R}$ by removing the subtree $\mathcal{T}_j$, still contains all the nodes in $\{ \, t_i^{(\ell)} \ : \ \ell \leq \tilde{L} \text{ and } i \leq k_\ell \, \}$. Furthermore, since $\frac{R}{2} \geq r_\ell$, the r_ℓ-neighbourhood of node $t_i^{(\ell)}$ in $\mathcal{F}_{h,R}^{-j}$ is isomorphic to the r_ℓ-neighbourhood of node $t_i^{(\ell)}$ in $\mathcal{F}_{h,R}$. Considering (13), 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 $\mathcal{F}_{h,R}^{-j}$ as in $\mathcal{F}_{h,R}$. We thus obtain from (13) that

$$\text{for each } \ell \leq \tilde{L}, \quad \mathcal{F}_{h,R}^{-j} \models \chi_\ell .$$

Let us now consider the formulas χ_ℓ with $\ell > \tilde{L}$. From (12) we know that $\mathcal{F}_{h,R} \not\models \chi_\ell$.

Claim 4.4. *Let $\mathcal{F}_{h,R} \not\models \chi_\ell$. Then, also $\mathcal{F}_{h,R}^{-j} \not\models \chi_\ell$.*

Proof. Aiming at a contradiction, let us assume that $\mathcal{F}_{h,R}^{-j} \models \chi_\ell$. Then, there are nodes $t_1, .., t_{k_\ell}$ in $\mathcal{F}_{h,R}^{-j}$ such that the formula

$$\bigwedge_{1 \leq i < j \leq k_\ell} \mathit{dist}(x_i, x_j) > 2r_\ell \ \wedge \bigwedge_{1 \leq i \leq k_\ell} \psi_\ell(x_i) \tag{14}$$

is satisfied in $\mathcal{F}_{h,R}^{-j}$ when interpreting each variable x_i with the node t_i. I.e., in $\mathcal{F}_{h,R}^{-j}$ the nodes $t_1, .., t_{k_\ell}$ are of pairwise distance greater than $2r_\ell$ and the r_ℓ-neighbourhood of each t_i in $\mathcal{F}_{h,R}^{-j}$ satisfies $\psi_\ell(t_i)$. Our aim now is to show that also $\mathcal{F}_{h,R} \models \chi_\ell$, contradicting the claim's assumption.

Of course, each of the nodes $t_1, .., t_{k_\ell}$ also occurs in $\mathcal{F}_{h,R}$, and the pairwise distance of these nodes in $\mathcal{F}_{h,R}$ is the same as in $\mathcal{F}_{h,R}^{-j}$, and thus $> 2r_\ell$. However, the r_ℓ-neighbourhood of some of the nodes may be different in $\mathcal{F}_{h,R}$ than in $\mathcal{F}_{h,R}^{-j}$. Namely, the r_ℓ-neighbourhood is different for exactly those nodes from

$\{t_1,..,t_{k_\ell}\}$ that belong to the $(r_\ell-1)$-neighbourhood of the node $a_{R\cdot(j+1)}$ (i.e., the node to which $\mathcal{T}_j$ is attached in $\mathcal{F}_{h,R}$). Of course, these are exactly the nodes in

$$V \;:=\; \{\, t_1,..,t_{k_\ell} \,\} \,\cap\, \{\, a_{R\cdot(j+1)+i} \,:\, -r_\ell < i < r_\ell \,\}.$$

Since the vertices $t_1,..,t_{k_\ell}$ have pairwise distance $> 2r_\ell$, we obtain that $|V| \leq 1$.

If $|V| = 0$, then the formula from (14) is obviously satisfied in $\mathcal{F}_{h,R}$ when interpreting the variables $x_1,..,x_{k_\ell}$ with the nodes $t_1,..,t_{k_\ell}$.

If $|V| = 1$ we can assume w.l.o.g. that $V = \{t_1\}$. Note that the r_ℓ-neighbourhood of t_1 in $\mathcal{F}_{h,R}^{-j}$ is just a path of $2r_\ell + 1$ vertices with t_1 being in the middle (whereas the r_ℓ-neighbourhood of t_1 in $\mathcal{F}_{h,R}$ contains vertices from $\mathcal{T}_j$). To satisfy the formula in (14), it suffices to find a node t_1' in $\mathcal{F}_{h,R}$ which has distance $> 2r_\ell$ to each of the nodes $t_2,..,t_{k_\ell}$ and whose r_ℓ-neighbourhood in $\mathcal{F}_{h,R}$ is a path of $2r_\ell + 1$ vertices with t_1' being in the middle.

Since $k_\ell - 1 < H < \mathit{Tower}(h)$ (see (10)) and $\frac{R}{2} > 2r_\ell$ (see (11)), we can find an $i \in \{0,..,\mathit{Tower}(h)\}$ such that none of the nodes $t_2,..,t_{k_\ell}$ belongs to the $2r_\ell$-neighbourhood of the node $a_{R\cdot i+(R/2)}$. Choosing $t_1' := a_{R\cdot i+(R/2)}$, we have found a node in $\mathcal{F}_{h,R}$ that has distance $> 2r_\ell$ to each of the nodes $t_2,..,t_{k_\ell}$ and whose r_ℓ-neighbourhood in $\mathcal{F}_{h,R}$ is isomorphic to the r_ℓ-neighbourhood of t_1 in $\mathcal{F}_{h,R}^{-j}$. Altogether, we thus obtain that the formula from (14) is satisfied in $\mathcal{F}_{h,R}$ when interpreting the variables $x_1, x_2,..,x_{k_\ell}$ with the nodes $t_1', t_2,..,t_\ell$. Hence, $\mathcal{F}_{h,R} \models \chi_\ell$, contradicting the claim's assumption and thus completing the proof of Claim 4.4. □

Proceeding with the proof of Theorem 4.2, note that we are now in a situation where we know the following:

$$\text{for each } \ell \leq \tilde{L},\quad \mathcal{F}_{h,R}^{-j} \models \chi_\ell\,, \qquad \text{and} \qquad \text{for each } \ell > \tilde{L},\quad \mathcal{F}_{h,R}^{-j} \not\models \chi_\ell\,. \tag{15}$$

From (15) and (12) we thus obtain that $\mathcal{F}_{h,R}^{-j}$ satisfies exactly the same basic local sentences from $\{\chi_1,\ldots,\chi_L\}$ as $\mathcal{F}_{h,R}$. Since ψ is a Boolean combination of the sentences $\chi_1,..,\chi_L$, we thus have that $\mathcal{F}_{h,R}^{-j} \models \psi \iff \mathcal{F}_{h,R} \models \psi$. This, however, is a contradiction to (9). Altogether, the proof of Theorem 4.2 is complete. □

To conclude this section let us mention that an easy reduction shows that Theorem 4.2 and Theorem 4.3 still hold when replacing $\mathfrak{T}$ and $\mathfrak{F}_{\leq h}$ by the class $\mathfrak{T}^u$ and $\mathfrak{F}_{\leq h}^u$ of *undirected* trees, respectively, *undirected* forests of height at most h.

5 Lower bounds for the complexity of the theory of trees

The *first-order theory* $\mathrm{Th}(\mathfrak{C})$ of a class $\mathfrak{C}$ of structures is the set of all first-order sentences that hold in all structures in $\mathfrak{C}$. It is a classical topic of mathematical logic to determine the decidability of various theories, which led to seminal results such as the decidability of the theory of the field of real numbers [26] or the undecidability of arithmetic [28, 4]. It is known since at least the 1960s that the first-order theory of the class of trees is decidable. This is known for both the class of finite trees and the class of all trees. Furthermore, the decidability result remains true if the trees are labelled, and whether they are directed or undirected. For simplicity, we focus on the class $\mathfrak{T}$ of finite unlabelled directed trees in the following, but it is easy to see that our lower bound result also holds for all other variants. The only important assumption for our lower bound result below is that the trees are unranked and not of bounded degree. Actually, the decidability result even holds for the monadic second-order theory of trees [27] (for infinite trees, this is a deep result due to Rabin [22]). It is proved by translating monadic-second order sentences into tree-automata. An example of a class of finite structures with an undecidable monadic second-order theory and a decidable first-order theory is the class of finite initial segments of Presburger arithmetic. (That is, the class of all structures whose universe is $\{0,\ldots,n\}$ and that have one ternary relation, the graph of the addition restricted to $\{0,\ldots,n\}$. The decidability of the first-order theory of this class follows from the decidability of Presburger arithmetic [21], and the undecidability of monadic second-order logic from the fact that multiplication is monadic second-order definable from addition and the strong undecidability results known for arithmetic.) However, natural classes of graphs such as planar graphs or graphs of bounded degree, which are well-behaved with respect to first-order logic in other respects [2, 10], have an undecidable first-order theory. To see this, observe that we can easily use labelled grids to encode the sequence of configurations of a run of a Turing machine. The labels can be replaced by little gadgets, and the degree can be reduced from four to three by replacing grids by walls. This yields the following:

Fact 5.1. $\mathrm{Th}(\mathfrak{P}_3)$ *is undecidable, where* $\mathfrak{P}_3$ *denotes the class of all finite planar graphs of degree* ≤ 3.

Our main result in this section is a lower bound for the complexity of the theory of trees:

Theorem 5.2. *There is no elementary algorithm deciding* $\mathrm{Th}(\mathfrak{T})$.

Proof. We only sketch the proof. We prove that for every Turing machine M and every n there is a first-order sentence $\varphi_{M,n}$ of length polynomial in n and the size of the description of M such that $\varphi_{M,n}$ is satisfiable in $\mathfrak{T}$ if and only if M halts in less than *Tower*(n) steps when started with empty input. Clearly, this implies the desired lower bound.

Let M be a 1-tape Turing machine and $n \in \mathbb{N}$. Without loss of generality we may assume that the states of M and the tape symbols are elements of $[0, n{-}1]$.

We can encode a run of M of length at most ℓ by a set of 5-tuples (t, u, s, q, h), where $t, u, h \in [0, \ell{-}1]$, $s, q \in [0, n{-}1]$. Here t is a step of the computation, s is the symbol of the u-th tape cell in step t, q is the state, and h is the head position. Using standard coding techniques, we can encode each set of such 5-tuples, and hence each run of M of length at most ℓ, by a single integer less than 2^{2^ℓ}, provided ℓ is sufficiently large.

Using the encodings of natural numbers by trees introduced in Section 3, the first-order formulas defined there to speak about the trees, and standard techniques for describing Turing machine computations by logical formulas, it is not hard to construct a formula $\varphi_{M,n}$ with the following properties:

- If $\varphi_{M,n}$ is satisfiable, then (up to isomorphism) the only model of $\varphi_{M,n}$ is the tree $\mathcal{T}(m)$, where $m = $ *Tower*$(n{+}3){-}1$.
- $\varphi_{M,n}$ is satisfiable (i.e., holds in $\mathcal{T}(m)$) if and only if there is a number $j <$ *Tower*$(n{+}2)$ such that j is the encoding of an accepting run of the machine M of length at most *Tower*(n).
- The length of $\varphi_{M,n}$ is polynomial in n.

Observe that in $\mathcal{T}(m)$, where $m =$ *Tower*$(n{+}3){-}1$, we have all numbers $i \in [0,$ *Tower*$(n{+}2){-}1]$ "available", represented by the children of the root. Furthermore, the formulas less_{n+2} and bit_{n+2} from Lemma 3.4 give us sufficient arithmetic on these numbers. Thus in $\mathcal{T}(m)$ we have an initial segment of arithmetic of length *Tower*$(n{+}2)$ represented in such a way that we can access it by first-order formulas of length polynomial in n. This allows us to decode the run of the Turing machine represented by the number j. ☐

6 Failure of Feferman-Vaught theorems for formula size

The classical Feferman-Vaught theorem [8] 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 for $i = 1, 2$ the structures $\mathcal{A}_i$ and $\mathcal{B}_i$ 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 for $i = 1, 2$ the structures $\mathcal{A}_i$ and $\mathcal{B}_i$ 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 *Tower*$(\mathcal{O}(\ell))$ on the function f. Maybe surprisingly, this upper bound is essentially tight:

Theorem 6.1. *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 $\varphi_h := \forall x\Big(\mathit{encoding}_h(x) \to \big(\mathit{max}_h(x) \vee \exists y\ \mathit{succ}_h(x,y)\big)\Big)$.
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 *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. □

7 Existential preservation on forests — lower bounds for the size of formulas

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. A 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 also $\mathcal{B} \models \varphi$.

The well-known *Łoś-Tarski Theorem* (see e.g. [16]) 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.

It is known that the *Łoś-Tarski Theorem* fails when shifting the attention from the class of *all* structures to the class of all *finite* structures. I.e., there are first-order sentences that are preserved under extensions on the class of all finite structures, but not equivalent to any existential first-order sentence ([25, 15]).

On the other hand, [1] exposed "well-behaved" classes of finite structures for which a Łoś-Tarski like theorem holds. For example, it was shown that every first-order sentence φ that is preserved under extensions on the class of *finite acyclic structures* is equivalent, over this class, to an existential first-order sentence ψ. The proof given in [1] leads to an algorithm which on input φ produces a corresponding existential sentence ψ; the size of the resulting sentence ψ, however, may be non-elementarily larger than the size of the original sentence φ. The main result of the present section, Theorem 7.1, shows that this increase in formula size is not just an artifact of the proof given in [1], but that indeed the size of the shortest equivalent existential sentence may be non-elementarily larger than the size of the original formula φ.

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 7.1. *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 7.1, we therefore introduce the following encoding of numbers by $\{L, X\}$-labelled trees. The remainder of Section 7 is devoted to the proof of Theorem 7.1.

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 7.2. For each natural number $h \geq 1$ and each $n \in \{0, 1, .., \mathit{Tower}(h){-}1\}$, 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.

These trees are illustrated in Figure 3, which shows all the trees $\tilde{\mathcal{T}}_1(n)$ and $\tilde{\mathcal{T}}_2(n)$, and Figure 4 which shows $\tilde{\mathcal{T}}_3(5)$.

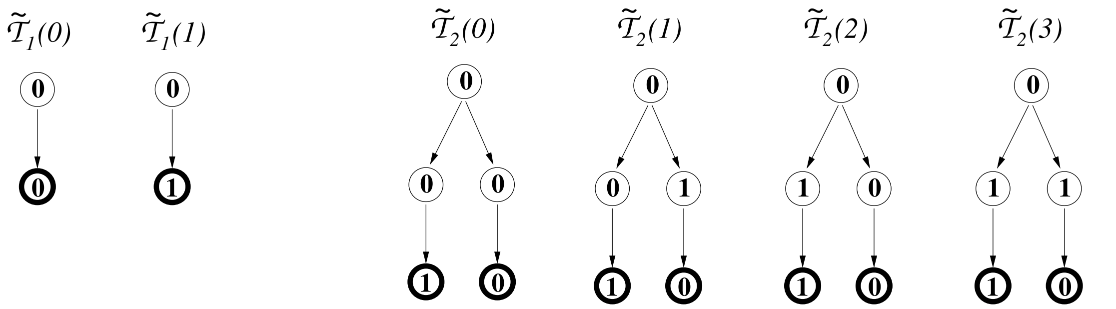

Figure 3: The trees $\tilde{\mathcal{T}}_1(0)$, $\tilde{\mathcal{T}}_1(1)$ and the trees $\tilde{\mathcal{T}}_2(0)$, $\tilde{\mathcal{T}}_2(1)$, $\tilde{\mathcal{T}}_2(2)$, $\tilde{\mathcal{T}}_2(3)$. Vertices in L are indicated by bold circles, vertices in X are labelled **1**, vertices outside X are labelled **0**.

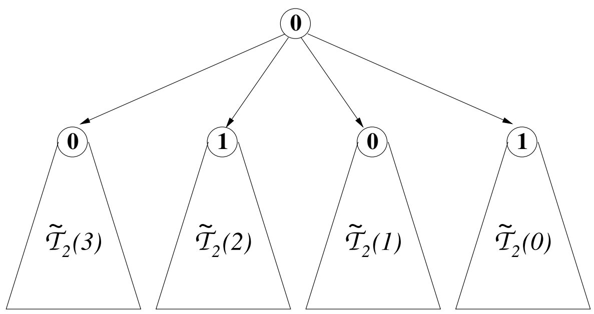

Figure 4: The tree $\tilde{\mathcal{T}}_3(5)$

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 will be given below. Prior to that, however, we need the following (easy) lemma.

Lemma 7.3 (The sentence $\mathit{forest}_{\leq h'}$). *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'.*

Proof. The formula $\mathit{forest}_{\leq h'}$ expresses the following:

1. There is no path of length $h'+1$, i.e.

$$\forall x_1\, \forall x_2 \cdots \forall x_{h'+1}\, \neg \bigwedge_{i=1}^{h'} E(x_i, x_{i+1})$$

   (note that the above formula also implies that there is no directed cycle, since the existence of a directed cycle implies that there are arbitrarily long paths).

2. Every node has in-degree at most 1, i.e.

$$\forall x\, \forall y\, \forall y' \Big(\big(E(y,x) \wedge E(y',x)\big) \rightarrow y = y' \Big)$$

   (note that along with 1. this implies that there is no "undirected cycle", i.e., no cycle in the undirected graph on vertex set A which, for each $u, v \in A$, has an undirected edge between u and v if, and only if, $(u,v) \in E^{\mathcal{A}}$ or $(v,u) \in E^{\mathcal{A}}$).

3. Every node that is labelled L has out-degree 0, i.e.

$$\forall x\, \forall y\, \Big(L(x) \rightarrow \neg E(x,y) \Big)$$

It should be straightforward to verify that the resulting formula $\mathit{forest}_{\leq h'}$ has the desired properties. ☐

Definition 7.4 (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 $\mathit{Rep}_1^{\mathcal{A}}(x) \in \{0,1\} = \{0, .., \mathit{Tower}(1){-}1\}$ as follows:

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

- 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 $i \in \{0, .., \mathit{Tower}(h){-}1\}$, $\mathrm{bit}(i,n) = 1 \iff X^{\mathcal{A}}(y_i)$.

The following notion of h-*inconsistency* can be viewed as a counterpart of 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 7.5 (h-inconsistency). 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:

- 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.

For example, let us consider the structure $\mathcal{A} := \tilde{T}_h(n)$ from Definition 7.2: The root t of $\mathcal{A}$ is the only node that is h-good in $\mathcal{A}$, and it represents the number $\mathit{Rep}_h^{\mathcal{A}}(t) = n$. Furthermore, each child of t is $(h{-}1)$-good in $\mathcal{A}$. Finally, there is no $\tilde{h} \leq h$ such that $\mathcal{A}$ is $\tilde{h}$-inconsistent, and thus $\mathcal{A}$ is *not* $(\leq h)$-inconsistent.

Lemma 7.6. *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 $\mathcal{A}$:* $\mathcal{A} \models \varphi_h \iff$

$$\Big(\mathcal{A} \models \neg \mathit{forest}_{\leq h}\Big) \text{ or } \Big(\mathcal{A} \text{ is } (\leq h)\text{-inconsistent}\Big) \text{ or } \Big(\text{there exists a node } x \text{ that is } h\text{-good in } \mathcal{A}\Big).$$

Before proving Lemma 7.6, we need the following two lemmas.

Lemma 7.7. *There is a constant $c \geq 1$ such that for every $h \geq 1$ there are* FO(τ)*-formulas $\mathit{good}_h(x)$, $\mathit{min}_h(x)$, $\mathit{max}_h(x)$, $\mathit{eq}_h(x)$, $\mathit{less}_h(x,y)$, and $\mathit{succ}_h(x,y)$, each of size at most $c \cdot 12^h$, such that the following is true: For every τ-structure $\mathcal{A}$ for which there is a $h' \geq h$ such that $\mathcal{A} \models \mathit{forest}_{\leq h'}$ and for all nodes $t, u \in A$ we have*

1. $\mathcal{A} \models \mathit{good}_h(t) \iff$ *t is h-good in $\mathcal{A}$*
2. $\mathcal{A} \models \mathit{min}_h(t) \iff$ *t is h-good in $\mathcal{A}$ and $\mathit{Rep}_h^{\mathcal{A}}(t) = 0$*
3. $\mathcal{A} \models \mathit{max}_h(t) \iff$ *t is h-good in $\mathcal{A}$ and $\mathit{Rep}_h^{\mathcal{A}}(t) = \mathit{Tower}(h){-}1$*
4. $\mathcal{A} \models \mathit{eq}_h(t,u) \iff$ *t and u are h-good in $\mathcal{A}$ and $\mathit{Rep}_h^{\mathcal{A}}(t) = \mathit{Rep}_h^{\mathcal{A}}(u)$*
5. $\mathcal{A} \models \mathit{less}_h(t,u) \iff$ *t and u are h-good in $\mathcal{A}$ and $\mathit{Rep}_h^{\mathcal{A}}(t) < \mathit{Rep}_h^{\mathcal{A}}(u)$*
6. $\mathcal{A} \models \mathit{succ}_h(t,u) \iff$ *t and u are h-good in $\mathcal{A}$ and $\mathit{Rep}_h^{\mathcal{A}}(t){+}1 = \mathit{Rep}_h^{\mathcal{A}}(u)$*

Proof. We define all these formulas simultaneously by induction on h. We start with $h = 1$ and define the following formulas:

1. $\mathit{good}_1(x) \ := \ \exists y\, \Big(E(x,y) \wedge L(y) \wedge \forall y' \big(\big(E(x,y') \wedge L(y')\big) \to \big(X(y') \leftrightarrow X(y)\big)\big)\Big)$
2. $\mathit{min}_1(x) \ := \ \mathit{good}_1(x) \wedge \exists y \big(E(x,y) \wedge L(y) \wedge \neg X(y)\big)$
3. $\mathit{max}_1(x) \ := \ \mathit{good}_1(x) \wedge \exists y \big(E(x,y) \wedge L(y) \wedge X(y)\big)$
4. $\mathit{eq}_1(x,y) \ := \ \mathit{good}_1(x) \wedge \mathit{good}_1(y) \wedge \big(\mathit{min}_1(x) \leftrightarrow \mathit{min}_1(y)\big)$
5. $\mathit{less}_1(x,y) \ := \ \mathit{good}_1(x) \wedge \mathit{good}_1(y) \wedge \big(\mathit{min}_1(x) \wedge \mathit{max}_1(y)\big)$
6. $\mathit{succ}_1(x,y) \ := \ \mathit{less}_1(x,y)\,.$

It is straightforward to see that all these formulas have the intended meaning. Let c be larger than the maximum size of these formulas. Then, each of the formulas has size $\leq c < c \cdot 12^1$. For the induction step from h to $h{+}1$ we choose the following formulas:

$$\begin{aligned} \mathit{good}_{h+1}(x) \ := \ & \exists y \Big(E(x,y) \wedge \mathit{min}_h(y) \Big) \ \wedge \ \exists y \Big(E(x,y) \wedge \mathit{max}_h(y) \Big) \ \wedge \\ & \forall y \Big(\big(E(x,y) \wedge \mathit{good}_h(y) \wedge \neg \mathit{max}_h(y) \big) \rightarrow \exists y' \big(E(x,y') \wedge \mathit{succ}_h(y,y') \big) \Big) \ \wedge \\ & \forall y \, \forall z \, \Big(\big(E(x,y) \wedge E(x,z) \wedge \mathit{eq}_h(y,z) \big) \rightarrow \big(X(z) \leftrightarrow X(y) \big) \Big) . \end{aligned}$$

The first two lines of this formula express that for every $i \in \{0,..,\mathit{Tower}(h){-}1\}$, x has a h-good child y_i with $\mathit{Rep}_h^{\mathcal{A}}(y_i) = i$. The last line expresses that for all h-good children y, z of x with $\mathit{Rep}_h^{\mathcal{A}}(y) = \mathit{Rep}_h^{\mathcal{A}}(z)$ the following is true: $X^{\mathcal{A}}(z) \leftrightarrow X^{\mathcal{A}}(y)$. Thus, the formula $\mathit{good}_{h+1}(x)$ expresses that x is $(h{+}1)$-good in $\mathcal{A}$. Furthermore,

$$\begin{aligned} ||\mathit{good}_{h+1}|| \ \leq \ & c + ||\mathit{min}_h|| + 2 \cdot ||\mathit{max}_h|| + ||\mathit{good}_h|| + ||\mathit{succ}_h|| + ||\mathit{eq}_h|| \\ \leq \ & c + 6 \cdot c \cdot 12^h \ \leq \ 7 \cdot c \cdot 12^h \ < \ c \cdot 12^{h+1} . \end{aligned}$$

We note for further use that $||\mathit{good}_{h+1}|| \ \leq \ 7 \cdot c \cdot 12^h$.

Next we choose

$$\mathit{min}_{h+1}(x) \ := \ \mathit{good}_{h+1}(x) \ \wedge \ \forall y \Big(\big(E(x,y) \wedge \mathit{good}_h(y) \big) \rightarrow \neg X(y) \Big) .$$

This formula expresses that x is $(h{+}1)$-good in $\mathcal{A}$ and that for $n := \mathit{Rep}_{h+1}^{\mathcal{A}}(x)$ we have $\mathrm{bit}(i,n) = 0$ for all $i \in \{0,..,\mathit{Tower}(h){-}1\}$, i.e., $n = 0$. Furthermore,

$$||\mathit{min}_{h+1}|| \ \leq \ c + ||\mathit{good}_{h+1}|| + ||\mathit{good}_h|| \ \leq \ c + 7 \cdot c \cdot 12^h + c \cdot 12^h \ \leq \ c \cdot 12^{h+1} .$$

The formula $\mathit{max}_{h+1}(x)$ is chosen analogously:

$$\mathit{max}_{h+1}(x) \ := \ \mathit{good}_{h+1}(x) \ \wedge \ \forall y \Big(\big(E(x,y) \wedge \mathit{good}_h(y) \big) \rightarrow X(y) \Big) .$$

This formula expresses that x is $(h{+}1)$-good in $\mathcal{A}$ and that for $n := \mathit{Rep}_{h+1}^{\mathcal{A}}(x)$ we have $\mathrm{bit}(i,n) = 1$ for all $i \in \{0,..,\mathit{Tower}(h){-}1\}$, i.e., $n = \mathit{Tower}(h{+}1){-}1$. Furthermore,

$$||\mathit{max}_{h+1}|| \ \leq \ c + ||\mathit{good}_{h+1}|| + ||\mathit{good}_h|| \ \leq \ c + 7 \cdot c \cdot 12^h + c \cdot 12^h \ \leq \ c \cdot 12^{h+1} .$$

The formula $\mathit{eq}_{h+1}(x,y)$ is chosen as follows:

$$\begin{aligned} \mathit{eq}_{h+1}(x,y) \ := \ & \forall u \Big(\big(u{=}x \ \vee \ u{=}y \big) \rightarrow \mathit{good}_{h+1}(u) \Big) \wedge \\ & \forall x' \, \forall y' \, \Big(\big(E(x,x') \wedge E(y,y') \wedge \mathit{eq}_h(x',y') \big) \rightarrow \big(X(x') \leftrightarrow X(y') \big) \Big) . \end{aligned}$$

The first line of this formula expresses that x and y are $(h{+}1)$-good in $\mathcal{A}$. The second line expresses that $\mathit{Rep}_{h+1}^{\mathcal{A}}(x) = \mathit{Rep}_{h+1}^{\mathcal{A}}(y)$. Furthermore,

$$||\mathit{eq}_{h+1}|| \ \leq \ c + ||\mathit{good}_{h+1}|| + ||\mathit{eq}_h|| \ \leq \ c + 7 \cdot c \cdot 12^h + c \cdot 12^h \ \leq \ c \cdot 12^{h+1} .$$

The formula $\mathit{less}_{h+1}(x,y)$ is chosen as follows:

$$\begin{aligned} & \mathit{less}_{h+1}(x,y) \ := \ \forall u \Big(\big(u{=}x \ \vee \ u{=}y \big) \rightarrow \mathit{good}_{h+1}(u) \Big) \ \wedge \\ & \quad \exists x' \, \exists y' \, \Big(E(x,x') \wedge E(y,y') \wedge \mathit{eq}_h(x',y') \wedge \neg X(x') \wedge X(y') \wedge \\ & \qquad \forall x'' \, \forall y'' \, \Big(\big(E(x,x'') \wedge E(y,y'') \wedge \mathit{eq}_h(x'',y'') \wedge \mathit{less}_h(x',x'') \big) \rightarrow \big(X(x'') \leftrightarrow X(y'') \big) \Big) \Big) . \end{aligned}$$

The first line of this formula expresses that x and y are $(h{+}1)$-good in $\mathcal{A}$. The remaining lines express that there exist h-good children x', y' of x, y with the following properties: $\mathit{Rep}_h^{\mathcal{A}}(x') = \mathit{Rep}_h^{\mathcal{A}}(y') =: i$ such that for $m := \mathit{Rep}_{h+1}^{\mathcal{A}}(x)$ and $n := \mathit{Rep}_{h+1}^{\mathcal{A}}(y)$ we have $\mathrm{bit}(i,m) = 0$, $\mathrm{bit}(i,n) = 1$, and for all $j > i$ we have $\mathrm{bit}(j,m) = \mathrm{bit}(j,n)$. Thus, the formula $\mathit{less}_{h+1}(x,y)$ expresses that $m < n$. Furthermore,

$$\begin{aligned} ||\mathit{less}_{h+1}|| &\leq c + ||\mathit{good}_{h+1}|| + 2\cdot||\mathit{eq}_h|| + ||\mathit{less}_h|| \\ &\leq c + 7\cdot c\cdot 12^h + 3\cdot c\cdot 12^h \;\leq\; c + 10\cdot c\cdot 12^h \;\leq\; c\cdot 12^{h+1}\,. \end{aligned}$$

Finally, the formula $\mathit{succ}_{h+1}(x,y)$ is chosen as follows:

$$\begin{aligned} \mathit{succ}_{h+1}(x,y) \;:=\;& \forall u\,\Big(\big(u{=}x \,\vee\, u{=}y\big) \to \mathit{good}_{h+1}(u)\Big) \;\wedge \\ & \exists x'\exists y'\,\Big(E(x,x') \wedge E(y,y') \wedge \mathit{eq}_h(x',y') \wedge \neg X(x') \wedge X(y') \,\wedge \\ & \quad \forall x''\forall y''\,\Big(\big(E(x,x'') \wedge E(y,y'') \wedge \mathit{eq}_h(x'',y'')\big) \to \\ & \qquad \Big(\big(\mathit{less}_h(x',x'') \to (X(x'') \leftrightarrow X(y''))\big) \wedge \big(\mathit{less}_h(x'',x') \to (X(x'') \wedge \neg X(y''))\big)\Big)\Big)\Big)\,. \end{aligned}$$

The first line of this formula expresses that x and y are $(h{+}1)$-good in $\mathcal{A}$. The remaining lines express that there exist h-good children x', y' of x, y with the following properties: $\mathit{Rep}_h^{\mathcal{A}}(x') = \mathit{Rep}_h^{\mathcal{A}}(y') =: i$ such that for $m := \mathit{Rep}_{h+1}^{\mathcal{A}}(x)$ and $n := \mathit{Rep}_{h+1}^{\mathcal{A}}(y)$ we have $\mathrm{bit}(i,m) = 0$, $\mathrm{bit}(i,n) = 1$, and for all $j > i$ we have $\mathrm{bit}(j,m) = \mathrm{bit}(j,n)$, and for all $j < i$ we have $\mathrm{bit}(j,m) = 1$ and $\mathrm{bit}(j,n) = 0$. Thus, the formula $\mathit{inc}_{h+1}(x,y)$ expresses that $m{+}1 = n$. Furthermore,

$$\begin{aligned} ||\mathit{succ}_{h+1}|| &\leq c + ||\mathit{good}_{h+1}|| + 2\cdot||\mathit{eq}_h|| + 2\cdot||\mathit{less}_h|| \\ &\leq c + 7\cdot c\cdot 12^h + 4\cdot c\cdot 12^h \;\leq\; c + 11\cdot c\cdot 12^h \;\leq\; c\cdot 12^{h+1}\,. \end{aligned}$$

This finally completes the proof of Lemma 7.7. □

Lemma 7.8. *There is a constant $d \geq 1$ such that for every $h \geq 1$ there are* FO(τ)*-sentences $\mathit{inconsistent}_h$ and $\mathit{inconsistent}_{\leq h}$, each of size $< d\cdot 12^h$, such that the following is true: For every τ-structure $\mathcal{A}$ for which there is a $h' \geq h$ such that $\mathcal{A} \models \mathit{forest}_{\leq h'}$ we have*

- *$\mathcal{A} \models \mathit{inconsistent}_h \iff \mathcal{A}$ is h-inconsistent*
- *$\mathcal{A} \models \mathit{inconsistent}_{\leq h} \iff \mathcal{A}$ is $(\leq h)$-inconsistent.*

Proof. For $h = 1$ we choose

$$\mathit{inconsistent}_1 \;:=\; \exists x\exists y\exists y'\,\Big(E(x,y) \wedge E(x,y') \wedge L(y) \wedge L(y') \wedge X(y) \wedge \neg X(y')\Big)\,.$$

Along Definition 7.5 one immediately sees that $\mathcal{A} \models \mathit{inconsistent}_h \iff \mathcal{A}$ is h-inconsistent.

For $h \geq 1$ we use Lemma 7.7 and choose

$$\mathit{inconsistent}_{h+1} \;:=\; \exists x\exists y\exists y'\,\Big(E(x,y) \wedge E(x,y') \wedge \mathit{eq}_h(y,y') \wedge X(y) \wedge \neg X(y')\Big)\,.$$

Using Lemma 7.7 one obtains that this formula expresses that there exist nodes x, y, y' such that y and y' are h-good children of x with $\mathit{Rep}_h^{\mathcal{A}}(y) = \mathit{Rep}_h^{\mathcal{A}}(y')$ and $X^{\mathcal{A}}(y)$ and $\neg X^{\mathcal{A}}(y')$. Thus, the formula $\mathit{inconsistent}_{h+1}$ expresses that $\mathcal{A}$ is $(h{+}1)$-inconsistent. Since $||\mathit{eq}_h|| \leq c\cdot 12^h$, we obtain that for a suitable constant $d \geq 1$ and for all $h \geq 1$ we have $||\mathit{inconsistent}_h|| < d\cdot 12^{h-1}$.

Finally, for each $h \geq 1$ we choose

$$\mathit{inconsistent}_{\leq h} \;:=\; \bigvee_{\tilde{h}=1}^{h} \mathit{inconsistent}_{\tilde{h}}\,.$$

It should be clear that this formula expresses that $\mathcal{A}$ is $(\leq h)$-inconsistent. Furthermore,

$$||\mathit{inconsistent}_{\leq h}|| \;\leq\; \sum_{\tilde{h}=1}^{h} d \cdot 12^{\tilde{h}-1} \;<\; d \cdot 12^{h}\,.$$

This completes the proof of Lemma 7.8. □

Finally, we are ready for the

Proof of Lemma 7.6:
We use Lemma 7.3, Lemma 7.7, and Lemma 7.8 and choose for every $h \geq 1$

$$\varphi_h \;:=\; \neg \mathit{forest}_{\leq h} \;\vee\; \mathit{inconsistent}_{\leq h} \;\vee\; \exists x\, \mathit{good}_h(x)\,.$$

We know that $||\mathit{forest}_{\leq h}|| = \mathcal{O}(h)$, $||\mathit{inconsistent}_{\leq h}|| \leq d{\cdot}12^h$, and $||\mathit{good}_h|| \leq c{\cdot}12^h$. Thus, $||\varphi_h|| = 2^{\mathcal{O}(h)}$.

Furthermore, for each τ-structure $\mathcal{A}$ we have $\mathcal{A} \models \varphi_h$ if and only if

$$\Big(\mathcal{A} \models \neg \mathit{forest}_{\leq h}\Big) \text{ or } \Big(\mathcal{A} \text{ is } (\leq h)\text{-inconsistent}\Big) \text{ or } \Big(\text{there exists a node } x \text{ that is } h\text{-good in } \mathcal{A}\Big).$$

Thus, the proof of Lemma 7.6 is complete. □

The next two lemmas will enable us to show that this sentence φ_h is preserved under extensions.

Lemma 7.9. *Let $h' \geq 1$ and let $\mathcal{A}$ be a τ-structure with $\mathcal{A} \models \mathit{forest}_{\leq h'}$. Let $\mathcal{B}$ be an extension of $\mathcal{A}$ with $\mathcal{B} \models \mathit{forest}_{\leq h'}$. Let $h \in \{1,..,h'\}$, and let x be a node in $\mathcal{A}$ that is h-good in $\mathcal{A}$. Then, at least one of the following statements is true:*

1. *x is h-good in $\mathcal{B}$ and $\mathit{Rep}_h^{\mathcal{B}}(x) = \mathit{Rep}_h^{\mathcal{A}}(x)$.*
2. *$\mathcal{B}$ is $(\leq h)$-inconsistent.*

Proof. By induction on h.

For $h = 1$ we know that x is 1-good in $\mathcal{A}$. I.e., x has a child y in $\mathcal{A}$ with $L^{\mathcal{A}}(y)$, and for all children y' of x in $\mathcal{A}$ with $L^{\mathcal{A}}(y')$ we have $X^{\mathcal{A}}(y') \leftrightarrow X^{\mathcal{A}}(y)$.

If x is 1-good in $\mathcal{B}$, then we must have $\mathit{Rep}_1^{\mathcal{B}}(x) = \mathit{Rep}_1^{\mathcal{A}}(x)$, because the node y is a child of x in $\mathcal{B}$ with $L^{\mathcal{B}}(y)$, and for all children y'' of x in $\mathcal{B}$ with $L^{\mathcal{B}}(y'')$ we have $X^{\mathcal{B}}(y') \leftrightarrow X^{\mathcal{B}}(y) \leftrightarrow X^{\mathcal{A}}(y)$.

On the other hand, if x is *not* 1-good in $\mathcal{B}$, then we know that in $\mathcal{B}$ the node x must have (at least) one further child z with $L^{\mathcal{B}}(z)$ and $X^{\mathcal{B}}(z) \leftrightarrow \neg X^{\mathcal{B}}(y)$. Then, however, $\mathcal{B}$ is 1-inconsistent.

Let us now consider the induction step from h to $h{+}1$. We know that x is $(h{+}1)$-good in $\mathcal{A}$, i.e., for every $i \in \{0,..,\mathit{Tower}(h){-}1\}$ there exists a child y_i of x in $\mathcal{A}$ that is h-good in $\mathcal{A}$ with $\mathit{Rep}_h^{\mathcal{A}}(y_i) = i$. Applying the induction hypothesis to each of the nodes y_i we obtain that at least one of the following statements is true:

1. $\mathcal{B}$ is $(\leq h)$-inconsistent.
2. y_i is h-good in $\mathcal{B}$ and $\mathit{Rep}_h^{\mathcal{B}}(y_i) = \mathit{Rep}_h^{\mathcal{A}}(y_i) = i$.

In case that $\mathcal{B}$ is $(\leq h)$-inconsistent, it is also $(\leq h{+}1)$-inconsistent, and we are done. There remains to consider the case where each y_i is h-good in $\mathcal{B}$ and $\mathit{Rep}_h^{\mathcal{B}}(y_i) = i$.

Now let z be an arbitrary child of x in $\mathcal{B}$ that is h-good in $\mathcal{B}$. Let $i_z := \mathit{Rep}_h^{\mathcal{B}}(z)$. If there is (at least) one such z with $X^{\mathcal{B}}(z) \leftrightarrow \neg X^{\mathcal{B}}(y_{i_z})$, then $\mathcal{B}$ is $(h{+}1)$-inconsistent and thus also $\leq(h{+}1)$-inconsistent. On the other hand, if all such z satisfy $X^{\mathcal{B}}(z) \leftrightarrow X^{\mathcal{B}}(y_{i_z})$, then x is $(h{+}1)$-good in $\mathcal{B}$. □

Lemma 7.10. *Let $h' \geq 1$ and let $\mathcal{A}$ be a τ-structure with $\mathcal{A} \models \mathit{forest}_{\leq h'}$. Let $\mathcal{B}$ be an extension of $\mathcal{A}$ with $\mathcal{B} \models \mathit{forest}_{\leq h'}$. Let $h \in \{1,..,h'\}$. Then the following is true: If $\mathcal{A}$ is $(\leq h)$-inconsistent then also $\mathcal{B}$ is $(\leq h)$-inconsistent.*

Proof. Let $\mathcal{A}$ be $(\leq h)$-inconsistent. Then there is a $\tilde{h} \in \{1, .., h\}$ such that $\mathcal{A}$ is $\tilde{h}$-inconsistent.

In case that $\tilde{h} = 1$, there thus are nodes x, y, y' in $\mathcal{A}$ such that y, y' are children of x in $\mathcal{A}$ which both belong to $L^{\mathcal{A}}$ such that $X^{\mathcal{A}}(y)$ and $\neg X^{\mathcal{A}}(y')$. Since $\mathcal{B}$ is an extension of $\mathcal{A}$, we know that y, y' are children of x in $\mathcal{B}$ which both belong to $L^{\mathcal{B}}$ such that $X^{\mathcal{B}}(y)$ and $\neg X^{\mathcal{B}}(y')$. Thus $\mathcal{B}$ is 1-inconsistent and hence also $(\leq h)$-inconsistent.

In case that $\tilde{h} > 1$, we know that there are nodes x, y, y' in $\mathcal{A}$ such that y, y' are children of x in $\mathcal{A}$ which have the following properties: y and y' both are $(\tilde{h}{-}1)$-good in $\mathcal{A}$ with $i := \textit{Rep}^{\mathcal{A}}_{\tilde{h}-1}(y) = \textit{Rep}^{\mathcal{A}}_{\tilde{h}-1}(y')$, and we have $X^{\mathcal{A}}(y)$ and $\neg X^{\mathcal{A}}(y')$. From Lemma 7.9 we obtain that at least one of the following statements is true:

1. y and y' are $(\tilde{h}{-}1)$-good in $\mathcal{B}$ and $i = \textit{Rep}^{\mathcal{B}}_{\tilde{h}-1}(y) = \textit{Rep}^{\mathcal{B}}_{\tilde{h}-1}(y')$.

2. $\mathcal{B}$ is $(\leq \tilde{h}{-}1)$-inconsistent.

In the second case we know that $\mathcal{B}$ is $(\leq h)$-inconsistent and thus we are done. In the first case we know that x, y, y' are nodes in $\mathcal{B}$ such that y, y' are children of x in $\mathcal{B}$ which have the following properties: y and y' both are $(\tilde{h}{-}1)$-good in $\mathcal{B}$ with $\textit{Rep}^{\mathcal{B}}_{\tilde{h}-1}(y) = \textit{Rep}^{\mathcal{B}}_{\tilde{h}-1}(y')$, and we have $X^{\mathcal{B}}(y)$ and $\neg X^{\mathcal{B}}(y')$. Thus, the nodes x, y, y' witness that $\mathcal{B}$ is $\tilde{h}$-inconsistent and thus also $(\leq h)$-inconsistent. □

We are now ready to prove Theorem 7.1. **Proof of Theorem 7.1:**
For each $h \geq 1$ let φ_h be the FO(τ)-sentence of length $2^{\mathcal{O}(h)}$ obtained from Lemma 7.6.
By applying Lemma 7.10 and Lemma 7.9 one easily obtains the following:

Claim 7.11. *φ_h is preserved under extensions on the class of all τ-structures.*

Proof. We have to show that φ_h is preserved under extensions on the class of all τ-structures. To this end let $\mathcal{A}$ be a τ-structure with $\mathcal{A} \models \varphi_h$, and let $\mathcal{B}$ be an extension of $\mathcal{A}$. Our aim is to show that $\mathcal{B} \models \varphi_h$.
Since $\mathcal{A} \models \varphi_h$, we know from Lemma 7.6 that

$$\Big(\mathcal{A} \models \neg \textit{forest}_{\leq h}\Big) \ \text{ or } \ \Big(\mathcal{A} \text{ is } (\leq h)\text{-inconsistent}\Big) \ \text{ or } \ \Big(\text{there exists a node } x \text{ that is } h\text{-good in } \mathcal{A}\Big).$$

We can thus consider three cases:

Case 1: $\mathcal{A} \models \neg \textit{forest}_{\leq h}$.
From Lemma 7.3 we know that $\textit{forest}_{\leq h}$ is a *universal* FO-formula. Hence, $\neg\textit{forest}_{\leq h}$ is an *existential* FO-formula. In particular, this implies that the formula $\neg\textit{forest}_{\leq h}$ is preserved under extensions. We thus know that $\mathcal{B} \models \neg\textit{forest}_{\leq h}$ and hence, also $\mathcal{B} \models \varphi_h$.

Case 2: $\mathcal{A} \models \textit{forest}_{\leq h}$, and $\mathcal{A}$ is $(\leq h)$-inconsistent.
If $\mathcal{B} \models \neg\textit{forest}_{\leq h}$, then certainly $\mathcal{B} \models \varphi_h$. On the other hand, if $\mathcal{B} \models \textit{forest}_{\leq h}$, then Lemma 7.10 tells us that also $\mathcal{B}$ is $(\leq h)$-inconsistent and hence $\mathcal{B} \models \varphi_h$.

Case 3: $\mathcal{A} \models \textit{forest}_{\leq h}$, and $\mathcal{A}$ is *not* $(\leq h)$-inconsistent.
Then we know that there exists a node x in $\mathcal{A}$ that is h-good in $\mathcal{A}$. If $\mathcal{B} \models \neg\textit{forest}_{\leq h}$, then certainly $\mathcal{B} \models \varphi_h$. On the other hand, if $\mathcal{B} \models \textit{forest}_{\leq h}$, then we know from Lemma 7.9 that x is h-good in $\mathcal{B}$ or $\mathcal{B}$ is $(\leq h)$-inconsistent. Thus, $\mathcal{B} \models \varphi_h$.

In summary, we hence obtain that φ_h is preserved under extensions on the class of all τ-structures. □

It remains to show that every existential FO(τ)-sentence ψ that is equivalent to φ_h on the class $\mathfrak{T}_{\leq h}$ is of size at least $\textit{Tower}(h{-}1)$. For contradiction, let us assume that φ_h is equivalent on $\mathfrak{T}_{\leq h}$ to an existential FO-sentence ψ of size $< \textit{Tower}(h{-}1)$. Then, ψ_h has a number $k < \textit{Tower}(h{-}1)$ of existential quantifiers and is w.l.o.g. of the form

$$\exists z_1 \, \exists z_2 \cdots \exists z_k \;\; \chi(z_1, .., z_k)$$

where χ is a Boolean combination of atomic τ-formulas. (For this, note that transforming an *existential* first-order formula into prenex normal form does not essentially increase the size of the formula.)

Consider the structure $\mathcal{T} := \tilde{\mathcal{T}}_h(0)$ from Definition 7.2 and recall that the root node t of $\mathcal{T}$ is h-good in $\mathcal{T}$. Thus, $\mathcal{T} \models \varphi_h$. Furthermore, $\mathcal{T}$ belongs to $\mathfrak{T}_{\leq h}$, and hence $\mathcal{T} \models \psi$. Therefore, there are nodes $a_1, .., a_k$ in $\mathcal{T}$ such that $\mathcal{T} \models \chi(a_1, .., a_k)$.

Since the root of $\mathcal{T}$ has $Tower(h-1)$ children, and $k < Tower(h-1)$, there must be (at least) one child u of the root t of $\mathcal{T}$ such that none of the nodes $a_1, .., a_k$ belongs to the subtree $\mathcal{T}_u$. Due to the definition of $\mathcal{T} = \tilde{\mathcal{T}}_h(0)$ we know that there must be an $i \in \{0, .., Tower(h-1)-1\}$ such that $\mathcal{T}_u$ corresponds to $\tilde{\mathcal{T}}_{h-1}(i)$. Along Definition 7.4 it is straightforward to see that u is $(h-1)$-good in $\mathcal{T}$ and $Rep^{\mathcal{T}}_{h-1}(u) = i$.

Let $\mathcal{T}^{-i}$ be the structure obtained from $\mathcal{T}$ by deleting the entire subtree $\mathcal{T}_u$. Then, each of the nodes $a_1, .., a_k$ belong to $\mathcal{T}^{-i}$ and thus, $\mathcal{T}^{-i} \models \chi(a_1, .., a_k)$ (because χ is a Boolean combination of atomic formulas). We hence have that $\mathcal{T}^{-i} \models \psi$. On the other hand, it is not difficult to see the following:

Claim 7.12. $\mathcal{T}^{-i} \not\models \varphi_h$.

Proof. First of all, $\mathcal{T}^{-i} \models forest_{\leq h}$, because $\mathcal{T}^{-i}$ is an induced substructure of $\mathcal{T}$ (recall that $forest_{\leq h}$ is a *universal* formula and thus preserved under induced substructures).
Furthermore, along Definition 7.5 it is straightforward to see that $\mathcal{T}^{-i}$ is not $(\leq h)$-inconsistent.
Finally, there is no h-good node in $\mathcal{T}^{-i}$: To see this, note that due to the height of nodes, the root of $\mathcal{T}^{-i}$ is the only candidate for a h-good node. This root, however, does not have a $(h-1)$-good child v with $Rep^{\mathcal{T}^{-i}}_{h-1}(v) = i$, and thus the root of $\mathcal{T}^{-i}$ cannot be h-good. In summary, we thus obtain that $\mathcal{T}^{-i} \not\models \varphi_h$. ☐

Altogether, we now know that $\mathcal{T}^{-i} \models \psi$ but $\mathcal{T}^{-i} \not\models \varphi_h$, contradicting our assumption that ψ is equivalent to φ_h on $\mathfrak{T}_{\leq h}$. Thus, the proof of Theorem 7.1 finally is complete. ☐

8 Structures of bounded degree — elementary upper bounds

All the non-elementary lower bounds in previous sections depended heavily on the fact that we considered 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 of the theorems 4.2, 5.2, 6.1, and 7.1. 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 [7], one obtains the following elementary *upper* bound, which we set out to prove next.

Lemma 8.1. *Let τ be a finite relational vocabulary, let $d \in \mathbb{N}$, and let $\mathfrak{D}_d$ be the class of all τ-structures whose Gaifman graph has degree $\leq d$. Then there is a 2-fold exponential function $f : \mathbb{N} \to \mathbb{N}$ such that for all structures $\mathcal{A}, \mathcal{B} \in \mathfrak{D}_d$ and all $m \in \mathbb{N}$ the following is true: If $\mathcal{A}$ and $\mathcal{B}$ satisfy the same basic local* FO(τ)*-sentences of length $\leq f(m)$, then $\mathcal{A} \equiv_m \mathcal{B}$, i.e., $\mathcal{A}$ and $\mathcal{B}$ satisfy the same* FO(τ)*-sentences of quantifier rank $\leq m$.*

Proof. The proof is an adaption of the proof of Lemma 2.5.2 in the textbook [7]. In fact, the present proof is formulated in such a way that it will be convenient for the reader to read this proof and the original proof in [7] in parallel. We use exactly the same notation as [7]. In contrast to the proof in [7], however, we use the following:

(i) $\cong$ instead of $\equiv_{g(j)}$. I.e., we let I_j be the set of all partial isomorphisms $\overline{a} \mapsto \overline{b}$ from $\mathcal{A}$ to $\mathcal{B}$ such that length($\overline{a}$) = length($\overline{b}$) $\leq m-j$ and

$$(\mathcal{S}^{\mathcal{A}}(7^j, \overline{a}), \overline{a}) \;\cong\; (\mathcal{S}^{\mathcal{B}}(7^j, \overline{b}), \overline{b})$$

(The notation of [7] is as follows: $\mathcal{S}^{\mathcal{A}}(7^j, \overline{a})$ denotes the induced substructure of $\mathcal{A}$ on the 7^j-neighbourhood of $\overline{a}$. Furthermore, $\mathcal{A} \cong \mathcal{B}$ means that $\mathcal{A}$ is isomorphic to $\mathcal{B}$.)

(ii) For every $a \in A$ and every $j \leq m$ we let $\psi^j_a(x)$ be a FO(τ)-formula which specifies the isomorphism type of the 7^j-neighbourhood of a. Precisely, we want that for every structure $\mathcal{C} \in \mathfrak{D}_d$ and every $c \in C$ we have

$$\mathcal{C} \models \psi^j_a(c) \iff (\mathcal{S}^{\mathcal{C}}(7^j, c), c) \cong (\mathcal{S}^{\mathcal{A}}(7^j, a), a)$$

Note that since $\mathcal{A} \in \mathfrak{D}_d$, the 7^j-neighbourhood of a consists of at most $1+d+d^2+\cdots+d^{7^j} \leq (d+1)^{7^j}$ vertices. Thus, the isomorphism type of the 7^j-neighbourhood of a can be described by a FO(τ)-formula $\psi_a^j(x)$ of size polynomial in $(d+1)^{7^j}$. Since d is fixed and j never gets larger than m, we thus have that there is a 2-fold exponential function $\tilde{f} : \mathbb{N} \to \mathbb{N}$ such that $||\psi_a^j(x)|| \leq \tilde{f}(m)$, for all $a \in A$ and all $j \leq m$.

We start with a number $m \in \mathbb{N}$ and two structures $\mathcal{A}, \mathcal{B} \in \mathfrak{D}_d$ which satisfy the same basic local sentences of size $\leq f(m)$, where f is a suitable 2-fold exponential function (the precise choice of which will be determined later).

Following the proof in [7], we suppose $0 \leq j < m$, and $\overline{a} \mapsto \overline{b} \in I_{j+1}$. Due to our particular choice of I_{j+1} (as specified in item (i) above), instead of equation (1) from [7], in our setting the assumption $\overline{a} \mapsto \overline{b} \in I_{j+1}$ means that

$$(1) \qquad (\mathcal{S}^{\mathcal{A}}(7^{j+1}, \overline{a}), \overline{a}) \;\cong\; (\mathcal{S}^{\mathcal{B}}(7^{j+1}, \overline{b}), \overline{b})$$

Our aim now is to show that for every $a \in A$ there exists a $b \in B$ such that $\overline{a}a \mapsto \overline{b}b \in I_j$, i.e.,

$$(\mathcal{S}^{\mathcal{A}}(7^j, \overline{a}a), \overline{a}a) \;\cong\; (\mathcal{S}^{\mathcal{B}}(7^j, \overline{b}b), \overline{b}b).$$

We follow the case distinction of [7]:

Case 1. $a \in S^{\mathcal{A}}(2 \cdot 7^j, \overline{a})$.
Then, of course, $S(7^j, a) \subseteq S(7^{j+1}, \overline{a})$. Due to (1) we therefore know that there must be a $b \in S^{\mathcal{B}}(2 \cdot 7^j, \overline{b})$ such that $(\mathcal{S}^{\mathcal{A}}(7^j, \overline{a}a), \overline{a}a) \cong (\mathcal{S}^{\mathcal{B}}(7^j, \overline{b}b), \overline{b}b)$.

Case 2. $a \notin S^{\mathcal{A}}(2 \cdot 7^j, \overline{a})$, i.e., $S^{\mathcal{A}}(7^j, \overline{a}) \cap S^{\mathcal{A}}(7^j, a) = \emptyset$.
Choose $e \in \mathbb{N}$ to be the *largest possible* number for which the following is true:

($*$) there are e elements $x_1, .., x_e$ in $S^{\mathcal{A}}(2 \cdot 7^j, \overline{a})$ of pairwise distance $> 4 \cdot 7^j$ whose 7^j-neighbourhoods (in $\mathcal{A}$) are all isomorphic to $(\mathcal{S}^{\mathcal{A}}(7^j, a), a)$.

Note that e must be of size $e \leq |\overline{a}| \leq m - (j+1) \leq m{+}1$.

From (1) we immediately obtain that this e is also the *largest possible* number for which

($**$) there are e elements $x_1, .., x_e$ in $S^{\mathcal{B}}(2 \cdot 7^j, \overline{b})$ of pairwise distance $> 4 \cdot 7^j$ whose 7^j-neighbourhoods (in $\mathcal{B}$) are all isomorphic to $(\mathcal{S}^{\mathcal{A}}(7^j, a), a)$.

Now consider the following basic local sentence

$$\chi \;:=\; \exists x_1 \cdots \exists x_e \exists x_{e+1} \Big(\bigwedge_{1 \leq \ell < k \leq e+1} \mathit{dist}(x_\ell, x_k) > 4 \cdot 7^j \;\;\wedge\; \bigwedge_{\ell=1}^{e+1} \psi_a^j(x_\ell) \Big)$$

I.e., χ expresses that there are $e{+}1$ elements of pairwise distance $> 4 \cdot 7^j$ whose 7^j-neighbourhoods are all isomorphic to $(\mathcal{S}^{\mathcal{A}}(7^j, a), a)$.

Since $e, j \leq m{+}1$ and since $||\psi_a^j(x)|| \leq \tilde{f}(m)$, for a 2-fold exponential function $\tilde{f}$ (cf., item (ii) at the beginning of this proof), we know that $||\chi|| \leq f(m)$, for a suitable 2-fold exponential function $f : \mathbb{N} \to \mathbb{N}$.

From the lemma's assumption we know that $\mathcal{A}$ and $\mathcal{B}$ satisfy the same basic local sentences of length $\leq f(m)$. Thus we have

$$\mathcal{A} \models \chi \iff \mathcal{B} \models \chi$$

Case 2.1. $\mathcal{A} \not\models \chi$ (and $\mathcal{B} \not\models \chi$).
Then, due to ($*$), *all* elements a' in $\mathcal{A}$ whose 7^j-neighbourhood is isomorphic to $(\mathcal{S}^{\mathcal{A}}(7^j, a), a)$ must belong to $S^{\mathcal{A}}(6 \cdot 7^j, \overline{a})$. (To see this, note that if a' was an element in $\mathcal{A}$ which is outside $S^{\mathcal{A}}(6 \cdot 7^j, \overline{a})$, then this a' together with the e elements from ($*$) would witness that $\mathcal{A} \models \chi$.)
In particular, we conclude that a itself must belong to $S^{\mathcal{A}}(6 \cdot 7^j, \overline{a})$ and hence $S^{\mathcal{A}}(7^j, a) \subseteq S^{\mathcal{A}}(7^{j+1}, \overline{a})$.
We then obtain from (1) that there is a $b \in S^{\mathcal{B}}(6 \cdot 7^j, \overline{b})$ with $(\mathcal{S}^{\mathcal{A}}(7^j, \overline{a}a), \overline{a}a) \cong (\mathcal{S}^{\mathcal{B}}(7^j, \overline{b}b), \overline{b}b)$.

Case 2.2. $\mathcal{A} \models \chi$ (and $\mathcal{B} \models \chi$).
Due to $\mathcal{B} \models \chi$, together with $(**)$, we know that there must be a b in $\mathcal{B}$ with $b \notin S^{\mathcal{B}}(2 \cdot 7^j, \bar{b})$ such that $(\mathcal{S}^{\mathcal{B}}(7^j, b), b) \cong (\mathcal{S}^{\mathcal{A}}(7^j, a), a)$. In particular, $S^{\mathcal{B}}(7^j, \bar{b}) \cap S^{\mathcal{B}}(7^j, b) = \emptyset$ and, since we are in *Case 2*, we also have $S^{\mathcal{A}}(7^j, \bar{a}) \cap S^{\mathcal{B}}(7^j, a) = \emptyset$. Furthermore, from (1) we obtain that $(\mathcal{S}^{\mathcal{B}}(7^j, \bar{b}), \bar{b}) \cong (\mathcal{S}^{\mathcal{A}}(7^j, \bar{a}), \bar{a})$. In total we thus conclude that $(\mathcal{S}^{\mathcal{B}}(7^j, \bar{b}b), \bar{b}b) \cong (\mathcal{S}^{\mathcal{A}}(7^j, \bar{a}a), \bar{a}a)$. This completes the proof of Lemma 8.1 □

Theorem 8.2. *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 $\mathfrak{D}_d$.*

Proof. Let m be the quantifier rank of the given sentence φ. For sure, $m \leq ||\varphi||$. We use Lemma 8.1 and note that there are at most $M = 2^{\mathcal{O}(f(m))}$ different basic local sentences $\chi_1, \ldots, \chi_M$ of size $\leq f(m)$ each. For each $I \subseteq \{1, .., M\}$ we let ψ_I be the following sentence

$$\psi_I \ := \ \bigwedge_{i \in I} \chi_i \ \wedge \bigwedge_{i \in \{1,..,M\} \setminus I} \neg \chi_i \, .$$

Furthermore, we let J be the collection of all $I \subseteq \{1, .., M\}$ for which there exists a structure $\mathcal{A}_I \in \mathfrak{D}_d$ such that $\mathcal{A}_I \models \psi_I$ and $\mathcal{A}_I \models \varphi$. Finally, we choose

$$\psi \ := \ \bigvee_{I \in J} \psi_I \, .$$

Using Lemma 8.1 it is straightforward to see that ψ is equivalent to φ for all structures in $\mathfrak{D}_d$.

Furthermore, for each $I \subseteq \{1, .., M\}$ we have $||\psi_I|| = \mathcal{O}(M \cdot f(m)) = 2^{\mathcal{O}(f(m))}$. Since $|J| \leq 2^M$ we thus have $||\psi|| = 2^{2^{f(m)}}$. Recall from Lemma 8.1 that f is 2-fold exponential in m. Thus, $||\psi||$ is 4-fold exponential in m. □

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. This can be proved by observing that there is an elementary upper bound on the number of binary labelled trees of a given height, and then using Hanf's theorem (see [7, 19]) to obtain a normal form for first-order formulas based on the isomorphism types of neighbourhoods. The size of formulas in this normal form is elementary in the size of the original formulas, and the complexity of the translation into the normal form is also elementary. The satisfiability of formulas in this normal form can then be decided by automata theoretic techniques.

Refining the methods of [1], one also obtains an elementary upper bound for the following variant of the Łoś-Tarski Theorem.

Theorem 8.3.

1. *There is a 5-fold exponential function f such that any formula φ of length m 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 formula of length at most $f(m)$.*
2. *There is a 3-fold exponential function g such that any formula φ of length m that is preserved under extensions on the class of acyclic structures in $\mathfrak{D}_2$ is equivalent, on this class, to an existential first-order formula of length at most $g(m)$.*

Let us mention that in all the above cases for structures of bounded degree we can prove at least 2-fold exponential lower bounds.

The aim, in the rest of this section is to give a proof of Theorem 8.3.

For simplicity, we assume a vocabulary τ consisting of one binary relation E and any number of unary relations. The results generalise without any great difficulty to vocabularies consisting of any number of binary relations. Moreover, we assume that E is interpeted as a symmetric relation. This will allow us to elide the distinction between the structure and its underlying Gaifman graph.

Let $\mathfrak{F}$ be the class of acyclic τ-structures, and let $\mathfrak{S}$ be the subclass of $\mathfrak{F}$ consisting of structures where each element has degree at most 2. Thus, structures in $\mathfrak{F}$ are coloured forests while structures in $\mathfrak{S}$ can be seen as disjoint unions of strings. We also write $\mathfrak{F}_d$ for the subclass of $\mathfrak{F}$ where each element has degree at most $d+1$. In particular, $\mathfrak{F}_2$ can be seen as the class of disjoint unions of coloured binary trees.

Let φ be a formula closed under extensions on $\mathfrak{S}$. We say that $\mathcal{A}$ is a *minimal model* of φ if $\mathcal{A} \models \varphi$ and every proper induced substructure $\mathcal{A}'$ of $\mathcal{A}$ is such that $\mathcal{A}' \not\models \varphi$. It is well-known that on any class $\mathfrak{C}$ of finite structures that is closed under taking substructures, a formula φ that is closed under extensions on $\mathfrak{C}$ is equivalent over $\mathfrak{C}$ to an existential formula if, and only if, φ has finitely many minimal models in $\mathfrak{C}$. For our purposes, we will need the following strengthening of this fact.

Lemma 8.4. *Let $\mathcal{C}$ be a class of finite τ-structures that is closed under substructures and φ be a τ-sentence that is preserved under extensions on $\mathcal{C}$. If there is an integer N such that all minimal models of φ have at most N elements, then there is an existential sentence ψ that is equivalent to φ on $\mathfrak{C}$ and such that the length of ψ is $\mathcal{O}(N^{|\varphi|})$, where $|\varphi|$ denotes the length of φ.*

Proof. Let $x_1, \ldots, x_k$ be an enumeration of all variables that appear in φ and suppose, without loss of generality, that no variable in φ is bound in more than one place. Let $y_1, \ldots, y_N$ be a collection of N new variables distinct from $x_1, \ldots, x_k$. We construct from the sentence φ a new *quantifier-free* formula φ^* with free variables $y_1, \ldots, y_N$. More generally, let σ be a function from $\{1, \ldots, k\}$ to $\{1, \ldots, N\}$ and ψ be a formula with free variables among $x_1, \ldots, x_k$. We define the formula $(\psi)^*_\sigma$ by induction as follows.

- If ψ is atomic, then $(\psi)^*_\sigma$ is $\psi[x_i/y_{\sigma(i)}]$, i.e. the formula obtained by replacing all occurrences of variables x_i by $y_{\sigma(i)}$.
- If ψ is $\neg\psi'$, $(\psi)^*_\sigma$ is $\neg(\psi')^*_\sigma$ and similarly if ψ is $\psi_1 \wedge \psi_2$ then $(\psi)^*_\sigma$ is $(\psi_1)^*_\sigma \wedge (\psi_2)^*_\sigma$.
- If ψ is $\exists x_i \psi'$, $(\psi)^*_\sigma$ is $\bigvee_{1 \leq j \leq N} (\psi')^*_{\sigma[i \mapsto j]}$ where $\sigma[i \mapsto j]$ is the function σ' that agrees with σ on all values except i and $\sigma'(i) = j$.

Now, for the sentence φ, let $\hat{\varphi}$ denote the sentence by existentially quantifying the variables $y_1, \ldots, y_N$ in the formula $(\varphi)^*_\sigma$ for an arbitrary σ (since φ has no free variables, the choice of σ does not make a difference). It is easily seen that $\mathcal{A} \models \hat{\varphi}$ if, and only if, $\mathcal{A}$ contains a substructure generated by at most N elements which satisfies φ. However, as φ is preserved under extensions and all its minimal models have at most N elements, this is equivalent to the statement that $\mathcal{A} \models \varphi$. We conclude that $\hat{\varphi}$ is equivalent to φ.

For a bound on the length of $\hat{\varphi}$, note that the only length increasing step in the inductive translation is the one for the quantifier, which increases the size of the formula by a factor of N. Since the number of quantifiers is bounded by the length of φ, the result follows. □

Thus, to establish upper bounds on the length of existential formulas, it suffices to establish upper bounds on the sizes of minimal models.

Let $\mathcal{A}$ and $\mathcal{B}$ be structures and $\overline{a}, \overline{b}$ be tuples of at most m elements from $\mathcal{A}$ and $\mathcal{B}$ respectively. We write $(\mathcal{A}, \overline{a}) \equiv_m (\mathcal{B}, \overline{b})$ to denote that the first-order quantifier-rank m-type of $\overline{a}$ in $\mathcal{A}$ is the same as the first-order m-type of $\overline{b}$ in $\mathcal{B}$. The equivalence relation $\equiv_m$ is characterized by Ehrenfeucht-Fraïssé games (see, for instance, [7]). These can be used to show that the relation is a congruence with respect to disjoint union with a multiplicity threshold of m. A precise statement of this property is given in the following lemma. We write $\mathcal{A} \oplus \mathcal{B}$ to denote the disjoint union of the structures $\mathcal{A}$ and $\mathcal{B}$ and $n\mathcal{A}$ to denote the disjoint union of n copies of $\mathcal{A}$ (see [7, Prop. 2.3.10] for a proof).

Lemma 8.5. *Let $\mathcal{A}_1$, $\mathcal{A}_2$, $\mathcal{B}_1$, and $\mathcal{B}_2$ be structures, and let m, n and n' be integers.*

1. *If $\mathcal{A}_1 \equiv_m \mathcal{B}_1$ and $\mathcal{A}_2 \equiv_m \mathcal{B}_2$ then $\mathcal{A}_1 \oplus \mathcal{A}_2 \equiv_m \mathcal{B}_1 \oplus \mathcal{B}_2$.*
2. *If $n, n' \geq m$ and $\mathcal{A} \equiv_m \mathcal{B}$ then $n\mathcal{A} \equiv_m n'\mathcal{B}$.*

We will use a slightly more general version of this lemma, which is established by a similar game based proof. To state it, we first require some notation. Given two structures $\mathcal{A}$ and $\mathcal{B}$ and a tuple of elements $\overline{c} \in A \cap B$, we write $\mathcal{A} \oplus_{\overline{c}} \mathcal{B}$ to denote the structure obtained from the disjoint union of $\mathcal{A}$ and $\mathcal{B}$ by identifying the two copies of $\overline{c}$. If $\mathcal{A}$ is a structure and $\overline{a}$ a tuple of elements, we write $\bigoplus^n_{\overline{a}} \mathcal{A}$ to denote the structure obtained from the disjoint union of n copies of $\mathcal{A}$ by identifying all distinct copies of $\overline{a}$.

Lemma 8.6. *If $(\mathcal{A}, \overline{ac}) \equiv_m (\mathcal{A}', \overline{a}'\overline{c}')$ and $(\mathcal{B}, \overline{bc}) \equiv_m (\mathcal{B}', \overline{b}'\overline{c}')$ then $(\mathcal{A} \oplus_{\overline{c}} \mathcal{B}, \overline{abc}) \equiv_m (\mathcal{A}' \oplus_{\overline{c}'} \mathcal{B}', \overline{a}'\overline{b}'\overline{c}')$.*

Lemma 8.7. *If $n, n' \geq m$ and $(\mathcal{A}, \overline{a}) \equiv_m (\mathcal{B}, \overline{b})$ then $\bigoplus_{\overline{a}}^{n} \mathcal{A} \equiv_m \bigoplus_{\overline{b}}^{n'} \mathcal{B}$.*

8.1 Strings

The *Hanf type* of radius r of a structure $\mathcal{A}$ is the multiset of isomorphism types of r-neighbourhoods of elements in $\mathcal{A}$. We say that two structures $\mathcal{A}$ and $\mathcal{B}$ are *Hanf equivalent* with radius r and threshold q, written $\mathcal{A} \simeq_{r,q} \mathcal{B}$, if, for every $a \in A$, *either* the number of occurrences of the isomorphism type of $N_r^{\mathcal{A}}(a)$ in the Hanf type of $\mathcal{A}$ is the same as that in the Hanf type of $\mathcal{B}$, *or* it is at least q, and conversely for every element $b \in B$. This allows us to state Hanf's locality theorem (see [19, Theo. 4.24] for a proof).

Theorem 8.8 (Hanf Locality). *For every vocabulary τ and every m there are $r \leq 3^m$ and $q \leq m$ such that for any pair of τ-structures $\mathcal{A}$ and $\mathcal{B}$ if $\mathcal{A} \simeq_{r,q} \mathcal{B}$ then $\mathcal{A} \equiv_m \mathcal{B}$.*

This theorem immediately gives us upper bounds on the index of the equivalence relation $\equiv_m$ on the classes of structures $\mathfrak{S}$ and $\mathfrak{F}_d$.

Corollary 8.9. *For any fixed vocabulary τ, the index of $\equiv_m$ on $\mathfrak{S}$ is bounded by a 3-fold exponential function in m.*

Proof. There are at most $n = s^{(2r+1)}$ isomorphism types of r-neighbourhoods among structures in $\mathfrak{S}$, where $s = 2^{|\tau|}$ is a bound on the number of different atomic types of a single element. Thus, this number is bounded by a 2-fold exponential function $f(m)$. Now, by Theorem 8.8, the number of $\equiv_m$-classes is at most $(m+1)^{f(m)}$. □

Corollary 8.10. *For any fixed vocabulary τ and any fixed d, the index of $\equiv_m$ on $\mathfrak{F}_d$ is bounded by a 4-fold exponential function in m.*

Proof. The proof is analogous to that of Corollary 8.9 except now the r-neighbourhood of an element may have as many as $\mathcal{O}(d^r)$ elements and therefore the number of isomorphism types of neighbourhood of radius r is 2-fold exponential in r. □

Theorem 8.8 is used in [1] to establish the following lemma.

Lemma 8.11. *For every vocabulary τ and every $m > 0$ there is a p such that if $\mathcal{A}$ is a connected τ-structure in $\mathfrak{S}$ with $|A| > p$, then there is a disjoint extension $\mathcal{B}$ of $\mathcal{A}$ and a proper substructure $\mathcal{A}'$ of $\mathcal{A}$ such that $\mathcal{A}' \equiv_m \mathcal{B}$.*

In particular, it is shown that if r and q are obtained as in Theorem 8.8, then taking $p = nl(q+l)$, where $l = 2r(n+1)+1$ and n is the number of distinct isomorphism types of $2r$-neighbourhoods in $\mathfrak{S}$, suffices. In particular, this implies that no minimal model of φ in $\mathfrak{S}$ can contain a path longer than p.

Theorem 8.12. *If φ is a formula of length m that is preserved under extensions on $\mathfrak{S}$, then φ is equivalent over $\mathfrak{S}$ to an existential formula whose length is bounded by a 3-fold exponential function in m.*

Proof. Let q be the quantifier rank of φ and τ its vocabulary. For any r, there are at most $n = s^{(4r+1)}$ isomorphism types of $2r$-neighbourhoods among structures in $\mathfrak{S}$, where $s = 2^{|\tau|}$ is a bound on the number of different atomic types of a single element. Taking $r \leq 3^q$ and and $p = nl(q+l)$, where $l = 2r(n+1)+1$, we have $p = \mathcal{O}(r^2n^3 + qrn^2)$. Putting in the bounds on r and n gives us $p = \mathcal{O}(3^{2q}s^{3(4 \cdot 3^q+1)})$. This provides an upper bound on the length of the longest path in a minimal model of φ. Since both q and $|\tau|$ are at most m (we can safely ignore any symbols in τ that do not appear in φ), we have that p is at most $2^{2^{f(m)}}$ for some polynomial function f and sufficiently large values of m.

Now, let g be a polynomial such that the index of $\equiv_m$ is at most $2^{2^{2^{g(m)}}}$ on $\mathfrak{S}$ (such a g exists by Lemma 8.9) and suppose $\mathcal{A}$ is any model of φ in $\mathfrak{S}$ with more than $N = m(2^{2^{2^{g(m)}}})(2^{2^{f(m)}})$ elements. If $\mathcal{A}$ contains a path longer than $2^{2^{f(m)}}$, then $\mathcal{A}$ is not minimal. Suppose then that there is no path longer than

$2^{2^{f(m)}}$. Since $\mathcal{A}$ is a disjoint union of strings, this means that it consists of more than $m(2^{2^{2^{g(m)}}})$ disjoint components. By the choice of g, this means that there is a set of more than m components that are pairwise $\equiv_m$-equivalent. Let $\mathcal{A}'$ be the structure obtained by removing one of these components. By Lemmas 8.5 and 8.7 we have that $\mathcal{A}' \equiv_m \mathcal{A}$. Thus, we have a substructure of $\mathcal{A}$ which must also be a model of φ, and thus $\mathcal{A}$ is not minimal.

We have shown that any minimal model of φ in $\mathfrak{S}$ has at most N elements. Combining this with Lemma 8.4 establishes the result □

8.2 Forests

In [1] an upper bound was also obtained on the size of minimal models of a sentence φ preserved under extensions on the class $\mathfrak{F}$ of acyclic structures. That bound was obtained as a function of the number of $\equiv_m$ types, where m is the quantifier rank of φ. As we have noted, this number is not bounded by any elementary function of m. However, when we fix the degree d and consider the class of structures $\mathfrak{F}_d$, we obtain tighter bounds as we now explore.

For a structure $\mathcal{A} \in \mathfrak{F}_d$ and a distinguished element a of $\mathcal{A}$, define the $\equiv_m$-type of $(\mathcal{A}, a)$ to be the set of all formulas $\varphi(x)$ of length at most m such that $\mathcal{A} \models \varphi[a]$. As noted in Corollary 8.10 the number of types is bounded by a 4-fold exponential function of m. Let $\theta_1, \dots, \theta_n$ be an enumeration of all types. We refer to a as *the distinguished element* of $(\mathcal{A}, a)$. We define a new vocabulary τ' which consists of the binary relation E and a unary relation T_i for each θ_i.

Let $\mathcal{A}$ be a connected τ'-structure that is acyclic and with degree at most 2. We construct from $\mathcal{A}$ a τ-structure $\tilde{\mathcal{A}}$ as follows: each element $a \in A$ with $T_i(a)$ is replaced by a structure $\mathbf{T}_a$ of type θ_i. Moreover, for the binary relation E, $(b, c) \in E^{\tilde{\mathcal{A}}}$ if, and only if, *either* b and c are in the same structure $\mathbf{T}_a$ and $(b, c) \in E^{\mathbf{T}_a}$ *or* b is the distinguished element of $\mathbf{T}_a$, c is the distinguished element of $\mathbf{T}_{a'}$ and $(a, a') \in E^{\mathcal{A}}$. We call $\tilde{\mathcal{A}}$ a τ-companion of $\mathcal{A}$. The structure $\tilde{\mathcal{A}}$ is not uniquely determined by $\mathcal{A}$ as there are, in general, many structures of type θ_i. However, we can establish the following.

Lemma 8.13. *If $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{A}}'$ are two τ-companions of $\mathcal{A}$, then $\tilde{\mathcal{A}} \equiv_m \tilde{\mathcal{A}}'$.*

Proof. The proof is by induction on the length l of $\mathcal{A}$. Indeed if $l = 1$ then by definition $\tilde{\mathcal{A}}$ and $\tilde{\mathcal{A}}'$ are two trees from the same $\equiv_m$ equivalence class.

Suppose now that $\mathcal{A}$ is of length $l + 1$ and a is one of its endpoints. Let $\mathcal{B}$ be the structure induced by $\mathcal{A} \setminus \{a\}$. Then, we have $\tilde{\mathcal{A}} = \tilde{\mathcal{B}} \oplus_a \mathbf{T}$ for some τ-companion $\tilde{\mathcal{B}}$ of $\mathcal{B}$ and some tree $\mathbf{T}$. Similarly, $\tilde{\mathcal{A}}' = \tilde{\mathcal{B}}' \oplus_a \mathbf{T}'$. $\tilde{\mathcal{B}} \equiv_m \tilde{\mathcal{B}}'$ by induction hypothesis and $\mathbf{T} \equiv_m \mathbf{T}'$ by definition of τ-companions. Thus, we get the desired result by Lemma 8.6. □

This now allows us to establish the following.

Lemma 8.14. *Let $\mathcal{A}$ and $\mathcal{B}$ be connected, acyclic structures of degree at most 2 with the property that for each element there is a unique i such that T_i holds, and let m be an integer. If $\mathcal{A} \equiv_m \mathcal{B}$ then $\tilde{\mathcal{A}} \equiv_m \tilde{\mathcal{B}}$.*

Proof. Let $\mathbf{T}_1, \dots, \mathbf{T}_n$ be a system of unique representatives of the $\equiv_m$ equivalence classes of τ-trees. Let $\tilde{\mathcal{A}}'$ and $\tilde{\mathcal{B}}'$ be the τ companions of $\mathcal{A}$ and $\mathcal{B}$ that are obtained using only trees from this list. Then, by Lemma 8.13, $\tilde{\mathcal{A}} \equiv_m \tilde{\mathcal{A}}'$ and $\tilde{\mathcal{B}} \equiv_m \tilde{\mathcal{B}}'$. But now, a straightforward Ehrenfeucht-Fraïssé game shows that $\tilde{\mathcal{A}}' \equiv_m \tilde{\mathcal{B}}'$, which establishes the result. □

We are now ready to use this machinery to establish an upper bound on the size of minimal models. For the remainder of this section, fix a first-order sentence φ that is preserved under extensions on $\mathfrak{F}_d$ and let m be the length of φ.

Lemma 8.15. *There is a 4-fold exponential function f such that no minimal model of φ contains a path longer than $f(m)$.*

Proof. Let τ' be the vocabulary, as above, obtained by taking a unary relation for each $\equiv_m$ type of τ-tree and let p be the bound obtained from Theorem 8.12 for the vocabulary τ'. Then, we can show, using Lemma 8.14, that p bounds the length of any minimal model of φ in $\mathfrak{F}_d$ (this argument is essentially the same as in [1, Lemma 3.5]).

To estimate p, note that $|\tau'|$ is bounded by a 4-fold exponential function of m, by Corollary 8.10. Also, from the proof of Theorem 8.12 that $p = \mathcal{O}(3^{2m} s^{3(3^m+1)})$. Note that, since in the τ' structures we construct, each element has a unique unary relation in τ' that it belongs to, we can take $s = |\tau'|$. Together, these facts yield the desired result. □

Lemma 8.16. *There is a* 4*-fold exponential function* h *such that no minimal model of* φ *has more than* $h(m)$ *connected components.*

Proof. Let t be a 4-fold exponential function such that the number of $\equiv_m$ types in $\mathfrak{F}_d$ is bounded by $t(m)$ and suppose $\mathcal{A}$ has more than $m(t(m))$ components. Since there are at most $t(m)$ types, there must be some type that occurs more than m times among the components. Let $\mathcal{A}'$ be the structure obtained by deleting one of these components. It is an easy application of Lemma 8.5 to show that $\mathcal{A}' \equiv_m \mathcal{A}$. □

Theorem 8.17. *There is a* 5*-fold exponential function* t *such that any formula* φ *of length* m *that is preserved under extensions on* $\mathfrak{F}_d$ *is equivalent to an existential formula of length at most* $t(m)$*.*

Proof. Putting together Lemmas 8.15, and 8.16, we know that any minimal model of φ has at most $h(m)((d+1)^{f(m)})$ elements. Together with Lemma 8.4, this gives us the desired result □

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] A. Atserias, A. Dawar, and P. G. Kolaitis. On preservation under homomorphisms and unions of conjunctive queries. *Journal of the ACM*, 53:208–237, 2006.

[3] M. Benedikt and L. Segoufin. Towards a characterization of order-invariant queries over tame structures. In C.-H. Ong, editor, *Proceedings of the 19th International Workshop on Computer Science Logic*, volume 3634 of *Lecture Notes in Computer Science*, pages 276–291. Springer Verlag, 2005.

[4] A. Church. A note on the Entscheidungsproblem. *Journal of Symbolic Logic*, 1:40–41, 1936.

[5] N. J. Cutland. *Computability*. Cambridge University Press, 1980.

[6] 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.

[7] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.

[8] S. Feferman and R. Vaught. The first order properties of products of algebraic systems. *Fundamenta Mathematicae*, 47:57–103, 1959.

[9] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.

[10] M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *Journal of the ACM*, 48:1184–1206, 2001.

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] W. Hodges. *Model Theory*. Cambridge University Press, 1993.

[17] 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.

[18] L. Libkin. Logics with counting and local properties. *ACM Transactions on Computational Logic*, 1:33–59, 2000.

[19] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.

[20] 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.

[21] M. Presburger. Über die Vollständigkeit eines gewissen Systems der Arithmetik ganzer Zahlen in welchem die Addition als einzige Operation hervortritt. *Comptes Rendus I Congrès des Mathématiciens des Pays Slaves*, Warsaw, 1930.

[22] M. Rabin. Decidability of second order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969.

[23] B. Rossman. Existential positive types and preservation under homomorphisisms. In *20th IEEE Symposium on Logic in Computer Science*, pages 467–476, 2005.

[24] 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.

[25] W. W. Tait. A counterexample to a conjecture of Scott and Suppes. *Journal of Symbolic Logic*, 24:15–16, 1959.

[26] A. Tarski. *A Decision Method for Elementary Algebra and Geometry*. Rand, 2nd edition, 1951.

[27] 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.

[28] A. Turing. On computable numbers with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society, Series 2*, 42:230–265, 1936.