

## School of Informatics, University of Edinburgh

## Laboratory for Foundations of Computer Science

Comparing the succinctness of monadic query languages over finite trees
by
Martin Grohe, Nicole Schweikardt

Informatics Research Report EDI-INF-RR-0168

# Comparing the succinctness of monadic query languages over finite trees 

Martin Grohe, Nicole Schweikardt<br>Informatics Research Report EDI-INF-RR-0168<br>SCHOOL of INFORMATICS<br>Laboratory for Foundations of Computer Science

June 2003

Extended abstract appears in Proc. CSL'03


#### Abstract

: We study the succinctness of monadic second-order logic and a variety of monadic fixed point logics on trees. All these languages are known to have the same expressive power on trees, but some can express the same queries much more succinctly than others. For example, we show that, under some complexity theoretic assumption, monadic second-order logic is non-elementarily more succinct than monadic least fixed point logic, which in turn is nonelementarily more succinct than monadic datalog.

Succinctness of the languages is closely related to the combined and parameterized complexity of query evaluation for these languages.


Keywords : Finite Model Theory, Monadic Second-Order Logic, Fixed Point Logics, mu-Calculus, Monadic Datalog, Tree-like structures, Succinctness

Copyright (c) 2003 by The University of Edinburgh. All Rights Reserved

The authors and the University of Edinburgh retain the right to reproduce and publish this paper for non-commercial purposes.

Permission is granted for this report to be reproduced by others for non-commercial purposes as long as this copyright notice is reprinted in full in any reproduction. Applications to make other use of the material should be addressed in the first instance to Copyright Permissions, School of Informatics, The University of Edinburgh, 2 Buccleuch Place, Edinburgh EH8 9LW, Scotland.

# Comparing the succinctness of monadic query languages over finite trees 

Martin Grohe and Nicole Schweikardt*<br>Laboratory for Foundations of Computer Science<br>University of Edinburgh, Scotland, U.K.<br>email: \{grohe|v1nschwe\}@inf.ed.ac.uk


#### Abstract

We study the succinctness of monadic second-order logic and a variety of monadic fixed point logics on trees. All these languages are known to have the same expressive power on trees, but some can express the same queries much more succinctly than others. For example, we show that, under some complexity theoretic assumption, monadic second-order logic is non-elementarily more succinct than monadic least fixed point logic, which in turn is non-elementarily more succinct than monadic datalog. Succinctness of the languages is closely related to the combined and parameterized complexity of query evaluation for these languages.


Keywords: Finite Model Theory, Monadic Second-Order Logic, Fixed Point Logics, $\mu$-Calculus, Monadic Datalog, Tree-like structures, Succinctness

## 1 Introduction

A central topic in finite model theory has always been a comparison of the expressive power of different logics on finite relational structures. In particular, the expressive power of fragments of monadic second-order logic and various fixed-point logics has already been investigated in some of the earliest papers in finite model theory [Fag75,CH82]. One of the main motivations for such studies was an interest in the expressive power of query languages for relational databases.

In recent years, the focus in database theory has shifted from relational to semi-structured data and in particular data stored as XML-documents. A lot of current research in the database community is concerned with the design and implementation of XML query languages (see, for example, [FSW00,HP00,GK02] or the monograph [ABS99] for a general introduction into semi-structured data and XML). The languages studied in the present paper may be viewed as node-selecting query languages for XML. They all contain the core of the language XPath, which is an important building block of several major XML-related technologies. Recently, monadic datalog has been proposed as a nodeselecting query language with a nice balance between expressive power and very good algorithmic properties [GK02,Koc03].

XML-documents are best modelled by trees, or more precisely, finite labelled ordered unranked trees. It turns out that when studying node-selecting query languages for XML-documents, expressive power is not the central issue. Quite to the contrary: Neven and Schwentick [NSO2] proposed to take the expressive power of monadic second-order logic (MSO) as a benchmark for node-selecting XML-query languages and, in some sense, suggested that such languages should at least have the expressive power of MSO. However, even languages with the same expressive power may have vastly different complexities. For example, monadic datalog and MSO have the same expressive power over trees [GK02]. However, monadic datalog queries can be evaluated in time linear both in the size of the datalog program and the size of the input tree [GK02], and thus the combined complexity of monadic datalog is in polynomial time, whereas the evaluation of MSO queries is PSPACE complete. The difference becomes even more obvious if we look at parameterized complexity: Unless PTIME $\neq$ NP, there is no algorithm evaluating a monadic second-order query in time $f$ (size of query) $p$ (size of tree)

[^0]for any elementary function $f$ and polynomial $p$ [FG03]. Similar statements hold for the complexity of the satisfiability problem for monadic datalog and MSO over trees. The reason for this different behaviour is that even though the languages have the same expressive power on trees, in MSO we can express queries much more succinctly. Indeed, there is no elementary translation from a given MSOformula into an equivalent monadic datalog program. We also say that MSO is non-elementarily more succinct than monadic datalog. Just to illustrate the connection between succinctness and complexity, let us point out that if there was an elementary translation from MSO to monadic datalog, then there would be an algorithm evaluating a monadic second-order query in time $f$ (size of query) $p$ (size of tree) for an elementary function $f$ and a polynomial $p$.

In this paper, we study the succinctness (in the sense just described) of a variety of fixed point logics on finite trees. Our main results are the following:

1. MSO is non-elementarily more succinct than monadic least fixed point logic MLFP (see Theorem 3.2). Unfortunately, we are only able to prove this result under the odd, but plausible complexity theoretic assumption that for some $i \geqslant 1, \mathrm{NP}$ is not contained in $\operatorname{DTIME}\left(n^{\log ^{(i)}(n)}\right)$, where $\log ^{(i)}$ denotes the $i$ times iterated logarithm.
2. MLFP is non-elementarily more succinct than its 2-variable fragment MLFP $^{2}$ (see Corollary 6.16).
3. $\mathrm{MLFP}^{2}$ is exponentially more succinct than the full modal $\mu$-calculus, that is, the modal $\mu$-calculus with future and past modalities (see Theorem 4.5, Example 4.3, and Theorem 4.6).
4. The full modal $\mu$-calculus is at most exponentially more succinct than stratified monadic datalog, and conversely, stratified monadic datalog is at most exponentially more succinct than the full modal $\mu$-calculus (see Theorem 6.10 and 6.11). Furthermore, stratified monadic datalog is at most exponentially more succinct than monadic datalog (see Theorem 6.9).
The exact relationship between these three languages remains open.

Of course we are a not the first to study the succinctness of logics with the same expressive power. Most known results are about modal and temporal logics. The motivation for these results has not come from database theory, but from automated verification and model-checking. The setting, however, is very similar. For example, Kamp's well know theorem states that first-order logic and linear time temporal logic have the same expressive power on strings [Kam68], but there is no elementary translation from first-order logic to linear time temporal logic on strings. Even closer to our results, monadic secondorder logic and the modal $\mu$-calculus have the same expressive power on (ordered) trees, but again is well-known that there is no elementary translation from the former to the latter. Both of these results can be proved by simple automata theoretic arguments. More refined results are known for various temporal logics [Wil99,AI00,AI01,EVW02]. By and large, however, succinctness has received surprisingly little attention in the finite model theory community. Apart from automata theoretic arguments, almost no good techniques for proving lower bounds on formula sizes are known. A notable exception are Adler and Immerman's [AIO1] nice games for proving such lower bounds. Unfortunately, we found that these games (adapted to fixed point logic) were of little use in our context. So we mainly rely on automata theoretic arguments. An exception is the, complexity theoretically conditioned, result that MSO is nonelementarily more succinct than MLFP. To prove this result, we are building on a technique introduced in [FG03].

The paper is organised as follows: In Section 2 we fix the basic notations used throughout the paper. Section 3 concentrates on the translation from MSO to MLFP. In Section 4 we present our results concerning the two-variable fragment of MLFP and the full modal $\mu$-calculus. Section 5 compares MLFP with its extension by simultaneous least fixed point operators. In Section 6 we concentrate on monadic datalog, stratified monadic datalog, and their relations to finite automata and to MLFP. Finally, Section 7 concludes the paper by pointing out several open questions.

Detailed proofs of our results are provided in an appendix.

## 2 Preliminaries

### 2.1 Basic Notations

Given a set $\Sigma$ we write $\Sigma^{*}$ to denote the set of all finite strings over $\Sigma$, and we use $\varepsilon$ to denote the empty string. We use $\mathbb{N}$ to denote the set $\{0,1,2, \ldots\}$ of natural numbers. We use lg to denote the logarithm with respect to base 2 . With a function $f$ that maps natural numbers to real numbers we associate the corresponding function from $\mathbb{N}$ to $\mathbb{N}$ defined by $n \mapsto\lceil f(n)\rceil$. For simplicity we often simply write $f(n)$ instead of $\lceil f(n)\rceil$.

The function Tower : $\mathbb{N} \rightarrow \mathbb{N}$ is inductively defined via $\operatorname{Tower}(0):=1$ and $\operatorname{Tower}(h+1)=$ $2^{\text {Tower }(h)}$, for all $h \in \mathbb{N}$. I.e., $\operatorname{Tower}(h)$ is a tower of 2 s of height $h$.

We say that a function $f: \mathbb{N} \rightarrow \mathbb{N}$ has bound $f(m) \leqslant \operatorname{Tower}(o(h(m))$, for some function $h:$ $\mathbb{N} \rightarrow \mathbb{N}$, if there is a function $g \in o(h)$ and a $m_{0} \in \mathbb{N}$ such that for all $m \geqslant m_{0}$ we have $f(m) \leqslant$ Tower $(g(m))$.

Note that, in particular, every elementary function $f$ has bound $f(m) \leqslant \operatorname{Tower}(o(m))$. Indeed, for every elementary function $f$ there is a $h \in \mathbb{N}$ such that, for all $n \in \mathbb{N}, f(n)$ is less than or equal to the tower of 2 s of height $h$ with an $n$ on top.

### 2.2 Structures

A signature $\tau$ is a finite set of relation symbols and constant symbols. Each relation symbol $R \in \tau$ has a fixed arity $\operatorname{ar}(R)$. A $\tau$-structure $\mathcal{A}$ consists of a set $\mathcal{U}^{\mathcal{A}}$ called the universe of $\mathcal{A}$, an interpretation $c^{\mathcal{A}} \in \mathcal{U}^{\mathcal{A}}$ of each constant symbol $c \in \tau$, and an interpretation $R^{\mathcal{A}} \subseteq\left(\mathcal{U}^{\mathcal{A}}\right)^{a r(R)}$ of each relation symbol $R \in \tau$. All structures considered in this paper are assumed to have a finite universe.

The main focus of this paper lies on the class Trees of finite binary trees. Precisely, finite binary trees are particular structures over the signature

$$
\tau_{\text {Trees }}:=\left\{\text { Root, } 1^{\text {st }} \text { Child, } 2^{\text {nd }} \text { Child, Has-No-1 } 1^{\text {st }} \text { Child, Has-No- } 2^{\text {nd }} \text { Child }\right\},
$$

where Root, Has-No-1 ${ }^{\text {st }}$ Child, Has-No- $2^{\text {nd }}$ Child are unary relation symbols and $1^{\text {st }}$ Child, $2^{\text {nd }}$ Child are binary relation symbols. We define Trees to be the set of all $\tau_{\text {Trees-structures }} T$ that satisfy the following conditions:

1. $\mathcal{U}^{T} \subset\{1,2\}^{*}$ and for every string $s i \in \mathcal{U}^{T}$ with $i \in\{1,2\}$ we also have $s \in \mathcal{U}^{T}$.
2. Root ${ }^{T}$ consists of the empty string $\varepsilon$.
3. $1^{s t}$ Child ${ }^{T}$ consists of the pairs $(s, s 1)$, for all $s 1 \in \mathcal{U}^{T}$.
4. $2^{\text {nd }}$ Child ${ }^{T}$ consists of the pairs $(s, s 2)$, for all $s 2 \in \mathcal{U}^{T}$.
5. Has-No-1 $1^{s t}$ Child ${ }^{T}$ consists of all strings $s \in \mathcal{U}^{T}$ with $s 1 \notin \mathcal{U}^{T}$.
6. Has-No-2 ${ }^{\text {nd }}$ Child $^{T}$ consists of all strings $s \in \mathcal{U}^{T}$ with $s 2 \notin \mathcal{U}^{T}$.

For $T \in$ Trees and $t \in \mathcal{U}^{T}$ we write $T_{t}$ to denote the subtree of $T$ with root $t$.
A schema $\sigma$ is a set of unary relation symbols each of which is distinct from Has-No-1 ${ }^{\text {st }}$ Child, Has-No-2 $2^{\text {nd }}$ Child, Root. A $\sigma$-labelled tree is a $\left(\tau_{\text {Trees }} \cup \sigma\right)$-structure consisting of some $T \in$ Trees and additional interpretations $P^{T} \subseteq \mathcal{U}^{T}$ for all symbols $P \in \sigma$. We sometimes write label $(t)$ to denote the set $\left\{P \in \sigma: t \in P^{T}\right\}$ of labels at vertex $t$ in $T$.

We identify a string $w=w_{0} \cdots w_{n-1}$ of length $|w|=n \geqslant 1$ over an alphabet $\Sigma$ with a $\sigma$-labelled tree $T^{w}$ in the following way: We choose $\sigma$ to consist of a unary relation symbol $P_{a}$ for each letter $a \in$ $\Sigma$, we choose $T^{w}$ to be the (unique) element in Trees with universe $\mathcal{U}^{T^{w}}=\left\{\varepsilon, 1,11, \ldots, 1^{n-1}\right\}$, and we choose $P_{a}^{T^{w}}:=\left\{1^{i}: w_{i}=a\right\}$, for each $a \in \Sigma^{w}$. This corresponds to the conventional representation of strings by structures in the sense that $\left\langle\mathcal{U}^{T^{w}}, 1^{s t} \operatorname{Child},\left(P_{a}^{T^{w}}\right)_{a \in \Sigma}\right\rangle$ is isomorphic to the structure $\left\langle\{0, \ldots, n-1\}, \operatorname{Succ},\left(P_{a}^{w}\right)_{a \in \Sigma}\right\rangle$ where Succ denotes the binary successor relation on $\{0, . ., n-1\}$ and $P_{a}^{w}$ consists of all positions of $w$ that carry the letter $a$. When reasoning about strings in the context of first-order logic, we sometimes also need the linear ordering $<$ on $\{0, \ldots, n-1\}$ (respectively, the transitive closure of the relation $1^{\text {st }}$ Child $)$. In these cases we explicitly write $\mathrm{FO}(<)$ rather than FO to indicate that the linear ordering is necessary.

XML-documents are usually modelled as ordered unranked trees and not as binary trees. Here ordered refers to the fact that the order of the children of a vertex is given. However, a standard representation of ordered unranked trees as relational structures uses binary relations $1^{\text {st }}$ Child,Next-Sibling and unary relations Root, Leaf, Last-Sibling (for details, see [GK02]) and thus essentially represents ordered unranked trees as binary trees. Therefore, all our results also apply to ordered unranked trees.

### 2.3 Logics and Queries

We assume that the reader is familiar with first-order logic, for short: FO, and with monadic secondorder logic, for short: MSO (cf., e.g., the textbooks [EF99,Imm99]). We use $\mathrm{FO}(\tau)$ and $\mathrm{MSO}(\tau)$, respectively, to denote the class of all first-order formulas and monadic second-order formulas, respectively, of signature $\tau$. We write $\varphi\left(x_{1}, \ldots, x_{k}, X_{1}, \ldots, X_{\ell}\right)$ to indicate that the free first-order variables of the formula $\varphi$ are $x_{1}, \ldots, x_{k}$ and the free set variables are $X_{1}, \ldots, X_{\ell}$. Sometimes we use $\bar{x}$ and $\bar{X}$ as abbreviations for sequences $x_{1}, \ldots, x_{k}$ and $X_{1}, \ldots, X_{\ell}$ of variables.

A formula $\varphi(x)$ of signature $\tau$ defines the unary query which associates with every $\tau$-structure $\mathcal{A}$ the set of elements $a \in \mathcal{U}^{\mathcal{A}}$ such that $\mathcal{A} \models \varphi(a)$, i.e., $\mathcal{A}$ satisfies $\varphi$ when interpreting the free occurrences of the variable $x$ by the element $a$. A sentence $\varphi$ of signature $\tau$ (i.e., a formula that has no free variables) defines the Boolean query that associates the answer "yes" with all $\tau$-structures that satisfy $\varphi$ and the answer "no" with all other $\tau$-structures.

Apart from FO and MSO we will also consider monadic least fixed point logic MLFP (cf., e.g., [EF99]) which is the extension of first-order logic by unary least fixed point operators, defined as follows: Let $\tau$ be a signature and let $\varphi(x, X)$ be a formula of signature $\tau$ which is positive in the variable $X$, i.e., every atom of the form $X(z)$ occurs in $\varphi$ within an even number of negation symbols. $\varphi$ defines for every $\tau$-structure $\mathcal{A}$ a monotone operator ${ }^{1} F_{\mathcal{A}, \varphi}: 2^{\mathcal{U}^{\mathcal{A}}} \rightarrow 2^{\mathcal{U}^{\mathcal{A}}}$ via $F_{\mathcal{A}, \varphi}(A):=\{a \in$ $\left.\mathcal{U}^{\mathcal{A}}: \mathcal{A} \models \varphi(a, A)\right\}$, for every $A \subseteq \mathcal{U}^{\mathcal{A}}$. A set $A$ is called a fixed point (respectively, pre fixed point, respectively, post fixed point) of $\varphi$ in $\mathcal{A}$ iff $F_{\mathcal{A}, \varphi}(A)=A$ (respectively, $F_{\mathcal{A}, \varphi}(A) \subseteq A$, respectively, $\left.F_{\mathcal{A}, \varphi}(A) \supseteq A\right)$.
For all $s \in \mathbb{N}$ we define $L_{\mathcal{A}, \varphi}^{0}:=\emptyset, L_{\mathcal{A}, \varphi}^{s+1}:=F_{\mathcal{A}, \varphi}\left(L_{\mathcal{A}, \varphi}^{s}\right)$, and $G_{\mathcal{A}, \varphi}^{0}:=\mathcal{U}^{\mathcal{A}}, G_{\mathcal{A}, \varphi}^{s+1}:=F_{\mathcal{A}, \varphi}\left(G_{\mathcal{A}, \varphi}^{s}\right)$. Since $F_{\mathcal{A}, \varphi}$ is monotone we have $L_{\mathcal{A}, \varphi}^{s} \subseteq L_{\mathcal{A}, \varphi}^{s+1}$ and $G_{\mathcal{A}, \varphi}^{s} \supseteq G_{\mathcal{A}, \varphi}^{s+1}$, for all $s \in \mathbb{N}$. Since $\mathcal{U}^{\mathcal{A}}$ is finite, a fixed point will be reached eventually, i.e., there is a $s_{0}$ such that $L_{\mathcal{A}, \varphi}^{s_{0}}=L_{\mathcal{A}, \varphi}^{s_{0}+1}=: L_{\mathcal{A}, \varphi}^{\infty}$ and $G_{\mathcal{A}, \varphi}^{s_{0}}=G_{\mathcal{A}, \varphi}^{s_{0}+1}=: G_{\mathcal{A}, \varphi}^{\infty}$. We define $L_{\mathcal{A}, \varphi}^{\infty}$ to be the least fixed point and $G_{\mathcal{A}, \varphi}^{\infty}$ to be the greatest fixed point of $\varphi$ in $\mathcal{A}$. It is straightforward to see that $L_{\mathcal{A}, \varphi}^{\infty}$ is indeed contained in every pre fixed point of $\varphi$ in $\mathcal{A}$, and that every post fixed point of $\varphi$ in $\mathcal{A}$ is contained in $G_{\mathcal{A}, \varphi}^{\infty}$. The sets $L_{\mathcal{A}, \varphi}^{s}$, for $s \in \mathbb{N}$, are called the stages of the least fixed point of $\varphi$ in $\mathcal{A}$.

The logic $\operatorname{MLFP}(\tau)$ is the extension of $\operatorname{FO}(\tau)$ by least and greatest fixed point operators. I.e.: $\operatorname{MLFP}(\tau)$ contains $\mathrm{FO}(\tau)$ and is closed under Boolean connectives and first-order quantifications; and if $\varphi(x, X, \bar{y}, \bar{Y})$ is an $\operatorname{MLFP}(\tau)$-formula which is positive in the variable $X$ then $\left[\operatorname{LFP}_{x, X} \varphi\right](z)$ and $\left[\operatorname{GFP}_{x, X} \varphi\right](z)$ are $\operatorname{MLFP}(\tau)$-formulas such that for every $(\tau \cup\{\bar{y}, \bar{Y}\})$-structure $\mathcal{A}$ and every element $a \in \mathcal{U}^{\mathcal{A}}$ we have $\mathcal{A} \models\left[\operatorname{LFP}_{x, X} \varphi\right](a)$ iff $a \in L_{\mathcal{A}, \varphi}^{\infty}$, and $\mathcal{A} \models\left[\operatorname{GFP}_{x, X} \varphi\right](a)$ iff $a \in G_{\mathcal{A}, \varphi}^{\infty}$.

It is well-known that the greatest fixed point is the dual of the least fixed point, i.e., that the formula $\left[\operatorname{GFP}_{x, X} \varphi\right](z)$ is equivalent to the formula $\neg\left[\operatorname{LFP}_{x, X} \neg \varphi^{\prime}\right](z)$, where $\varphi^{\prime}$ is obtained from $\varphi$ by replacing every atom of the form $X(u)$ by the literal $\neg X(u)$. Therefore, every MLFP-formula $\Psi$ can easily be transformed into an equivalent MLFP-formula $\Psi^{\prime}$ in negation normal form, where negation symbols " $\neg$ " only occur directly in front of atomic subformulas.

### 2.4 Formula size and succinctness

In a natural way, we view formulas as finite trees, where leaves correspond to the atoms of the formulas and inner vertices correspond to Boolean connectives, quantifiers, and fixed-point operators. We define the size $\|\varphi\|$ of a formula $\varphi$ to be the number of vertices of the tree that corresponds to $\varphi$.

[^1]Note that this measure of formula size is a uniform cost measure in the sense that it accounts just 1 cost unit for each variable and relation symbol appearing in a formula, no matter what its index is. An alternative is to define the size of a formula as the length of a binary encoding of the formula. Such a logarithmic cost measure is, for example, used in [FG03]. Switching between a uniform and a logarithmic measure usually involves a logarithmic factor.

Definition 2.1 (Succinctness). Let $L_{1}$ and $L_{2}$ be logics, let $F$ be a class of functions from $\mathbb{N}$ to $\mathbb{N}$, and let $\mathcal{C}$ be a class of structures.
We say that $L_{1}$ is $F$-succinct in $L_{2}$ on $\mathcal{C}$ if there is a function $f \in F$ such that for every formula $\varphi_{1} \in L_{1}$ there is a formula $\varphi_{2} \in L_{2}$ of size $\left\|\varphi_{2}\right\| \leqslant f\left(\left\|\varphi_{1}\right\|\right)$ which is equivalent to $\varphi_{1}$ on all structures in $\mathcal{C}$. $\square$

Intuitively, a logic $L_{1}$ being $F$-succinct in a logic $L_{2}$ means that $F$ gives an upper bound for the size of $L_{1}$-formulas needed to express all of $L_{2}$. This definition may seem slightly at odds with the common use of the term "succinctness" in statements such as " $L_{2}$ is exponentially more succinct than $L_{1}$ " meaning that there is some $L_{2}$-formula that is not equivalent to any $L_{1}$-formula of subexponential size. In our terminology, we would rephrase this last statement as " $L_{1}$ is not $2^{o(n)}$-succinct in $L_{2}$ " (here we interpret subexponential as $2^{o(n)}$, but of course this is not the issue). The reason for defining $F$-succinctness the way we did is that it makes the formal statements of our results much more convenient. We will continue to use statements such as " $L_{2}$ is exponentially more succinct than $L_{1}$ " in informal discussions.

Example 2.2. MLFP is $\mathcal{O}(m)$-succinct in MSO on the class of all finite structures, because every formula $\left[\mathrm{LFP}_{x, X} \varphi(x, X, \bar{y}, \bar{Y})\right](z)$ is equivalent to $\forall X(X z \vee \exists x \neg X x \wedge \varphi(x, X, \bar{y}, \bar{Y}))$.

Example 2.3. MLFP is $\mathcal{O}(m)$-succinct in the fragment of MLFP whose formulas are in negation normal form on the class of all structures.

This allows us to usually work with formulas in negation normal form.

## 3 From MSO to MLFP

By the standard translation from MSO-logic to tree automata (cf., e.g., [Tho96]) one knows that every MSO-sentence $\Phi$ can be translated into a nondeterministic tree automaton with $\operatorname{Tower}(\mathcal{O}(\|\Phi\|))$ states that accepts exactly those labelled trees that satisfy $\Phi$. This leads to

Theorem 3.1 (Folklore). MSO-sentences are Tower $(\mathcal{O}(m))$-succinct in MLFP on the class of all labelled trees.

To show that we cannot do essentially better, i.e., that there is no translation from MSO to MLFP of size $\operatorname{Tower}(o(m))$ we need a complexity theoretic assumption that, however, does not seem to be too far-fetched. Let SAT denote the NP-complete satisfiability problem for propositional formulas in conjunctive normal form. Until now, all known deterministic algorithms that solve SAT have worstcase complexity $2^{\Omega(n)}$ (cf., [DGH $\left.{ }^{+} 02\right]$ ). Although not answering the P vs. NP question, the exposition of a deterministic algorithm for SAT with worst-case complexity $\leqslant n^{\lg n}$ would be a surprising and unexpected breakthrough in the SAT-solving community.

In the following, we write $\lg ^{(i)}$ to denote the $i$ times iterated logarithm, inductively defined by $\lg ^{(1)}(n):=\lg (n)$ and $\lg ^{(i+1)}(n):=\lg \left(\lg ^{(i)}(n)\right)$. Moreover, we we write $\lg ^{*}$ to denote the "inverse" of the Tower function, that is, the (unique) integer valued function with $\operatorname{Tower}\left(\lg ^{*}(n)-1\right)<n \leqslant$ $\operatorname{Tower}\left(\lg ^{*}(n)\right)$.

Theorem 3.2. Unless SAT is solvable by a deterministic algorithm that has, for every $i \in \mathbb{N}$, time bound $\|\gamma\| \|^{\mathbf{l g}^{(i)}(n)}$ (where $\gamma$ is the input formula and $n$ the number of propositional variables occurring in $\gamma$ ), MSO is not Tower $(o(m))$-succinct in MLFP on the class of all finite strings.

The overall proof idea is to assume that the function $f$ specifies the size of the translation from MSO to MLFP and to exhibit a SAT-solving algorithm which

- constructs a string $w$ that represents the SAT-instance $\gamma$,
- constructs an MSO-formula $\Phi(z)$ of extremely small size that, when evaluated in $w$, specifies a canonical satisfying assignment for $\gamma$ (if $\gamma$ is satisfiable at all),
- tests, for all MLFP-formulas $\Psi(z)$ of size $\leqslant f(\|\Phi\|)$, whether $\Psi$ specifies a satisfying assignment for $\gamma$.
Before presenting the proof in detail we provide the necessary notations and lemmas:
It is straightforward to see
Lemma 3.3. There is an algorithm that, given an MLFP-formula $\Psi(z)$, a string $w$, and a position $p$ in $w$, decides in time $|w|^{\mathcal{O}(| | \Psi| |)}$ whether $w \models \Psi(p)$.
Let us now concentrate on the construction of a string $w$ that represents a SAT-instance $\gamma$ and of an MSO-formula $\Phi(z)$ that specifies a canonical satisfying assignment of $\gamma$ (provided that $\gamma$ is satisfiable at all). Since we want $\Phi$ to be extremely short, we cannot choose $w$ to be the straightforward stringrepresentation of $\gamma$. Instead, we use the following, more complicated, representation of [FG03]:

For all $h \geqslant 1$ let $\Sigma_{h}:=\{0,1,\langle 1\rangle,</ 1\rangle, . .,\langle h>,</ h>\}$. The "tags" $<i>$ and $</ i>$ represent single letters of the alphabet and are just chosen to improve readability. For every $n \geqslant 1$ let $L(n)$ be the length of the binary representation of the number $n-1$, i.e., $L(0)=0, L(1)=1$, and $L(n)=$ $\lfloor\lg (n-1)\rfloor+1$, for all $n \geqslant 2$. By $\operatorname{bit}(i, n)$ we denote the $i$-th bit of the binary representation of $n$, i.e., $\operatorname{bit}(i, n)$ is 1 if $\left\lfloor\frac{n}{2^{i}}\right\rfloor$ is odd, and $\operatorname{bit}(i, n)$ is 0 otherwise.
We encode every number $n \in \mathbb{N}$ by a string $\mu_{h}(n)$ over the alphabet $\Sigma_{h}$, where $\mu_{h}(n)$ is inductively defined as follows: $\left.\mu_{1}(0):=<1\right\rangle\langle/ 1\rangle$, and

$$
\left.\mu_{1}(n):=<1>\operatorname{bit}(0, n-1) \operatorname{bit}(1, n-1) \cdots \operatorname{bit}(L(n)-1, n-1)</ 1\right\rangle
$$

for $n \geqslant 1$. For $h \geqslant 2$ we let $\mu_{h}(0):=<\mathrm{h}></ \mathrm{h}>$ and

$$
\begin{aligned}
& \mu_{h}(n):=<\mathrm{h}> \\
& \mu_{h-1}(0) \operatorname{bit}(0, n-1) \\
& \mu_{h-1}(1) \operatorname{bit}(1, n-1) \\
& \mu_{h-1}(L(n)-1) \operatorname{bit}(L(n)-1, n-1)
\end{aligned}
$$

for $n \geqslant 1$. Here empty spaces and line breaks are just used to improve readability.
To encode a CNF-formula $\gamma$ by a string we use the alphabet

$$
\begin{aligned}
\Sigma_{h}^{\prime}:= & \Sigma_{h} \cup \\
& \{+,-,<l i t>,</ l i t>,<c l a u s e>,</ \text { clause>, <cnf>, </cnf>\} } \cup \\
& \{\star,<\text { ass }\rangle,</ \text { ass }>,<\text { val>, </val> }\}
\end{aligned}
$$

Let $i \in \mathbb{N}$ and let $X_{i}$ be a propositional variable. The literal $X_{i}$ is encoded by the string

$$
\mu_{h}\left(X_{i}\right):=<\text { lit }>\mu_{h}(i)+</ \text { lit }>,
$$

and the literal $\neg X_{i}$ is encoded by $\mu_{h}\left(\neg X_{i}\right):=<$ lit> $\mu_{h}(i)-</$ lit>. A clause $\delta:=\lambda_{1} \vee \cdots \vee \lambda_{r}$ of literals is encoded by

$$
\mu_{h}(\delta):=<\text { clause }>\mu_{h}\left(\lambda_{1}\right) \cdots \mu_{h}\left(\lambda_{r}\right)</ \text { clause }>.
$$

A CNF-formula $\gamma:=\delta_{1} \wedge \cdots \wedge \delta_{m}$ is encoded by the string

$$
\mu_{h}(\gamma):=<\mathrm{cnf}>\mu_{h}\left(\delta_{1}\right) \cdots \mu_{h}\left(\delta_{m}\right)</ \mathrm{cnf}>
$$

We write $\operatorname{CNF}(n)$ to denote the class of all CNF-formulas the propositional variables of which are among $X_{0}, \ldots, X_{n-1}$. To provide the "infrastructure" for specifying a truth assignment to the variables $X_{0}, \ldots, X_{n-1}$, we use the string

$$
\begin{aligned}
& \mu_{h}\left(X_{0}, \ldots, X_{n-1}\right):=\text { <ass }> \\
& \\
& \text { <val> } \mu_{h}(0) \star</ \text { val> } \\
& \text { <val> } \mu_{h}(1) \star</ \text { val> } \\
& \vdots \\
& \\
& \text { <val> } \mu_{h}(n-1) \star</ \text { val> } \\
& \text { </ass }>
\end{aligned}
$$

Remark 3.4. There is a 1-1-correspondence between assignments $\alpha:\left\{X_{0}, . ., X_{n-1}\right\} \rightarrow\{$ true,false $\}$, on the one hand, and sets $P$ of positions of $\mu_{h}\left(X_{0}, \ldots, X_{n-1}\right)$ that carry the letter $\star$, on the other hand: Such a set $P$ specifies the assignment $\alpha^{P}$ that, for each $i<n$, maps the variable $X_{i}$ to the value true iff the $\star$-position directly after the substring $\mu_{h}(i)$ in $\mu_{h}\left(X_{0}, . ., X_{n-1}\right)$ belongs to $P$. Conversely, a given assignment $\alpha$ specifies the set $P^{\alpha}$ consisting of exactly those $\star$-positions of $\mu_{h}\left(X_{0}, \ldots, X_{n-1}\right)$ that occur directly after a substring $\mu_{h}(i)$ where $\alpha\left(X_{i}\right)=$ true.

Finally, we encode a formula $\gamma \in \operatorname{CNF}(n)$ by the string

$$
\mu_{h}(\gamma, \star):=\mu_{h}(\gamma) \mu_{h}\left(X_{0}, \ldots, X_{n-1}\right)
$$

$\mu_{h}(\gamma, \star)$ is the string $w$ that we will furtheron use as the representative of a SAT-instance $\gamma$. We use the following result of [FG03]:

## Lemma 3.5.

(a) There is an algorithm that, given $h \in \mathbb{N}$ and $\gamma \in C N F(n)$, computes (a binary representation of) the string $\mu_{h}(\gamma, \star)$ in time $\mathcal{O}\left(h \cdot(\lg h) \cdot(\lg n)^{2} \cdot(\|\gamma\|+n)\right)$ (cf., [FG03, Lemma 9]).
The string $\mu_{h}(\gamma, \star)$ has length $\left|\mu_{h}(\gamma, \star)\right|=\mathcal{O}\left(h \cdot(\lg n)^{2} \cdot(\|\gamma\|+n)\right)$.
(b) There is an algorithm that, given $h \in \mathbb{N}$, computes (the binary representation of) a $\mathrm{FO}(<)$-formula $\varphi_{h}(Z)$ in time $\mathcal{O}(h \cdot \lg h)$, such that for all $n \leqslant \operatorname{Tower}(h)$, for all $\gamma \in C N F(n)$, and for all sets $P$ of $\star$-positions in the string $\mu_{h}(\gamma, \star)$ we have

$$
\mu_{h}(\gamma, \star) \models \varphi_{h}(P) \quad \text { iff } \quad \alpha^{P} \text { is a satisfying assignment for } \gamma
$$

(cf., [FG03, Lemma 10]). The formula $\varphi_{h}(Z)$ has $\operatorname{siz}^{2}\left\|\varphi_{h}(Z)\right\|=\mathcal{O}(h)$.
Given a $\operatorname{CNF}(n)$-formula $\gamma$ and its representative $\mu_{h}(\gamma, \star)$, we now specify a canonical satisfying assignment of $\gamma$, provided that $\gamma$ is satisfiable at all. As observed in Remark 3.4, every assignment $\alpha:\left\{X_{0}, \ldots, X_{n-1}\right\} \rightarrow\{$ true, false $\}$ corresponds to a set $P^{\alpha}$ of positions in $\mu_{h}(\gamma, \star)$ that carry the letter $\star . P^{\alpha}$, again, can be identified with the $0-1$-string of length $\left|\mu_{h}(\gamma, \star)\right|$ that carries the letter 1 exactly at those positions that belong to $P^{\alpha}$. Now, the lexicographic ordering of these strings gives us a linear ordering on the set of all assignments $\alpha:\left\{X_{0}, \ldots, X_{n-1}\right\} \rightarrow\{$ true, false $\}$. As the canonical satisfying assignment of $\gamma$ we choose the lexicographically smallest satisfying assignment.

Lemma 3.6. There is an algorithm that, given $h \in \mathbb{N}$, computes (the binary representation of) an MSO-formula $\Phi_{h}(z)$ in time $\mathcal{O}(h \cdot \lg h)$, such that for all $n \leqslant \operatorname{Tower}(h)$, for all $\gamma \in C N F(n)$, and for all positions $p$ of $\mu_{h}(\gamma, \star)$ that carry the letter $\star$, we have

$$
\begin{aligned}
\mu_{h}(\gamma, \star) \models \Phi_{h}(p) \text { iff } \begin{array}{l}
\text { in the lexicographically smallest satisfying assignment } \\
\text { for } \gamma \text {, the propositional variable corresponding to posi- } \\
\text { tion } p \text { is assigned the value true. }
\end{array}
\end{aligned}
$$

The formula $\Phi_{h}(z)$ has size $\left\|\Phi_{h}\right\|=\mathcal{O}(h)$.
Finally, we are ready for the Proof of Theorem 3.2:

## Proof of Theorem 3.2.

Let $f: \mathbb{N} \rightarrow \mathbb{N}$ be a function such that there is, for every MSO-formula $\Phi(z)$, a MLFP-formula $\Psi(z)$ of size $\|\Psi\| \leqslant f(\|\Phi\|)$ which defines the same query as $\Phi$ on the class of all finite strings (recall that such an $f$ does indeed exist, because MSO and MLFP have the same expressive power over the class of finite strings).
Consider the algorithm displayed in Figure 1, which decides if the input formula $\gamma$ is satisfiable. The correctness of this algorithm directly follows from Lemma 3.6 and from the fact that at least one of the formulas $\Psi(z)$ of size $\leqslant f\left(\left\|\Phi_{h}\right\|\right)$ defines the same query as $\Phi_{h}(z)$.

[^2]```
Input: a SAT-instance }\gamma\mathrm{ in CNF
1. Count the number }n\mathrm{ of propositional variables occurring in }\gamma\mathrm{ , and
    modify }\gamma\mathrm{ in such a way that only the propositional variables }\mp@subsup{X}{0}{},..,\mp@subsup{X}{n-1}{}\mathrm{ occur in it.
2. Compute h:= 矢*}(n)\mathrm{ , i.e., choose }h\in\mathbb{N}\mathrm{ such that Tower(h-1)<n}\leqslant\mathrm{ Tower(h).
3. Construct the string }\mp@subsup{\mu}{h}{}(\gamma,\star)\mathrm{ that represents }\gamma\mathrm{ (see Lemma }3.5\mathrm{ (a)).
4. Construct an MSO-formula }\mp@subsup{\Phi}{h}{}(z)\mathrm{ that has the following property:
    Whenever }p\mathrm{ is a position in }\mp@subsup{\mu}{h}{}(\gamma,\star)\mathrm{ that carries the letter }\star\mathrm{ , we have
\mu}(\gamma,\star)\models\mp@subsup{\Phi}{h}{}(p)\mathrm{ iff in the lexicographically smallest satisfying assignment
                                    for }\gamma\mathrm{ , the propositional variable corresponding to posi-
                                    tion p}\mathrm{ is assigned the value true
    (cf., Lemma 3.6).
5. For all MLFP-formulas }\Psi(z)\mathrm{ of size |}||\leqslantf(|\mp@subsup{\Phi}{h}{}|)\mathrm{ do:
    (a) Initialise the assignment }\alpha:=\emptyset
    (b) For all positions p in }\mp@subsup{\mu}{h}{}(\gamma,\star)\mathrm{ that carry the letter }\star\mathrm{ do
        check whether }\mp@subsup{\mu}{h}{}(\gamma,\star)\models\Psi(p)
        if so, then insert the propositional variable corresponding to p into \alpha.
    (c) Check whether }\alpha\mathrm{ is a satisfying assignment for }\gamma\mathrm{ ;
        if so, then STOP with output " }\gamma\mathrm{ is satisfiable via assignment }\alpha\mathrm{ ".
6. STOP with output " }\gamma\mathrm{ is not satisfiable".
```

Fig. 1. A SAT-solving algorithm.

It remains to determine the worst-case running time of the algorithm. Let $\gamma$ be an input CNFformula for the algorithm, let $n$ be the number of propositional variables of $\gamma$, and let $h:=\lg ^{*}(n)$. The steps $1-4$ of the algorithm will be performed within a number of steps polynomial in $\|\gamma\|$, and the MSO-formula $\Phi_{h}(z)$ produced in step 4 will have size $\left\|\Phi_{h}\right\| \leqslant c \cdot h$, for a suitable constant $c \in \mathbb{N}$ (cf., Lemma 3.5 (a) and Lemma 3.6).
The loop in step 5 will be performed for $2^{c_{1} \cdot f\left(\left\|\Phi_{h}\right\|\right) \cdot \lg \left(f\left(\left\|\Phi_{h}\right\|\right)\right)}$ times, for a suitable constant $c_{1} \in \mathbb{N}$. To see this, note that formulas of length $\leqslant f\left(\left\|\Phi_{h}\right\|\right)$ use at most $f\left(\left\|\Phi_{h}\right\|\right)$ different first-order variables and at most $f\left(\left\|\Phi_{h}\right\|\right)$ different set variables. I.e., these formulas can be viewed as strings of length $f\left(\left\|\Phi_{h}\right\|\right)$ over an alphabet of size $c_{2}+2 \cdot f\left(\left\|\Phi_{h}\right\|\right)$, for a suitable constant $c_{2} \in \mathbb{N}$. Therefore, the number of such formulas is $\leqslant\left(c_{2}+2 \cdot f\left(\left\|\Phi_{h}\right\|\right)\right)^{f\left(\left\|\Phi_{h}\right\|\right)} \leqslant 2^{c_{1} \cdot f\left(\left\|\Phi_{h}\right\|\right) \cdot \lg \left(f\left(| | \Phi_{h} \|\right)\right)}$.
Each performance of the loop in step 5 will take a number of steps polynomial in

$$
\left|\mu_{h}(\gamma, \star)\right|^{\mathcal{O}\left(f\left(\left\|\Phi_{h}\right\|\right)\right)} \leqslant\left(c_{3} \cdot h \cdot(\lg n)^{2} \cdot\|\gamma\|\right)^{c_{4} \cdot f(c \cdot h)},
$$

for suitable constants $c_{3}, c_{4} \in \mathbb{N}$ (cf., Lemma 3.3 and Lemma 3.5 (a)). Altogether, for suitable constants $c, d \in \mathbb{N}$, the algorithm will perform the steps $1-6$ within $\|\gamma\|^{d \cdot f(c \cdot h) \cdot \lg (f(c \cdot h))}$ steps.

Now let us suppose that $f$ has bound $f(m) \leqslant \operatorname{Tower}(o(m))$. From Lemma 3.7 below we then obtain that our SAT-solving algorithm has, for every $i \in \mathbb{N}$, time bound $\|\gamma\|^{\lg ^{(i)}(n)}$. This finally completes the proof of Theorem 3.2.

Lemma 3.7. Let $f: \mathbb{N} \rightarrow \mathbb{N}$ be a function with bound $f(m) \leqslant \operatorname{Tower}(o(m))$, and let $c, d \in \mathbb{N}$. For every $i \in \mathbb{N}$ there is an $n_{0} \in \mathbb{N}$ such that for all $n \geqslant n_{0}$ we have

$$
d \cdot f\left(c \cdot \lg ^{*}(n)\right) \cdot \lg \left(f\left(c \cdot \lg ^{*}(n)\right)\right) \leqslant \lg ^{(i)}(n) .
$$

## 4 The Two-Variable Fragment of MLFP and the Full Modal $\mu$-Calculus

Defining the 2-variable fragment of MLFP requires some care, as the following theorem illustrates. To the best of our knowledge, this theorem has not appeared in the literature before, although other
researchers have noticed that just restricting least fixed-point logic to formulas with at most two firstorder variables does not yield a logic that behaves as a "proper 2-variable logic" [GOR99,GO99,Dzi96]. Note that the following theorem holds for arbitrary finite structures and not just for trees.

Theorem 4.1. Let $\tau$ be a signature that consists of relation symbols of arity $\leqslant 2$. Every $\operatorname{MLFP}(\tau)-$ sentence $\varphi$ is equivalent to an $\operatorname{MLFP}(\tau)$-sentence $\varphi^{\prime}$ of size $\left\|\varphi^{\prime}\right\| \in \mathcal{O}(\|\varphi\|)$ with just two first-order variables.

The following restriction MLFP ${ }^{2}$ of MLFP is more interesting. MLFP ${ }^{2}$ is the monadic version of the logic $\mathrm{FP}^{2}$ considered by [GOR99,GO99], which is denoted by $\widehat{\mathrm{FP}^{2}}$ in [Dzi96].
Definition 4.2 (MLFP $^{\mathbf{2}}$ ). An MLFP-formula is strict if no parameters are allowed in the fixed points, that is, if for all subformulas of the form $\left[L F P_{x, X} \varphi\right](y), x$ is the only free first-order variable of $\varphi$.

We define 2-variable monadic least fixed point logic, denoted by MLFP ${ }^{2}$ to be the restriction of MLFP to strict formulas with at most 2 first-order variables.
We first note that MLFP ${ }^{2}$, and actually $\mathrm{FO}^{2}$, the two variable fragment of first-order logic, is doubly exponentially more succinct than nondeterministic automata on the class of all finite strings:
Example 4.3. Let $\sigma:=\left\{L, R, P_{1}, \ldots, P_{n}\right\}$ and

$$
\varphi_{n}:=\forall x\left(L x \rightarrow \exists y\left(R y \wedge \bigwedge_{i=1}^{n}\left(P_{i} x \leftrightarrow P_{i} y\right)\right)\right)
$$

We claim that every nondeterministic finite automaton accepting precisely those strings over alphabet $2^{\sigma}$ that satisfy $\varphi$ has at least $2^{2^{n}}$ states. To see this, for every $S \subseteq 2^{\{1, . ., n\}}$, we define strings $X_{n}(S)$ and $Y_{n}(S)$ such that

- $L^{X_{n}(S)}=\mathcal{U}^{X_{n}(S)}$ and $R^{Y_{n}(S)}=\mathcal{U}^{Y_{n}(S)}$
- For all $s \in S$ there exists an $x \in \mathcal{U}^{X_{n}(S)}$ and an $y \in \mathcal{U}^{Y_{n}(S)}$ such that $s=\left\{i \mid x \in P_{i}^{X_{n}(S)}\right\}=\left\{i \mid y \in P_{i}^{Y_{n}(S)}\right\}$.
Let $W_{n}(S, T):=X_{n}(S) Y_{n}(T)$ be the concatenation of $X_{n}(S)$ and $Y_{n}(T)$. Then $W_{n}(S, T) \models$ $\varphi \Longleftrightarrow S \subseteq T$. Clearly, a nondeterministic finite automaton accepting precisely those strings $W_{n}(S, T)$ with $\bar{S} \subseteq T$ needs at least $2^{2^{n}}$ states.
Let us return to binary trees now. Following Vardi [Var98], we define the full modal $\mu$-calculus $\mathrm{FL}_{\mu}$ on binary trees as follows:

Definition $4.4\left(\mathbf{F L}_{\mu}\right)$. For each schema $\sigma$, an $\mathrm{FL}_{\mu}$-formula of schema $\sigma$ is either:

- true, false, $P$, or $\neg P$, where $P \in \sigma \cup\left\{\right.$ Root, Has-No-1 ${ }^{\text {st }}$ Child, Has-No- $2^{\text {nd }}$ Child $\}$;
- $\Phi_{1} \wedge \Phi_{2}$ or $\Phi_{1} \vee \Phi_{2}$, where $\Phi_{1}$ and $\Phi_{2}$ are $\mathrm{FL}_{\mu}$-formulas of schema $\sigma$;
- X, where $X$ is a propositional variable;
$-\langle R\rangle \Phi$ or $[R] \Phi$, where $R \in\left\{1^{\text {st }}\right.$ Child, $2^{\text {nd }}$ Child, $1^{\text {st }}$ Child $^{-1}, 2^{\text {nd }}$ Child $\left.^{-1}\right\}$ and $\Phi$ is an $\mathrm{FL}_{\mu^{-}}$ formula of schema $\sigma$;
- $\mu X . \Phi$ or $\nu X . \Phi$, where $X$ is a propositional variable and $\Phi$ is an $\mathrm{FL}_{\mu}$-formula of schema $\sigma$.

Instead of formally defining the semantics of $\mathrm{FL}_{\mu}$, we just give a translation of $\mathrm{FL}_{\mu}$ into MLFP ${ }^{2}$. For a $\sigma$-labelled tree $T$ and a node $t \in \mathcal{U}^{T}$, we write $(T, t) \models \Phi$ to denote that the $\mathrm{FL}_{\mu}$-formula $\Phi$ holds at $t$ in $T$. We identify propositional variables in $\mathrm{FL}_{\mu}$-formulas and set variables in MLFP ${ }^{2}$-formulas. For every $\mathrm{FL}_{\mu}$-formula $\Phi$ we define an $\operatorname{MLFP}^{2}$-formula $\varphi(x)$ in such a way that for all trees $T$ and nodes $t \in \mathcal{U}^{T}$ we have $(T, t) \models \Phi \Longleftrightarrow T \models \varphi(t)$.

- If $\Phi=$ true then $\varphi(x):=(x=x) . \quad$ If $\Phi=$ false then $\varphi(x):=\neg(x=x)$. If $\Phi=P$ then $\varphi(x):=P(x) . \quad$ If $\Phi=\neg P$ then $\varphi(x):=\neg P(x)$.
- If $\Phi=X$ then $\varphi(x):=X(x)$.
- If $\Phi=\Phi_{1} \vee \Phi_{2}$ then $\varphi(x):=\varphi_{1}(x) \vee \varphi_{2}(x)$.

If $\Phi=\Phi_{1} \wedge \Phi_{2}$ then $\varphi(x):=\varphi_{1}(x) \wedge \varphi_{2}(x)$.

- If $\Phi:=\langle R\rangle . \Psi$ for $R \in\left\{1^{\text {st }}\right.$ Child, $2^{\text {nd }}$ Child $\}$ then $\varphi(x):=\exists y(R(x, y) \wedge \psi(y))$, where $\psi(y)$ is obtained from $\psi(x)$ by simultaneously replacing all occurrences of $x$ by $y$ and vice-versa. Similarly, if $\Phi:=\left\langle R^{-1}\right\rangle . \Psi$ for $R \in\left\{1^{\text {st }}\right.$ Child, $2^{n d}$ Child $\}$ then $\varphi(x):=\exists y(R(y, x) \wedge \psi(y))$.
- If $\Phi:=[R] . \Psi$ for $R \in\left\{1^{\text {st }}\right.$ Child, $2^{\text {nd }}$ Child $\}$ then $\varphi(x):=\forall y(R(x, y) \rightarrow \psi(y))$. Similarly, if $\Phi:=\left[R^{-1}\right] . \Psi$ for $R \in\left\{1^{\text {st }}\right.$ Child, $2^{\text {nd }}$ Child $\}$ then $\varphi(x):=\forall y(R(y, x) \rightarrow \psi(y))$.
- If $\Phi:=\mu X . \Psi$ then $\varphi(x):=\left[\operatorname{LFP}_{x, X} \psi(x)\right](x)$.

If $\Phi:=\nu X . \Psi$ then $\varphi(x):=\left[\operatorname{GFP}_{x, X} \psi(x)\right](x)$.
What we have seen above is that $\mathrm{FL}_{\mu}$ is $\mathcal{O}(m)$-succinct in MLFP ${ }^{2}$. Our next result is that there also is a reverse translation from $\mathrm{MLFP}^{2}$ to $\mathrm{FL}_{\mu}$ which only incurs an exponential blow-up in size:
Theorem 4.5. MLFP ${ }^{2}$ is $2^{\text {poly }(m)}$-succinct in $\mathrm{FL}_{\mu}$ on the class of labelled trees. More precisely: There is a number $c \in \mathbb{N}$ such that for every $\operatorname{MLFP}^{2}$-formula $\varphi(x)$ there is an $\mathrm{FL}_{\mu}$-formula $\Phi$ of size $\|\Phi\| \leqslant$ $2^{\|\varphi\|^{c}}$ such that for all labelled trees $T$ and all nodes $t \in \mathcal{U}^{T}, T \models \varphi(t)$ iff $(T, t) \models \Phi$.

Theorem 4.6 (Vardi [Var98]). For every formula $\Phi$ of the full modal $\mu$-calculus $\mathrm{FL}_{\mu}$ there is a nondeterministic tree automaton of size $2^{\text {poly }(| | \Phi| |)}$ that accepts exactly those labelled trees in which $\Phi$ holds at the root.

As a matter of fact, Vardi [Var98] proved a stronger version of this theorem for infinite trees and parity tree automata. But on finite trees, a parity acceptance condition can always be replaced by a normal acceptance for finite tree automata.
The Theorems 4.5 and 4.6 directly imply
Corollary 4.7. For every $\mathrm{MLFP}^{2}$-formula $\varphi(x)$ there is a nondeterministic tree automaton of size $2^{2^{\text {pobl }(\|\varphi\|)}}$ that accepts exactly those labelled trees in which $\varphi$ holds at the root.

## 5 Simultaneous Least Fixed Point Logic

In this section we consider the simultaneous least fixed point operator which is defined as follows (cf., e.g., the textbook [EF99]):

Let $\tau$ be a signature and let $\varphi_{1}\left(x_{1}, X_{1}, \ldots, X_{n}\right), \ldots, \varphi_{n}\left(x_{n}, X_{1}, \ldots, X_{n}\right)$ be formulas over the signature $\tau$, each of which is positive in all the variables $X_{1}, \ldots, X_{n}$. For every $\tau$-structure $\mathcal{A}$ every formula $\varphi_{i}$ defines a monotone operator $F_{\mathcal{A}, \varphi_{i}}:\left(2^{\mathcal{U}^{\mathcal{A}}}\right)^{n} \rightarrow 2^{\mathcal{U}^{\mathcal{A}}}$ via $F_{\mathcal{A}, \varphi_{i}}\left(A_{1}, \ldots, A_{n}\right):=\left\{a \in \mathcal{U}^{\mathcal{A}}: \mathcal{A} \models\right.$ $\left.\varphi_{i}\left(a, A_{1}, \ldots, A_{n}\right)\right\}$. A tuple $\left(A_{1}, \ldots, A_{n}\right)$ is called a simultaneous fixed point of $\left(\varphi_{1}, \ldots, \varphi_{n}\right)$ in $\mathcal{A}$ iff, for all $i \leqslant n, F_{\mathcal{A}, \varphi_{i}}\left(A_{1}, \ldots, A_{n}\right)=A_{i}$.
For all $i \leqslant n$ and $s \in \mathbb{N}$ we define $L_{\mathcal{A}, \varphi_{i}}^{0}:=\emptyset$ and $L_{\mathcal{A}, \varphi_{i}}^{s+1}:=F_{\mathcal{A}, \varphi_{i}}\left(L_{\mathcal{A}, \varphi_{1}}^{s}, \ldots, L_{\mathcal{A}, \varphi_{n}}^{s}\right)$. Since $F_{\mathcal{A}, \varphi_{i}}$ is monotone we have $L_{\mathcal{A}, \varphi_{i}}^{s} \subseteq L_{\mathcal{A}, \varphi_{i}}^{s+1}$ for all $s \in \mathbb{N}$; and since $\mathcal{A}$ is finite a fixed point will be reached eventually, i.e., there is a $s_{0}$ such that $L_{\mathcal{A}, \varphi_{i}}^{s_{0}}=L_{\mathcal{A}, \varphi_{i}}^{s_{0}+1}=: L_{\mathcal{A}, \varphi_{i}}^{\infty}$ for all $i \leqslant n$. We define $\left(L_{\mathcal{A}, \varphi_{1}}^{\infty}, \ldots, L_{\mathcal{A}, \varphi_{n}}^{\infty}\right)$ to be the simultaneous least fixed point of $\left(\varphi_{1}, \ldots, \varphi_{n}\right)$ in $\mathcal{A}$. It is straightforward to see that the simultaneous least fixed point is included in every simultaneous fixed point $\left(A_{1}, \ldots, A_{n}\right)$ of $\left(\varphi_{1}, \ldots, \varphi_{n}\right)$ in $\mathcal{A}$, i.e., $L_{\mathcal{A}, \varphi_{i}}^{\infty} \subseteq A_{i}$, for all $i \leqslant n$ (cf., e.g., [EF99, Lemma 8.1.17]).

The logic $\operatorname{Sim}-\operatorname{MLFP}(\tau)$ is the extension of $\operatorname{MLFP}(\tau)$ by simultaneous least fixed point operators. I.e.: $\operatorname{Sim}-\operatorname{MLFP}(\tau)$ contains $\operatorname{MLFP}(\tau)$ and is closed under Boolean connectives and first-order quantifications; and if $n \in \mathbb{N}$ and $\varphi_{1}\left(x_{1}, X_{1}, \ldots, X_{n}, \bar{y}, \bar{Y}\right), \ldots, \varphi_{n}\left(x_{n}, X_{1}, \ldots, X_{n}, \bar{y}, \bar{Y}\right)$ are $\operatorname{Sim-MLFP}(\tau)-$ formulas each of which is positive in the variables $X_{1}, \ldots, X_{n}$ then, for every $i \leqslant n$,

$$
\left[\operatorname{Sim}^{\left.-\operatorname{LFP}_{x_{1}, X_{1}, . ., x_{n}, X_{n}} \varphi_{1}, \ldots, \varphi_{n}\right]_{X_{i}}(x)}\right.
$$

is a $\operatorname{Sim}-\operatorname{MLFP}(\tau)$-formula such that for every $(\tau \cup\{\bar{y}, \bar{Y}\})$-structure $\mathcal{A}$ and every element $a \in \mathcal{U}^{\mathcal{A}}$


Remark 5.1. It is straightforward to see that Sim-MLFP is $\mathcal{O}(m)$-succinct in MSO.

Every simultaneous least fixed point can be expressed by an MLFP-formula of exponential size:
Proposition 5.2. Let $\tau$ be a signature, let $n, s \in \mathbb{N}$, and let, for each $i \leqslant n, \varphi_{i}\left(x_{i}, X_{1}, \ldots, X_{n}\right)$ be an $\operatorname{MLFP}(\tau)$-formula of size $\left\|\varphi_{i}\right\| \leqslant s$ and positive in $X_{1}, \ldots, X_{n}$. For every $i \leqslant n$ there is a $\operatorname{MLFP}(\tau)$-formula $\Phi_{X_{i}}(x)$ of size $\left\|\Phi_{X_{i}}\right\| \leqslant s^{n}$ such that, for all $\tau$-structures $\mathcal{A}, \mathcal{A} \vDash \forall x \Phi_{X_{i}}(x) \leftrightarrow$ $\left[\operatorname{Sim}-\operatorname{LFP}_{x_{1}, X_{1}, . ., x_{n}, X_{n}} \varphi_{1}, \ldots, \varphi_{n}\right]_{X_{i}}(x)$.
Furthermore, if $\varphi_{1}, \ldots, \varphi_{n}$ are in MLFP ${ }^{2}$, then also $\Phi_{X_{i}}$ is, for every $i \leqslant n$.
Using Proposition 5.2 one easily obtains the following result by induction on the construction of Sim-MLFP-formulas:
Corollary 5.3. Sim-MLFP is $2^{2^{\mathcal{O}(m)}}$-succinct in MLFP on the class of all finite structures.
Together with Theorem 3.2 this leads to
Corollary 5.4. Unless SAT is solvable by a deterministic algorithm that has, for every $i \in \mathbb{N}$, time bound $\|\gamma\|^{\lg ^{(i)}(n)}$ (where $\gamma$ is the input formula and $n$ the number of propositional variables occurring in $\gamma$ ), MSO is not Tower $(o(m))$-succinct in Sim-MLFP on the class of all finite strings.

## 6 Monadic Datalog and Stratified Monadic Datalog

We assume that the reader is familiar with datalog, which may be viewed as logic programming without function symbols (cf., e.g., the textbook [AHV95]). A datalog program is monadic if all its IDBpredicates (i.e., its intensional predicates that appear in the head of some rule of the program) are unary. In this paper we restrict attention to monadic datalog programs that are interpreted over labelled trees. A monadic datalog program of schema $\sigma$ may use as EDB-predicates (i.e., extensional predicates which are determined by the structure the program is interpreted over) the predicates in $\tau_{\text {Trees }}$, the predicates in $\sigma$, and a predicate $\neg P$ for every $P \in \sigma$ which is interpreted as the complement of $P$. We use $\operatorname{IDB}(\mathcal{P})$ to denote the set of IDB-predicates of $\mathcal{P}$, and we write MonDatalog to denote the class of all monadic datalog programs.

More formally, a monadic datalog program $\mathcal{P}$ of schema $\sigma$ is a finite set of rules of the form $X(x) \leftarrow \gamma(x, \bar{y})$, where $\gamma$ is a conjunction of atomic formulas over the signature $\tau_{\text {Trees }} \cup \sigma \cup\{\neg P:$ $P \in \sigma\} \cup \operatorname{IDB}(\mathcal{P})$. We define the size $\|\mathcal{P}\|$ of $\mathcal{P}$ in the same way as we defined the size of formulas. Instead of formally defining the semantics of MonDatalog, we just give a translation of MonDatalog into Sim-MLFP on the class of all labelled trees:
Given a MonDatalog-program $\mathcal{P}$ of schema $\sigma$ with $\operatorname{IDB}(\mathcal{P})=\left\{X_{1}, \ldots, X_{n}\right\}$, we can assume w.l.o.g. that there is a unique first-order variable $x$ such that every rule in $\mathcal{P}$ is of the form $X_{i}(x) \leftarrow \gamma(x, \bar{y})$, for some $i \leqslant n$. Let $\gamma_{i, 1}(x, \bar{y}), \ldots, \gamma_{i, \ell_{i}}(x, \bar{y})$ be a list of all bodies of rules with head $X_{i}(x)$ in $\mathcal{P}$. Of course, the FO-formula $\varphi_{i}\left(x, X_{1}, \ldots, X_{n}\right):=\bigvee_{j=1}^{\ell_{i}} \exists \bar{y} \gamma_{i, j}(x, \bar{y})$ is positive in all the variables $X_{1}, \ldots, X_{n}$. When evaluated in a $\sigma$-labelled tree $T$, the program $\mathcal{P}$ defines the unary relations $\left(X_{i}\right)_{\mathcal{P}}^{\infty}(T)$, for $i \leqslant n$, to be the simultaneous least fixed point of $\left(\varphi_{1}, \ldots, \varphi_{n}\right)$ in $T$.

A MonDatalog-program $\mathcal{P}$ of schema $\sigma$, together with a designated goal predicate $X \in \operatorname{IDB}(\mathcal{P})$, defines the unary query which yields, for every $\sigma$-labelled tree $T$, the set $X_{\mathcal{P}}^{\infty}(T)$ of vertices of $T$. We say that $T$ belongs to the Boolean query defined by $(\mathcal{P}, X)$ iff the root of $T$ belongs to $X_{\mathcal{P}}^{\infty}(T)$.
Let us sum up what we have seen above:
Lemma 6.1. For every MonDatalog-program $\mathcal{P}$ with $\operatorname{IDB}(\mathcal{P})=\left\{X_{1}, \ldots, X_{n}\right\}$ there are FO-formulas $\varphi_{i}\left(x, X_{1}, \ldots, X_{n}\right)$, for every $i \in\{1, \ldots, n\}$, positive in $X_{1}, \ldots, X_{n}$, such that $\left\|\varphi_{i 1}\right\|+\ldots+\left\|\varphi_{n}\right\| \leqslant\|\mathcal{P}\|$ and $\left[\operatorname{Sim}-\operatorname{LFP}_{x, X_{1}, . ., x, X_{n}} \varphi_{1}, \ldots, \varphi_{n}\right]_{X_{i}}(x)$ defines the same unary query as $\left(\mathcal{P}, X_{i}\right)$ on the class of all labelled trees.

Gottlob and Koch [GK02] proposed the following useful normal form for monadic datalog:
Definition 6.2 (TMNF). A MonDatalog-program $\mathcal{P}$ of schema $\sigma$ is in tree marking normal form (TMNF, for short) iff each rule is of one of the four forms

$$
\text { 1. } \quad X(x) \leftarrow S(x)
$$

2. $\quad X(x) \leftarrow X^{\prime}(x) \wedge X^{\prime \prime}(x)$,
3. $X(x) \leftarrow R(x, y) \wedge X^{\prime}(y)$,
4. $\quad X(x) \leftarrow R(y, x) \wedge X^{\prime}(y)$, where
$X, X^{\prime}, X^{\prime \prime} \in \operatorname{IDB}(\mathcal{P}), R \in\left\{1^{\text {st }}\right.$ Child, $2^{\text {nd }}$ Child $\}$, and
$S \in \sigma \cup\{\neg P: P \in \sigma\} \cup\left\{\right.$ Root, Has-No-1 ${ }^{\text {st }}$ Child, Has-No- $2^{\text {nd }}$ Child $\}$.
Theorem 6.3 ([GK02, Theorem 4.11]). For every MonDatalog-program $\mathcal{P}$ and every $X \in \operatorname{IDB}(\mathcal{P})$ there is a TMNF-program $\mathcal{P}_{X}$ of size $\mathcal{O}(\|\mathcal{P}\|)$ with $X \in \operatorname{IDB}\left(\mathcal{P}_{X}\right)$ such that $\left(\mathcal{P}_{X}, X\right)$ defines the same unary query as $(\mathcal{P}, X)$ on the class of labelled trees.

One way of adding negation to datalog is to consider stratified datalog (cf., [AHV95]): A stratified monadic datalog program $\mathcal{P}$ of schema $\sigma$ is a finite set of rules of the form $X(x) \leftarrow \gamma(x, \bar{y})$, where $\gamma$ is a conjunction of atomic formulas over the signature $\tau_{\text {Trees }} \cup \sigma \cup\{\neg P: P \in \sigma\} \cup \operatorname{IDB}(\mathcal{P}) \cup\{\neg X$ : $X \in \operatorname{IDB}(\mathcal{P})\}$ that has the following property: There is a partition of $\mathcal{P}$ into sets $\mathcal{P}_{1}, \ldots, \mathcal{P}_{n}$, for some $n \in \mathbb{N}$, such that each $\mathcal{P}_{i}$ is a MonDatalog-program of schema $\sigma \cup \bigcup_{j<i} \operatorname{IDB}\left(\mathcal{P}_{j}\right)$. The programs $\mathcal{P}_{1}, \ldots, \mathcal{P}_{n}$ are called the strata of $\mathcal{P}$, and $\left(\mathcal{P}_{1}, \ldots, \mathcal{P}_{n}\right)$ is called a decomposition of $\mathcal{P}$ into strata.
We write S-MonDatalog to denote the class of all stratified monadic datalog programs.
In [GK02] it was shown that MonDatalog can define the same unary queries on the class of labelled trees as monadic second-order logic. In the remainder of this section we will compare the succinctness of MonDatalog, S-MonDatalog, $\mathrm{FL}_{\mu}$, MLFP, and a particular kind of tree automaton.

### 6.1 From MonDatalog to Finite Automata

Several mechanisms have been proposed in the literature for specifying unary queries by finite automata operating on labelled trees (cf. [NS02]). One such mechanism, introduced in [Nev99] and further investigated in [FGK03,Koc03], is the selecting tree automaton:

Definition 6.4 (STA). Let $\sigma$ be a schema. A selecting $\sigma$-tree automaton ( $\sigma$-STA, for short) is a tuple $\mathfrak{A}=\left(Q, 2^{\sigma}, F, \delta, S\right)$, where $S \subseteq Q$ is the set of selecting states and $\left(Q, 2^{\sigma}, F, \delta\right)$ is a conventional nondeterministic bottom-up tree automaton (cf., e.g., [Tho96]) with finite state space $Q$, input alphabet $2^{\sigma}$, accepting states $F \subseteq Q$, and transition function

$$
\delta: 2^{\sigma} \cup\left(\{1\} \times Q \times 2^{\sigma}\right) \cup\left(\{2\} \times Q \times 2^{\sigma}\right) \cup\left(Q \times Q \times 2^{\sigma}\right) \rightarrow 2^{Q}
$$

A run of $\mathfrak{A}$ on a $\sigma$-labelled tree $T$ is a mapping $\rho: \mathcal{U}^{T} \rightarrow Q$ that has the following property, for all vertices $t, t_{1}, t_{2} \in \mathcal{U}^{T}$ : if $t$ has no children then $\rho(t)=\delta(\operatorname{label}(t))$; if $1^{\text {st }} \operatorname{Child}\left(t, t_{1}\right) \wedge$ Has-No- $2^{\text {nd }}$ Child $(t)$ then $\rho(t) \in \delta\left(1, \rho\left(t_{1}\right)\right.$, label $\left.(t)\right)$; if $2^{\text {nd }} \operatorname{Child}\left(t, t_{2}\right) \wedge$ Has-No-1 ${ }^{\text {st }}$ Child $(t)$ then $\rho(t) \in \delta\left(2, \rho\left(t_{2}\right)\right.$, label $\left.(t)\right)$; if $1^{\text {st }}$ Child $\left(t, t_{1}\right) \wedge 2^{\text {nd }}$ Child $\left(t, t_{2}\right)$ then $\rho(t) \in \delta\left(\rho\left(t_{1}\right), \rho\left(t_{2}\right)\right.$, label $\left.(t)\right)$.

A run $\rho$ of $\mathfrak{A}$ on $T$ is said to be accepting if it maps the root of $T$ to a state in $F$. The unary query defined by $\mathfrak{A}$ is the query which maps every $\sigma$-labelled tree $T$ to the set of those vertices $t \in \mathcal{U}^{T}$ that satisfy the following condition: $\rho(t) \in S$ for every accepting run $\rho$ of $\mathfrak{A}$ on $T$.

It was shown in [FGK03,Nev99] that STAs can define exactly those unary queries on the class of labelled trees that are definable in monadic-second order logic.
Theorem 6.5 ([FGK03,GK02]). MonDatalog is $2^{\mathcal{O}(m)}$-succinct in STAs on the class of labelled trees. More precisely: Let $\mathcal{P}$ be a MonDatalog-program of schema $\sigma$, and let $X \in \operatorname{IDB}(\mathcal{P})$. There is a $\sigma$-STA $\mathfrak{A}$ with $2^{\mathcal{O}(\|\mathcal{P}\|)}$ states that defines the same unary query as $(\mathcal{P}, X)$ on the class of labelled trees. Given $\mathcal{P}$, the STA $\mathfrak{A}$ can be computed in time $2^{\mathcal{O}(\|\mathcal{P}\|)}$.
The next example shows that, asymptotically, the above construction is optimal:
Example 6.6. For every $k \in \mathbb{N}$ there is a MonDatalog-program of size $\mathcal{O}(k)$ that defines a query not definable by an STA with less than $2^{k}$ states:
Restricting attention to conventional nondeterministic finite automata (NFAs) that operate on finite strings, it is straightforward to see that there is no NFA $\mathfrak{A}$ with less than $2^{k}$ states such that, for all
$w_{1}, w_{2} \in\{0,1\}^{k}, \mathfrak{A}$ accepts the string $w_{1} 2 w_{2}$ if and only if $w_{1}=w_{2}$. On the other hand, the following MonDatalog-program $\mathcal{P}_{k}$ has size $\mathcal{O}(k)$ and has the property that, for all $w_{1}, w_{2} \in\{0,1\}^{k}$, $\left(X^{o k}\right)_{\mathcal{P}_{k}}^{\infty}\left(w_{1} 2 w_{2}\right)$ contains the maximum position of the string $w_{1} 2 w_{2}$ if and only if $w_{1}=w_{2}$. The program $\mathcal{P}_{k}$ consists of the following rules:

$$
\begin{aligned}
X_{0}^{\text {transfer }}(x) & \leftarrow P_{0}\left(y_{1}\right) \wedge \operatorname{Succ}\left(y_{1}, y_{2}\right) \wedge \cdots \wedge \operatorname{Succ}\left(y_{k-1}, y_{k}\right) \wedge \operatorname{Succ}\left(y_{k}, x\right) \\
X_{1}^{\text {transfer }}(x) & \leftarrow P_{1}\left(y_{1}\right) \wedge \operatorname{Succ}\left(y_{1}, y_{2}\right) \wedge \cdots \wedge \operatorname{Succ}\left(y_{k-1}, y_{k}\right) \wedge \operatorname{Succ}\left(y_{k}, x\right) \\
X^{\text {compare }}(x) & \leftarrow P_{0}(x) \wedge X_{0}^{\text {transfer }}(x) \\
X^{\text {compare }}(x) & \leftarrow P_{1}(x) \wedge X_{1}^{\text {transfer }}(x) \\
X^{o k}(x) & \leftarrow P_{2}(x) \\
X^{o k}(x) & \leftarrow X^{o k}(y) \wedge \operatorname{Succ}(y, x) \wedge X^{\text {compare }}(x) .
\end{aligned}
$$

### 6.2 From S-MonDatalog to MonDatalog

In this section we show that S-MonDatalog-programs can be translated into MonDatalog-programs of at most exponential size. It remains open if the exponential size is indeed necessary or if, on the contrary, for every S-MonDatalog-program $\mathcal{P}$ there exists an equivalent MonDatalog-program $\mathcal{P}^{\prime}$ of size polynomial in $\|\mathcal{P}\|$.

Lemma 6.7. For every $\sigma$-STA $\mathfrak{A}=\left(Q, 2^{\sigma}, F, \delta, S\right)$ there is a MonDatalog-program $\mathcal{P}$ of size $\mathcal{O}\left(|Q|^{3}\right.$. $\left.\left|2^{\sigma}\right|+|\sigma| \cdot\left|2^{\sigma}\right|\right)$ and a designated goal predicate $\bar{X} \in \operatorname{IDB}(\mathcal{P})$ such that $(\mathcal{P}, \bar{X})$ defines the complement of the query defined by $\mathfrak{A}$ on the class of all $\sigma$-labelled trees.

Using Theorem 6.5 and Lemma 6.7 one easily obtains
Proposition 6.8. For every MonDatalog-program $\mathcal{P}$ there is a MonDatalog-program $\mathcal{P}^{\prime}$ of size $2^{\mathcal{O}(\|\mathcal{P}\|)}$ with $\operatorname{IDB}\left(\mathcal{P}^{\prime}\right) \supseteq\{X, \bar{X}: X \in \operatorname{IDB}(\mathcal{P})\}$ such that, for every $X \in \operatorname{IDB}(\mathcal{P}),\left(\mathcal{P}^{\prime}, X\right)$ defines the same unary query as $(\mathcal{P}, X)$, and $\left(\mathcal{P}^{\prime}, \bar{X}\right)$ defines the complement of the unary query defined by $(\mathcal{P}, X)$ on the class of labelled trees.

Using the above proposition, it is not difficult to prove
Theorem 6.9. S-MonDatalog is $2^{\mathcal{O}(m)}$-succinct in MonDatalog on the class of labelled trees. More precisely: For every S-MonDatalog-program $\mathcal{P}$ there is a MonDatalog-program $\mathcal{P}^{\prime}$ of size $2^{\mathcal{O}(\|\mathcal{P}\|)}$ such that $\operatorname{IDB}\left(\mathcal{P}^{\prime}\right) \supseteq \operatorname{IDB}(\mathcal{P})$ and, for every $X \in \operatorname{IDB}(\mathcal{P}),\left(\mathcal{P}^{\prime}, X\right)$ defines the same unary query as $(\mathcal{P}, X)$ on the class of all labelled trees.
Given $\mathcal{P}$, the program $\mathcal{P}^{\prime}$ can be computed in time $2^{\mathcal{O}(\|\mathcal{P}\|)}$.

### 6.3 S-MonDatalog vs $\mathrm{FL}_{\mu}$

From Theorem 4.6 and Lemma 6.7 one directly obtains
Theorem 6.10. $\mathrm{FL}_{\mu}$ is $2^{\text {poly }(m)}$-succinct in $\mathrm{S}-\mathrm{MonDatalog}$ on the class of labelled trees. More precisely: For every $\mathrm{FL}_{\mu}$-formula $\Phi$ there is an S-MonDatalog-program $\mathcal{P}$ of size $2^{\text {poly }(| | \Phi \|)}$ and a predicate $X \in \operatorname{IDB}(\mathcal{P})$, such that, for all labelled trees $T$, the root of $T$ belongs to the unary query defined by $(\mathcal{P}, X)$ if, and only if, the root of $T$ satisfies $\Phi$.

Conversely, using Theorem 6.3, Lemma 6.1, and Proposition 5.2 one can show the following
Theorem 6.11. S-MonDatalog is $2^{\mathcal{O}(m \cdot \lg m)}$-succinct in $\mathrm{FL}_{\mu}$ on the class of labelled trees. More precisely: For every S-MonDatalog-program $\mathcal{P}$ and every $X \in \operatorname{IDB}(\mathcal{P})$, there is a $\mathrm{FL}_{\mu}$-formula $\Phi_{X}$ of size $2^{\mathcal{O}(\|\mathcal{P}\| \cdot \lg \|\mathcal{P}\|)}$ that defines the same unary query as $(\mathcal{P}, X)$ on the class of labelled trees.

It remains open whether the above bounds are optimal.

### 6.4 From MLFP to S-MonDatalog

Similarly to Theorem 3.1 one easily obtains
Theorem 6.12 (Folklore). MLFP-sentences are $\operatorname{Tower}(\mathcal{O}(m))$-succinct in S-MonDatalog on the class of labelled trees.

The aim of this section is to show that there are no essentially smaller translations from MLFP to S-MonDatalog. We will use the following well-known observation:
Proposition 6.13 (Folklore). There is no function $f: \mathbb{N} \rightarrow \mathbb{N}$ with bound $f(m) \leqslant \operatorname{Tower}(o(m))$ such that for every $\mathrm{FO}(<)$-sentence $\varphi$ there is a nondeterministic finite automaton $\mathfrak{A}$ with at most $f(\|\varphi\|)$ states that accepts exactly those strings that satisfy $\varphi$.
Using Proposition 6.13 and the results of the Sections 6.1 and 6.2, one obtains the following:
Theorem 6.14. There is no function $f: \mathbb{N} \rightarrow \mathbb{N}$ with bound $f(m) \leqslant \operatorname{Tower}(o(m))$ such that for every $\mathrm{FO}(<)$-sentence $\varphi$ there is a S-MonDatalog-program $\mathcal{P}$ of size $\|\mathcal{P}\| \leqslant f(\|\varphi\|)$ and a designated goal predicate $X \in \operatorname{IDB}(\mathcal{P})$ such that $(\mathcal{P}, X)$ defines the same Boolean query as $\varphi$ on the class of all finite strings.

Since $\mathrm{FO}(<)$ is included in MLFP, the above theorem directly implies the following:
Corollary 6.15. MLFP is not Tower $(o(m))$-succinct in S-MonDatalog on the class of all finite strings.

It remains open if this result remains valid when replacing MLFP with MLFP ${ }^{2}$. Note, however, that for the proof of Proposition 6.13 a small number $k$ of first-order variables suffices. I.e., Proposition 6.13 remains valid when replacing $\mathrm{FO}(<)$ with $\mathrm{FO}^{k}(<)$, and Corollary 6.15 remains valid when replacing MLFP with MLFP ${ }^{k}$.
Together with Corollary 4.7 and Lemma 6.7, the above Corollary 6.15 implies
Corollary 6.16. MLFP is not Tower $(o(m))$-succinct in MLFP $^{2}$ on the class of all finite strings.

## 7 Conclusion

We studied the succinctness of a number of fixed point logics on trees. We believe that the analysis of succinctness, which may be viewed as a refined, "quantitative" analysis of expressive power, is a very interesting topic that deserves much more attention.

Even though we were able to get a good overall picture of the succinctness of monadic fixed point logics on trees, a number of questions remain open. Let us just mention a few of them:

- The exact relationship between monadic datalog, stratified monadic datalog, and the full modal $\mu$-calculus remains unclear. In particular: Is the class of all queries whose complements can be defined by monadic datalog programs polynomially succinct in monadic datalog, or is there an exponential lower bound? (Recall that in Proposition 6.8 we prove an exponential upper bound.)
- Our proof that MSO is not Tower $(o(m))$-succinct in MLFP relies on a complexity theoretic assumption. Is it possible to prove this result without such an assumption?
- We have only considered the 2 -variable fragment of MLFP here. What about the $k$-variable fragments, for $k \geqslant 3$ ? Do they form a strict hierarchy with respect to succinctness?


## References

[ABS99] S. Abiteboul, P. Buneman, and D. Suciu. Data on the Web: From Relations to Semistructured Data and XML. Morgan Kaufmann, 1999.
[AHV95] Serge Abiteboul, Richard Hull, and Victor Vianu. Foundations of databases. Addison-Wesley, 1995.
[AIO0] N. Alechina and N. Immerman. Reachability logic: An efficient fragment of transitive closure logic. Logic Journal of the IGPL, 8(3):325-338, 2000.
[AI01] M. Adler and N. Immerman. An $n$ ! lower bound on formula size. In Proceedings of the 16th IEEE Symposium on Logic in Computer Science, pages 197-206, 2001.
[CH82] A. Chandra and D. Harel. Structure and complexity of relational queries. Journal of Computer and Systems Sciences, 25:99-128, 1982.
$\left[\mathrm{DGH}^{+} 02\right]$ E. Dantsin, A. Goerdt, E.A. Hirsch, R. Kannan, J. Kleinberg, C. Papadimitriou, P. Raghavan, and U. Schöning. A deterministic $(2-2 /(k+1))^{n}$ algorithm for $k$-SAT based on local search. Theoretical Computer Science, 289(1):69-83, 2002. Revised version of: Deterministic algorithms for $k$-SAT based on covering codes and local search, ICALP'00, LNCS volume 1853.
[Dzi96] S. Dziembowski. Bounded-variable fixpoint queries are PSPACE-complete. In D. van Dalen and M. Bezem, editors, Computer Science Logic, 10th International Workshop, CSL '96, volume 1258 of Lecture Notes in Computer Science, pages 89-105. Springer-Verlag, 1996.
[EF99] Heinz-Dieter Ebbinghaus and Jörg Flum. Finite model theory. Springer, New York, second edition, 1999.
[EVW02] K. Etessami, M. Y. Vardi, and Th. Wilke. First-order logic with two variables and unary temporal logic. Information and Computation, 179(2):279-295, 2002.
[Fag75] R. Fagin. Monadic generalized spectra. Zeitschrift für mathematische Logik und Grundlagen der mathematik, 21:89-96, 1975.
[FG03] Markus Frick and Martin Grohe. The complexity of first-order and monadic second-order logic revisited. Journal version of LICS'02 paper. Available at http://www.dcs.ed.ac.uk/home/ grohe/pub. html, 2003.
[FGK03] Markus Frick, Martin Grohe, and Christoph Koch. Query evaluation on compressed trees. In 18th IEEE Symposium on Logic in Computer Science (LICS'03), Ottawa, Canada, June 2003.
[FSW00] M. F. Fernandez, J. Siméon, and Ph. Wadler. An algebra for XML query. In S. Kapoor and S. Prasad, editors, Proceedings of the 20th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'00), volume 1974 of Lecture Notes in Computer Science, pages 11-45. Springer-Verlag, 2000.
[GK02] Georg Gottlob and Christoph Koch. Monadic datalog and the expressive power of web information extraction languages. Submitted, November 2002. Journal version of PODS'02 paper. Available as CoRR report arXiv:cs.DB/0211020.
[GO99] E. Grädel and M. Otto. On Logics with Two Variables. Theoretical Computer Science, 224:73-113, 1999.
[GOR99] E. Grädel, M. Otto, and E. Rosen. Undecidability Results on Two-Variable Logics. Archive for Mathematical Logic, 38:213-354, 1999.
[HP00] H. Hosoya and B. C. Pierce. XDuce: A typed XML processing language (preliminary report). In D. Suciu and G. Vossen, editors, International Workshop on the Web and Databases, 2000. Reprinted in The Web and Databases, Selected Papers, Springer LNCS volume 1997, 2001.
[Imm99] Neil Immerman. Descriptive complexity. Springer, New York, 1999.
[Kam68] H. Kamp. Tense Logic and the theory of linear order. PhD thesis, University of California, Los Angeles, 1968.
[Koc03] Christoph Koch. Efficient processing of expressive node-selecting queries on XML data in secondary storage: A tree-automata based approach. In Proceedings of the 29th Conference on Very Large Data Bases, 2003. To appear.
[Nev99] Frank Neven. Design and Analysis of Query Languages for Structured Documents - A Formal and Logical Approach. PhD thesis, Limburgs Universitair Centrum, 1999.
[NS02] Frank Neven and Thomas Schwentick. Query automata over finite trees. Theoretical Computer Science, 275(1-2):633-674, 2002. Journal version of PODS'00 paper.
[Tho96] Wolfgang Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, Handbook of formal languages, volume 3. Springer, New York, 1996.
[Var98] Moshe Y. Vardi. Reasoning about the past with two-way automata. In K.G. Larsen, S. Skyum, and G. Winskel, editors, 25th International Colloquium on Automata, Languages and Programming (ICALP'98), volume 1443 of Lecture Notes in Computer Science, pages 628-641. Springer-Verlag, 1998.
[Wil99] T. Wilke. CTL+ is exponentially more succinct than CTL. In C. P. Rangan, V. Raman, and R. Ramanujam, editors, Proceedings of the 19th Conference on Foundations of Software Technology and Theoretical Computer Science, volume 1738 of Lecture Notes in Computer Science, pages 110-121. Springer-Verlag, 1999.

## Appendix

## A Proofs omitted in Section 3

## Proof of Lemma 3.3.

It is straightforward to exhibit an algorithm $\mathcal{A}$ with time bound $|w|^{\mathcal{O}(\|\Psi\|)}$ such that the following is true: Given $k, \ell \in \mathbb{N}$, an MLFP-formula $\Psi\left(x_{1}, \ldots, x_{k}, X_{1}, \ldots, X_{\ell}\right)$, a string $w$, a sequence $p_{1}, \ldots, p_{k}$ of positions in $w$, and sets $P_{1}, \ldots, P_{\ell}$ of positions in $w, \mathcal{A}$ decides whether $w \vDash \Psi\left(p_{1}, \ldots, p_{k}, P_{1}, \ldots, P_{\ell}\right)$. $\mathcal{A}$ operates by recursion on the construction of $\Psi$. The only non-trivial case is when $\Psi$ is of the form $\left[\operatorname{LFP}_{y, Y} \Psi^{\prime}\left(y, Y, x_{1}, \ldots, x_{k}, X_{1}, \ldots, X_{\ell}\right)\right]\left(x_{i}\right)$. In this case, $\mathcal{A}$ computes the stages of the least fixed point of $\Psi^{\prime}$. I.e., $\mathcal{A}$ performs the following operations:

1. Initialise the set $L^{0}:=\emptyset$.
2. For $s:=1$ to $|w|$ do
(a) Initialise the set $L^{s}:=\emptyset$.
(b) For $q:=0$ to $|w|-1$ do check whether $w \models \Psi^{\prime}\left(q, L^{s-1}, p_{1}, \ldots, p_{k}, P_{1}, \ldots, P_{\ell}\right)$; if so, then insert $q$ into $L^{s}$.
3. Check whether $p_{i}$ belongs to $L^{|w|}$;
if so, then STOP with output "yes", otherwise STOP with output "no".
This computation takes $\mathcal{O}\left(|w|^{2} \cdot|w|^{\mathcal{O}\left(| | \Psi^{\prime}| |\right)}\right)=|w|^{\mathcal{O}(| | \Psi| |)}$ steps.

## Proof of Lemma 3.6.

First, use the algorithm of Lemma 3.5 (b) to construct the $\mathrm{FO}(<)$-formula $\varphi_{h}(Z)$.
It is straightforward to verify that the following formula has the desired properties:

$$
\begin{aligned}
\Phi_{h}(z):=\exists Z( & \left(\forall x Z x \rightarrow P_{\star} x\right) \wedge Z z \wedge \varphi_{h}(Z) \wedge \\
& \left.\forall Z^{\prime}\left(\left(\forall x Z^{\prime} x \rightarrow P_{\star} x\right) \wedge \varphi_{h}\left(Z^{\prime}\right)\right) \rightarrow Z \leqslant \operatorname{lex} Z^{\prime}\right)
\end{aligned}
$$

where $Z \leqslant$ lex $Z^{\prime}$ is an abbreviation for the $\mathrm{FO}(<)$-formula

$$
\left(\forall x Z x \leftrightarrow Z^{\prime} x\right) \vee \exists y\left(\neg Z y \wedge Z^{\prime} y \wedge \forall x\left(x<y \rightarrow\left(Z x \leftrightarrow Z^{\prime} x\right)\right)\right)
$$

Afterwards, we replace every occurrence of an atom of the form $x<y$ in $\Phi_{h}(z)$ by the MSO-formula $\exists Y\left(Y y \wedge \neg Y x \wedge \forall z_{1} \forall z_{2}\left(\operatorname{Succ}\left(z_{1}, z_{2}\right) \wedge Y z_{1}\right) \rightarrow Y z_{2}\right)$.

## Proof of Lemma 3.7.

Since $f$ has bound $f(m) \leqslant \operatorname{Tower}(o(m))$, we know that there is a $m_{0} \in \mathbb{N}$ and a function $g: \mathbb{N} \rightarrow \mathbb{N}$ with $g \in \omega(1)$ such that $f(m) \leqslant \operatorname{Tower}\left(\frac{m}{g(m)}\right)$, for all $m \geqslant m_{0}$. Therefore, for $h:=\lg ^{*}(n)$ and $m:=c \cdot h$ we have

$$
d \cdot f(m) \cdot \lg (f(m)) \leqslant d \cdot \operatorname{Tower}\left(\frac{m}{g(m)}\right) \cdot \operatorname{Tower}\left(\frac{m}{g(m)}-1\right) \leqslant \operatorname{Tower}\left(\frac{m}{g(m)}+1\right)
$$

On the other hand, $n \geqslant \operatorname{Tower}(h-1)$, and thus $\lg ^{(i)}(n) \geqslant \operatorname{Tower}(h-(i+1))$. It therefore suffices to show that $\frac{m}{g(m)}+1 \leqslant h-(i+1)$, i.e., that $\frac{m}{g(m)} \leqslant h-(i+2)$. Since $g \in \omega(1)$, there is some $m_{1} \geqslant m_{0}$ such that for all $m \geqslant m_{1}$ we have $g(m) \geqslant c \cdot(i+3)$. Hence,

$$
\frac{m}{g(m)}=\frac{c \cdot(h-(i+3))}{g(m)}+\frac{c \cdot(i+3)}{g(m)} \leqslant h-(i+3)+1=h-(i+2)
$$

This completes the proof of Lemma 3.7.

## B Proofs omitted in Section 4

## Proof of Theorem 4.1.

Without loss of generality we may assume that $\varphi$ is in negation normal form. Let $x_{1}, \ldots, x_{k}$ be the firstorder variables occurring in $\varphi$, and let $X_{1}, \ldots, X_{k}$ be set variables not occurring in $\varphi$. We claim that for every subformula $\psi\left(x_{i_{1}}, \ldots, x_{i_{\ell}}, \bar{Y}\right)$ there is a formula $\psi^{\prime}\left(X_{i_{1}}, \ldots, X_{i_{\ell}}, \bar{Y}\right)$ such that $\psi^{\prime}$ is positive in $X_{i_{1}}, \ldots, X_{i_{\ell}}, \bar{Y}$ and for all structures $\mathcal{A}$ and for all $a_{1}, \ldots, a_{\ell} \in \mathcal{U}^{\mathcal{A}}, \bar{B} \subseteq \mathcal{U}^{A}$ we have:

$$
\begin{equation*}
\mathcal{A} \models \psi^{\prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell}\right\}, \bar{B}\right) \Longleftrightarrow \mathcal{A} \models \psi\left(a_{1}, \ldots, a_{\ell}, \bar{B}\right) . \tag{*}
\end{equation*}
$$

$\psi^{\prime}$ is defined inductively as follows:

- If $\psi$ is a positive atom, say, $\psi=R x_{1} x_{2}$, then we let $\psi^{\prime}\left(X_{1}, X_{2}\right):=\exists x \in X_{1} \exists y \in X_{2} R x y$. If $\psi$ is a negated atom, say, $\psi=\neg R x_{1} x_{2}$, then we let $\psi^{\prime}\left(X_{1}, X_{2}\right):=\neg \forall x \in X_{1} \forall y \in X_{2} R x y$.
- If $\psi=\psi_{1} \vee \psi_{2}$, then we let $\psi^{\prime}:=\psi_{1}^{\prime} \vee \psi_{2}^{\prime}$. Conjunctions are handled analogously.
- If $\psi$ starts with an existential quantifier, say, $\psi=\exists x_{\ell} \chi\left(x_{1}, \ldots, x_{\ell}, \bar{Y}\right)$, then we let

$$
\psi^{\prime}:=\exists y \forall x\left[\operatorname{LFP}_{x, X_{\ell}} x=y \vee \chi^{\prime}\left(X_{1}, \ldots, X_{\ell}, \bar{Y}\right)\right](x) .
$$

Similarly, if $\psi$ starts with a universal quantifier, say, $\psi=\forall x_{\ell} \chi\left(x_{1}, \ldots, x_{\ell}, \bar{Y}\right)$, then we let

$$
\psi^{\prime}:=\forall y \forall x\left[\operatorname{LFP}_{x, X_{\ell}} x=y \vee \chi^{\prime}\left(X_{1}, \ldots, X_{\ell}, \bar{Y}\right)\right](x)
$$

- If $\psi$ starts with a least fixed point operator, say, $\psi=\left[\operatorname{LFP}_{x_{\ell}, Y} \chi\left(x_{1}, \ldots, x_{\ell}, Y, \bar{Y}\right)\right]\left(x_{\ell}\right)$ (note that by renaming the variables and using equality we can always assume that formulas starting with a least fixed point operator have this form), then we let

$$
\psi^{\prime}:=\exists x\left(X_{\ell}(x) \wedge\left[\operatorname{LFP}_{y, Y} \forall x\left[\operatorname{LFP}_{x, X_{\ell}} x=y \vee \chi^{\prime}\left(X_{1}, \ldots, X_{\ell}, Y, \bar{Y}\right)\right](x)\right](x)\right)
$$

The case that $\varphi$ starts with a greatest fixed point operator can be treated similarly.
To see that the resulting sentence $\varphi^{\prime}$ is indeed equivalent to $\varphi$, the crucial observation is the following: Suppose that $\chi\left(x_{1}, \ldots, x_{\ell}, \bar{Y}\right)$ and $\chi^{\prime}\left(X_{1}, \ldots, X_{\ell}, \bar{Y}\right)$ satisfy ( $*$ ). Let

$$
\chi^{\prime \prime}\left(X_{1}, \ldots, X_{\ell-1}, y, \bar{Y}\right):=\forall x\left[\operatorname{LFP}_{x, X_{\ell}} x=y \vee \chi^{\prime}\left(X_{1}, \ldots, X_{\ell}, \bar{Y}\right)\right](x)
$$

Then for all structures $\mathcal{A}$ and for all $a_{1}, \ldots, a_{\ell} \in \mathcal{U}^{\mathcal{A}}, \bar{B} \subseteq \mathcal{U}^{\mathcal{A}}$ we have:

$$
\begin{equation*}
\mathcal{A} \vDash \chi^{\prime \prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell-1}\right\}, a_{\ell}, \bar{B}\right) \Longleftrightarrow \mathcal{A} \vDash \chi\left(a_{1}, \ldots, a_{\ell}, \bar{B}\right) . \tag{**}
\end{equation*}
$$

To see this, let $a_{1}, \ldots, a_{\ell} \in \mathcal{U}^{\mathcal{A}}$ and $\bar{B} \subseteq \mathcal{U}^{\mathcal{A}}$. Suppose first that $\mathcal{A} \vDash \chi^{\prime \prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell-1}\right\}, a_{\ell}, \bar{B}\right)$. Then $\mathcal{A} \vDash \chi^{\prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell-1}\right\}, \emptyset, \bar{B}\right)$ or $\mathcal{A} \vDash \chi^{\prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell-1}\right\},\left\{a_{\ell}\right\}, \bar{B}\right)$, because otherwise the fixed point process associated with $\chi^{\prime \prime}$ remains empty. Since $X_{\ell}$ only occurs positively in $\chi^{\prime}, \mathcal{A} \models$ $\chi^{\prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell-1}\right\}, \emptyset, \bar{B}\right)$ implies that $\mathcal{A} \vDash \chi^{\prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell-1}\right\},\left\{a_{\ell}\right\}, \bar{B}\right)$, so we may assume the latter. Then by $(*), \mathcal{A} \vDash \chi\left(a_{1}, \ldots, a_{\ell}, \bar{B}\right)$.

Conversely, if $\mathcal{A} \vDash \chi\left(a_{1}, \ldots, a_{\ell}, \bar{B}\right)$ then, by $(*), \mathcal{A} \vDash \chi^{\prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell-1}\right\},\left\{a_{\ell}\right\}, \bar{B}\right)$. But this implies that all elements are in $X$ in the second stage of the fixed point process associated with $\chi^{\prime \prime}$ and hence $\mathcal{A} \models \chi^{\prime \prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell-1}\right\}, a_{\ell}, \bar{B}\right)$.

Now we are ready to prove $(*)$ by induction on $\psi$. It is easy to prove $(*)$ for atoms and negated atoms. The conjunction and disjunction step is trivial. In all other cases, $(*)$ follows immediately from $(* *)$. For example, let $\psi=\left[\operatorname{LFP}_{x_{\ell}, Y} \chi\left(x_{1}, \ldots, x_{\ell}, Y, \bar{Y}\right)\right]\left(x_{\ell}\right)$. Then

$$
\psi^{\prime}=\exists x\left(X_{\ell}(x) \wedge\left[\operatorname{LFP}_{y, Y} \chi^{\prime \prime}\left(X_{1}, \ldots, X_{\ell-1}, y, Y, \bar{Y}\right)\right](x)\right)
$$

Let $a_{1}, \ldots, a_{\ell-1} \in \mathcal{U}^{\mathcal{A}}$ and $\bar{B} \subseteq \mathcal{U}^{\mathcal{A}}$. By $(* *)$, for all $B \subseteq \mathcal{U}^{\mathcal{A}}$ we have

$$
\left\{a \in \mathcal{U}^{\mathcal{A}} \mid \mathcal{A} \models \chi^{\prime \prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell-1}\right\}, a, B, \bar{B}\right)\right\}=\left\{a \in \mathcal{U}^{\mathcal{A}} \mid \mathcal{A} \models \chi\left(a_{1}, \ldots, a_{\ell-1}, a, B, \bar{B}\right)\right\}
$$

Thus for all $a \in \mathcal{U}^{\mathcal{A}}$,

$$
\mathcal{A} \vDash\left[\operatorname{LFP}_{y, Y} \chi^{\prime \prime}\left(\left\{a_{1}\right\}, \ldots,\left\{a_{\ell-1}\right\}, y, \bar{B}\right)\right](a) \Longleftrightarrow \mathcal{A} \vDash\left[\operatorname{LFP}_{y, Y} \chi\left(a_{1}, \ldots, a_{\ell-1}, y, \bar{B}\right)\right](a)
$$

This implies (*) and completes the proof of Theorem 4.1.

## Proof of Theorem 4.5.

Let the arity of a formula be the number of free variables it has.
We say that an MLFP $^{2}$-formula $\varphi(x, \bar{X})$ of the form $\exists y \psi(x, y, \bar{X})$ is in normal form if it is in negation normal form and $\psi(x, y, \bar{X})$ is a disjunction of formulas of the form $\alpha \wedge \theta$ satisfying the following conditions:
(i) $\theta$ is a conjunction of at most unary formulas.
(ii) $\alpha$ is one of the following five basic binary formulas:

$$
\begin{aligned}
& \alpha^{1}=1^{\text {st }} \operatorname{Child}(x, y) \\
& \alpha^{2}=2^{\text {nd }} \operatorname{Child}(x, y) \\
& \alpha^{3}=1^{\text {st }} \operatorname{Child}(y, x), \\
& \alpha^{4}=2^{\text {nd }} \operatorname{Child}(y, x), \\
& \alpha^{5}=\neg 1^{\text {st }} \operatorname{Child}(x, y) \wedge \neg 2^{\text {nd }} \operatorname{Child}(x, y) \wedge \neg 1^{s t} \operatorname{Child}(y, x) \wedge \neg 2^{n d} \operatorname{Child}(y, x)
\end{aligned}
$$

Similarly, a formula of the form $\forall y \psi(x, y, \bar{X})$ is in normal form if is a conjunction of formulas $\delta$ of the form $\alpha \rightarrow \theta$, where
(i') $\theta$ is a disjunction of at most unary formulas.
(ii') $\alpha$ is one of the basic binary formulas $\alpha^{1}, \ldots, \alpha^{5}$.
Finally, an arbitrary MLFP ${ }^{2}$-formula is in normal form if all its subformulas of the form $\exists y \psi$ and $\forall y \psi$ are.
Step 1: For every MLFP ${ }^{2}$-formula $\varphi$ we construct an equivalent MLFP $^{2}$-formula $\varphi^{\prime}$ in normal form. The construction is by induction on $\varphi$, the only interesting cases being subformulas of the form $\exists y \psi$ and $\forall y \psi$. We only consider the universal case; the existential case can be treated similarly.

So let $\varphi(x, \bar{X})=\forall y \psi(x, y, \bar{X})$. We may view $\psi$ as a Boolean combination of atomic formulas and at most unary formulas. We bring this Boolean combination into conjunctive normal form; let $\psi^{\prime}$ be the resulting formula. Now let us look at the clauses $\gamma$ of $\psi^{\prime}$. Let $\theta$ be the disjunction of all at most unary literals in $\gamma$, and let $\alpha^{\prime}$ be the conjunction of the negations of the remaining literals. Then we may replace $\gamma$ with $\left(\alpha^{\prime} \rightarrow \theta\right)$.
$\alpha^{\prime}$ is a conjunction of literals of the form $(\neg) R(x, y)$ and $(\neg) R(y, x)$, where $R \in\left\{1^{\text {st }}\right.$ Child, $2^{\text {nd }}$ Child $\}$. If $\alpha^{\prime}$ is not satisfiable in any tree, then we may simply discard it from $\psi^{\prime}$, because ( $\alpha^{\prime} \rightarrow \theta$ ) always holds. If this makes $\psi^{\prime}$ empty, we replace the whole formula $\varphi$ by $(x=x)$. So let us assume that $\alpha^{\prime}$ is satisfiable in some tree. Then $\alpha^{\prime}$ is equivalent to a disjunction $\alpha^{\prime \prime}$ of some of the five basic binary formulas. Say, $\alpha^{\prime}$ is equivalent to $\bigvee_{j=1}^{\ell} \alpha^{i_{j}}$, for some $\ell \leqslant 5$. Then we replace $\gamma$ by the clauses $\left(\alpha^{i_{1}} \rightarrow \theta\right), \ldots,\left(\alpha^{i_{\ell}} \rightarrow \theta\right)$.

We do this for all clauses of $\psi^{\prime}$ and obtain a formula $\psi^{\prime \prime}$ equivalent to $\psi$ that has the desired form. This completes Step 1.

Before we translate MLFP ${ }^{2}$-formulas in normal form into $\mathrm{FL}_{\mu}$-formulas, we define a few auxiliary MLFP $^{2}$-formulas. Recall that for a tree $T$ and a node $t \in \mathcal{U}^{T}$, by $T_{t}$ we denote the subtree of $T$ with root $t$.
Step 2: For every $\mathrm{FL}_{\mu}$-formula $\Phi$ we construct $\mathrm{FL}_{\mu}$-formulas $\exists \Phi, \exists \downarrow \Phi$, and $\exists \uparrow \Phi$ of size $\mathcal{O}(\|\Phi\|)$ such that for all labelled trees $T$ and nodes $t \in \mathcal{U}^{T}$,

1. $(T, t) \vDash \exists \Phi$ if, and only if, there exists a $u \in \mathcal{U}^{T}$ such that $(T, u) \models \Phi$.
2. $(T, t) \vDash \exists \downarrow \Phi$ if, and only if, there exists a $u \in \mathcal{U}^{T_{t}}$ such that $(T, u) \models \Phi$.
3. $(T, t)=\exists \uparrow \Phi$ if, and only if, there exists a $u \in \mathcal{U}^{T} \backslash \mathcal{U}^{T_{t}}$ such that $(T, u) \models \Phi$.

Similarly, there are formulas $\forall \Phi, \forall \downarrow \Phi$, and $\forall \uparrow \Phi$ with the obvious meaning.
It is an easy exercise to construct the formulas. Note, however, that we make crucial use of the fact that we are working on trees.
Step 3: For every MLFP ${ }^{2}$-formula $\varphi$ in normal form we construct an $\mathrm{FL}_{\mu}$-formula $\Phi$ such that for all labelled trees $T$ and all nodes $t \in \mathcal{U}^{T}, T \models \varphi(t) \Longleftrightarrow(T, t) \models \Phi$.

The construction is by induction on $\varphi$.
For all unary subformulas that do not start with an existential or universal quantifier, we can simply revert the translation from $\mathrm{FL}_{\mu}$ to MLFP ${ }^{2}$.

For 0 -ary formulas $\varphi=\exists x \psi$, suppose $\Psi$ is the $\mathrm{FL}_{\mu}$-formula already constructed for $\psi$. We let $\Phi=\exists \Psi$. Similarly, for 0 -ary $\varphi=\forall x \psi$, we let $\Phi=\forall \Psi$.

As the first interesting case, let us assume that $\varphi(x, \bar{X})=\exists y \psi(x, y, \bar{X})$ is in normal form. Suppose that $\psi=\bigvee_{i=1}^{\ell}\left(\alpha_{i} \wedge \theta_{i}\right)$, where $\theta_{i}$ and $\alpha_{i}$ satisfy the conditions (i) and (ii) for all $i$. For every $i$ we define a formula $\Gamma_{i}$ equivalent to $\exists y\left(\alpha_{i} \wedge \theta_{i}\right)$, and then we let $\Phi=\bigvee_{i=1}^{\ell} \Gamma_{i}$.

So let $1 \leq i \leq \ell$. Let $\theta_{x}$ be the conjunction of all at most unary formulas of $\left(\alpha_{i} \wedge \theta_{i}\right)$ that either do not have any free first-order variables or whose only free first-order variable is $x$, and let $\theta_{y}$ be the conjunction of all at most unary subformulas of $\left(\alpha_{i} \wedge \theta_{i}\right)$ whose only free first-order variable is $y$. If $\alpha_{i}=R(x, y)$ for $R \in\left\{1^{s t}\right.$ Child, $2^{\text {nd }}$ Child $\}$, we let $\Gamma_{i}:=\Theta_{x} \wedge\langle R\rangle \Theta_{y}$, where $\Theta_{x}$ and $\Theta_{y}$ are the $\mathrm{FL}_{\mu^{-}}$ formulas already constructed for $\theta_{x}$ and $\theta_{y}$. Similarly, if $\alpha_{i}=R(y, x)$, we let $\Gamma_{i}:=\Theta_{x} \wedge\left\langle R^{-1}\right\rangle \Theta_{y}$. If $\alpha_{i}=\neg 1^{s t} \operatorname{Child}(x, y) \wedge \neg 2^{\text {nd }} \operatorname{Child}(x, y) \wedge \neg 1^{\text {st }} \operatorname{Child}(y, x) \wedge \neg 2^{\text {nd }} \operatorname{Child}(y, x)$, we let

$$
\begin{aligned}
\Gamma_{i}:=\Theta_{x} \wedge & \left(\left\langle 1^{s t} \text { Child }\right\rangle\left\langle 1^{\text {st }} \text { Child }\right\rangle \exists \downarrow \Theta_{y} \vee\left\langle 1^{\text {st }} \text { Child }\right\rangle\left\langle 2^{\text {nd }} \text { Child }\right\rangle \exists \downarrow \Theta_{y} \vee\right. \\
& \left\langle 2^{2 n} \text { Child }\right\rangle\left\langle 1^{\text {st }} \text { Child }\right\rangle \exists \downarrow \Theta_{y} \vee\left\langle 2^{n d} \text { Child }\right\rangle\left\langle 2^{n d} \text { Child }\right\rangle \exists \downarrow \Theta_{y} \vee \\
& \left\langle 1^{1{ }^{s t}} \text { Child }{ }^{-1}\right\rangle \exists \uparrow \Theta_{y} \vee\left\langle 2^{\text {nd }} \text { Child }^{-1}\right\rangle \exists \uparrow \Theta_{y} .
\end{aligned}
$$

The case $\varphi(x, \bar{X})=\forall y \psi(x, y, \bar{X})$ can be treated similarly. This completes Step 3 .
It remains to prove that the size of $\Phi$ is at most exponential in the size of $\varphi$. The (operator) depth of a formula is naturally defined as the height of its parse tree.

By induction on $d$, we prove that for every at most unary subformula $\psi$ of $\varphi$ of depth $d$ we have

$$
\|\Psi\| \in 2^{\mathcal{O}(d \cdot\|\psi\|)}
$$

Again, the only interesting cases are unary subformulas of the form $\exists y \chi$ or $\forall y \chi$. So suppose that $\psi$ has either of these forms. Let $\xi_{1}, \ldots, \xi_{m}$ be the at most unary subformulas of $\chi$ that are not atoms such that $\chi$ can be written as a Boolean combination of $\xi_{1}, \ldots, \xi_{m}$ and atomic subformulas. Let $\Xi_{1}, \ldots, \Xi_{m}$ be the $\mathrm{FL}_{\mu}$-formulas corresponding to $\xi_{1}, \ldots, \xi_{m}$. By induction hypothesis, for $1 \leq i \leq m$ we have

$$
\left\|\Xi_{i}\right\| \leq 2^{\mathcal{O}\left((d-1) \cdot\left\|\xi_{i}\right\|\right)} \leq 2^{\mathcal{O}((d-1) \cdot\|\psi\|)}
$$

An inspection of our construction shows that

$$
\|\Psi\| \leq 2^{\mathcal{O}(\|\psi\|)} \max _{1 \leq i \leq m}\left\|\Xi_{i}\right\| \leq 2^{\mathcal{O}(\|\psi\|)} 2^{\mathcal{O}((d-1) \cdot\|\psi\|)}=2^{\mathcal{O}(d \cdot\|\psi\|)}
$$

## C Proofs omitted in Section 5

## Proof of Proposition 5.2.

By induction on $n$. The base case $n=1$ is trivial. For $n>1$ and a fixed $i \leqslant n$, the induction hypothesis
gives us for every $j \in\{1, \ldots, n\} \backslash\{i\}$ a $\operatorname{MLFP}\left(\tau \cup\left\{X_{i}\right\}\right)$-formula $\tilde{\Phi}_{X_{j}}(x)$ of size $\leqslant s^{n-1}$ such that, for all $\tau \cup\left\{X_{i}\right\}$-structures $\mathcal{B}, \mathcal{B} \models \forall x \tilde{\Phi}_{X_{j}}(x) \leftrightarrow$

$$
\left[\operatorname{Sim}^{-L F P} x_{x_{1}, X_{1}, ., x_{i-1}, X_{i-1}, x_{i+1}, X_{i+1}, ., x_{n}, X_{n}} \varphi_{1}, \ldots, \varphi_{i-1}, \varphi_{i+1}, \ldots \varphi_{n}\right]_{X_{j}}(x) .
$$

Define $\tilde{\varphi}_{i}\left(x_{i}, X_{i}\right)$ to be the MLFP-formula obtained from $\varphi_{i}\left(x_{i}, X_{1}, \ldots, X_{n}\right)$ by replacing every atom of the form $X_{j}(x)$, for $j \neq i$, with the formula $\tilde{\Phi}_{X_{j}}(x)$. It should be clear that $\left\|\tilde{\varphi}_{i}\right\|<s^{n}$. To complete the proof of Proposition 5.2 it therefore suffices to show the following

Claim. For all $\tau$-structures $\mathcal{A}$ we have

$$
\mathcal{A} \equiv \forall x\left[\operatorname{LFP}_{x_{i}, X_{i}} \tilde{\varphi}_{i}\right](x) \leftrightarrow\left[\operatorname{Sim}-\operatorname{LFP}_{x_{1}, X_{1}, ., x_{n}, X_{n}} \varphi_{1}, \ldots, \varphi_{n}\right]_{X_{i}}(x)
$$

For proving this claim let, for every $S \subseteq \mathcal{U}^{\mathcal{A}}$ and every $j \neq i$,

$$
\begin{aligned}
A_{j}(S) & :=\left\{a \in \mathcal{U}^{\mathcal{A}}:\langle\mathcal{A}, S\rangle \models \tilde{\Phi}_{X_{j}}(a)\right\} \\
A_{i}^{(\ell)} & :=\left\{a \in \mathcal{U}^{\mathcal{A}}: \mathcal{A} \models\left[\operatorname{LFP}_{x_{i}, X_{i}} \tilde{\varphi}_{i}\right](a)\right\} \\
A_{i}^{(r)} & :=\left\{a \in \mathcal{U}^{\mathcal{A}}: \mathcal{A} \models\left[\operatorname{Sim}^{(L F P} P_{x_{1}, X_{1}, ., x_{n}, X_{n}} \varphi_{1}, \ldots, \varphi_{n}\right]_{X_{i}}(a)\right\} \\
A_{j}^{(r)} & :=\left\{a \in \mathcal{U}^{\mathcal{A}}: \mathcal{A} \models\left[\operatorname{Sim}-\operatorname{LFP}_{x_{1}, X_{1}, \ldots, x_{n}, X_{n}} \varphi_{1}, \ldots, \varphi_{n}\right]_{X_{j}}(a)\right\} .
\end{aligned}
$$

Our aim is to show that $A_{i}^{(\ell)}=A_{i}^{(r)}$.
Since $A_{i}^{(\ell)}$ is a fixed point of $\tilde{\varphi}_{i}$, we have $A_{i}^{(\ell)}=$

$$
\left\{a \in \mathcal{U}^{\mathcal{A}}: \mathcal{A} \models \varphi_{i}\left(a, A_{1}\left(A_{i}^{(\ell)}\right), \ldots, A_{i-1}\left(A_{i}^{(\ell)}\right), A_{i}^{(\ell)}, A_{i+1}\left(A_{i}^{(\ell)}\right), \ldots, A_{n}\left(A_{i}^{(\ell)}\right)\right)\right\}
$$

From the definition of $\tilde{\Phi}_{X_{j}}$ we furthermore know that $A_{j}\left(A_{i}^{(\ell)}\right)=$

$$
\left\{a \in \mathcal{U}^{\mathcal{A}}: \mathcal{A} \vDash \varphi_{j}\left(a, A_{1}\left(A_{i}^{(\ell)}\right), \ldots, A_{i-1}\left(A_{i}^{(\ell)}\right), A_{i}^{(\ell)}, A_{i+1}\left(A_{i}^{(\ell)}\right), \ldots, A_{n}\left(A_{i}^{(\ell)}\right)\right)\right\}
$$

I.e., $\left(A_{1}\left(A_{i}^{(\ell)}\right), \ldots, A_{i-1}\left(A_{i}^{(\ell)}\right), A_{i}^{(\ell)}, A_{i+1}\left(A_{i}^{(\ell)}\right), \ldots, A_{n}\left(A_{i}^{(\ell)}\right)\right.$ is a simultaneous fixed point of the formulas $\left(\varphi_{1}, \ldots, \varphi_{n}\right)$ in $\mathcal{A}$, the least fixed point of which is $\left(A_{1}^{(r)}, \ldots, A_{n}^{(r)}\right)$. In particular, this implies that $A_{i}^{(\ell)} \supseteq A_{i}^{(r)}$.

On the other hand, $\left(A_{1}^{(r)}, \ldots, A_{i-1}^{(r)}, A_{i+1}^{(r)}, \ldots, A_{n}^{(r)}\right)$ is a simultaneous fixed point of $\left(\varphi_{1}, \ldots, \varphi_{i-1}\right.$, $\left.\varphi_{i+1}, \ldots, \varphi_{n}\right)$ in $\left\langle\mathcal{A}, A_{i}^{(r)}\right\rangle$, the least fixed point of which is $\left(A_{1}\left(A_{i}^{(r)}\right), \ldots, A_{i-1}\left(A_{i}^{(r)}\right), A_{i+1}\left(A_{i}^{(r)}\right)\right.$, $\left.\ldots, A_{n}\left(A_{i}^{(r)}\right)\right)$.
Therefore, $A_{j}\left(A_{i}^{(r)}\right) \subseteq A_{j}^{(r)}$, for all $j \neq i$, and hence $\left\{a \in \mathcal{U}^{\mathcal{A}}:\left\langle\mathcal{A}, A_{i}^{(r)}\right\rangle \vDash \tilde{\varphi}_{i}(a)\right\}=$

$$
\begin{aligned}
& \left\{a \in \mathcal{U}^{\mathcal{A}}: \mathcal{A} \models \varphi_{i}\left(a, A_{1}\left(A_{i}^{(r)}\right), \ldots, A_{i-1}\left(A_{i}^{(r)}\right), A_{i}^{(r)}, A_{i+1}\left(A_{i}^{(r)}\right), \ldots, A_{n}\left(A_{i}^{(r)}\right)\right)\right\} \subseteq \\
& \left\{a \in \mathcal{U}^{\mathcal{A}}: \mathcal{A} \models \varphi_{i}\left(a, A_{1}^{(r)}, \ldots, A_{i-1}^{(r)}, A_{i}^{(r)}, A_{i+1}^{(r)}, \ldots, A_{n}^{(r)}\right)\right\}=A_{i}^{(r)} .
\end{aligned}
$$

I.e., $A_{i}^{(r)}$ is a pre fixed point ${ }^{3}$ of $\tilde{\varphi}_{i}$ on $\mathcal{A}$, and therefore $A_{i}^{(\ell)} \subseteq A_{i}^{(r)}$.

This finally completes the proof of the above claim and of Proposition 5.2.

## D Proofs omitted in Section 6

## Proof of Theorem 6.5.

From Theorem 6.3 we know that $\mathcal{P}$ can be translated in time $\mathcal{O}(\|\mathcal{P}\|)$ into a TMNF-program $\mathcal{P}^{\prime}$ of size $\mathcal{O}(\|\mathcal{P}\|)$ such that $\left(\mathcal{P}^{\prime}, X\right)$ defines the same unary query as $(\mathcal{P}, X)$ on the class of labelled trees. Let $\ell$ be the number of IDB-predicates of $\mathcal{P}^{\prime}$. In [FGK03, Example 20], a $\sigma$-STA with $2^{\ell}$ states was

[^3]constructed, defining the same unary query as $\left(\mathcal{P}^{\prime}, X\right)$ on the class of labelled trees.

## Proof of Lemma 6.7.

The program $\mathcal{P}$ operates according to the following evaluation algorithm of [FGK03, Proposition 21]: In a bottom-up pass of the input tree $T$, the set $\operatorname{Reach}(t)$ is computed for every vertex $t$ of $T$, where $\operatorname{Reach}(t)$ consists of all states $q$ such that there is a partial run $\rho$ of $\mathfrak{A}$ on $T$ with $\rho(t)=q$. Afterwards, in a top-down pass the set $\operatorname{Acc}(t)$ is computed for every $t \in \mathcal{U}^{T}$, where $\operatorname{Acc}(t)$ consists of all states $q$ such that there is an accepting run $\rho$ of $\mathfrak{A}$ on $T$ with $\rho(t)=q$. Clearly, $t$ does not belong to the query defined by $\mathfrak{A}$ if and only if $\operatorname{Acc}(t)$ contains a state in $Q \backslash S$.

To store the sets $\operatorname{Reach}(\cdot)$ and $\operatorname{Acc}(\cdot)$, respectively, the program $\mathcal{P}$ will use predicates $X_{q}^{\text {Reach }}$ and $X_{q}^{A c c}$, respectively, for every $q \in Q$. The intended meaning of these predicates is, that for all $\sigma$-labelled trees $T, t \in \mathcal{U}^{T}$, and $q \in Q$ we shall have $t \in\left(X_{q}^{\text {Reach }}\right)_{\mathcal{P}}^{\infty}(T)$ iff $q \in \operatorname{Reach}(t)$, and $t \in\left(X_{q}^{\text {Acc }}\right)_{\mathcal{P}}^{\infty}(T)$ iff $q \in \operatorname{Acc}(t)$.
$\mathcal{P}$ has one further predicate, $X^{n o}$ for storing the vertices that do not belong to the query defined by $\mathfrak{A}$. Now, the definition of $\mathcal{P}$ is straightforward:

1. For every $a \in 2^{\sigma}, \mathcal{P}$ contains the rule $\operatorname{Input}_{a}(x) \leftarrow \bigwedge_{P \in a} P(x) \wedge \bigwedge_{P \in \sigma \backslash a} \neg P(x)$.
2. For every $q \in Q \backslash S, \mathcal{P}$ contains the rule $X^{n o}(x) \leftarrow X_{q}^{\text {Acc }}(x)$.
3. For all $q \in F, \mathcal{P}$ contains the rule $X_{q}^{A c c}(x) \leftarrow X_{q}^{\text {Reach }}(x) \wedge \operatorname{Root}(x)$.
4. For all $a \in 2^{\sigma}$ and $p \in \delta(a), \mathcal{P}$ contains the rule $X_{p}^{\text {Reach }}(x) \leftarrow$ Input $_{a}(x) \wedge$ Has-No- $1^{s t}$ Child $(x) \wedge$ Has-No-2 $2^{\text {nd }}$ Child $(x)$.
5. For all $q \in Q, a \in 2^{\sigma}$, and $p \in \delta(1, q, a), \mathcal{P}$ contains the rules $X_{p}^{\text {Reach }}(x) \leftarrow \operatorname{Input}_{a}(x) \wedge 1^{\text {st }} \operatorname{Child}(x, y) \wedge \operatorname{Has}-\operatorname{No}-2^{\text {nd }} \operatorname{Child}(x) \wedge X_{q}^{\text {Reach }}(y)$, and $X_{q}^{\text {Acc }}(y) \leftarrow X_{q}^{\text {Reach }}(y) \wedge 1^{s t}$ Child $(x, y) \wedge \operatorname{Input}_{a}(x) \wedge \operatorname{Has}-$ No- $2^{\text {nd }}$ Child $(x) \wedge X_{p}^{\text {Acc }}(x)$.
6. For all $q \in Q, a \in 2^{\sigma}$, and $p \in \delta(2, q, a), \mathcal{P}$ contains the rules
$X_{p}^{\text {Reach }}(x) \leftarrow \operatorname{Input}_{a}(x) \wedge \operatorname{Has}-N o-1^{\text {st }} \operatorname{Child}(x) \wedge 2^{\text {nd }} \operatorname{Child}(x, y) \wedge X_{q}^{\text {Reach }}(y)$, and
$X_{q}^{\text {Acc }}(y) \leftarrow X_{q}^{\text {Reach }}(y) \wedge 2^{\text {nd }}$ Child $(x, y) \wedge \operatorname{Input}_{a}(x) \wedge$ Has-No-1 $1^{\text {st }}$ Child $(x) \wedge X_{p}^{\text {Acc }}(x)$.
7. For all $q_{1}, q_{2} \in Q, a \in 2^{\sigma}$ and $p \in \delta\left(q_{1}, q_{2}, a\right), \mathcal{P}$ contains the rules
$X_{p}^{\text {Reach }}(x) \leftarrow \operatorname{Input}_{a}(x) \wedge 1^{\text {st }} \operatorname{Child}\left(x, y_{1}\right) \wedge X_{q_{1}}^{\text {Reach }}\left(y_{1}\right) \wedge 2^{\text {nd }} \operatorname{Child}\left(x, y_{2}\right) \wedge X_{q_{2}}^{\text {Reach }}\left(y_{2}\right)$,
$X_{q_{1}}^{\text {Acc }}\left(y_{1}\right) \leftarrow X_{q_{1}}^{\text {Reach }}\left(y_{1}\right) \wedge 1^{\text {st }}$ Child $\left(x, y_{1}\right) \wedge 2^{q_{1} \text { nd }}$ Child $\left(x, y_{2}\right) \wedge$ Input $_{a}(x) \wedge X_{p}^{\text {Acc }}(x) \wedge X_{q_{2}}^{\text {Reach }}\left(y_{2}\right)$,
$X_{q_{2}}^{\text {Acc }}\left(y_{2}\right) \leftarrow X_{q_{2}}^{\text {Reach }}\left(y_{2}\right) \wedge 1^{\text {st }} \operatorname{Child}\left(x, y_{1}\right) \wedge 2^{\text {nd }} \operatorname{Child}\left(x, y_{2}\right) \wedge$ Input $_{a}(x) \wedge X_{p}^{\text {Acc }}(x) \wedge X_{q_{1}}^{\text {Reach }}\left(y_{1}\right)$.
One can easily verify that $\left(\mathcal{P}, X^{\text {no }}\right)$ defines the complement of the query defined by $\mathfrak{A}$ on the class of $\sigma$-labelled trees. Furthermore, $\|\mathcal{P}\|=\mathcal{O}\left(|Q|^{3} \cdot\left|2^{\sigma}\right|+|\sigma| \cdot\left|2^{\sigma}\right|\right)$.

## Proof of Proposition 6.8.

Let $\sigma$ be the schema of $\mathcal{P}$, and let $X \in \operatorname{IDB}(\mathcal{P})$. From Theorem 6.5 we obtain a $\sigma$-STA $\mathfrak{A}_{X}$ with $|Q|=2^{\mathcal{O}(\|\mathcal{P}\|)}$ states defining the same query as $(\mathcal{P}, X)$. Lemma 6.7 gives us a MonDatalog-program $\overline{\mathcal{P}_{X}}$ and a predicate $\bar{X} \in \operatorname{IDB}\left(\overline{\mathcal{P}_{X}}\right)$ such that $\left(\overline{\mathcal{P}_{X}}, \bar{X}\right)$ defines the complement of the query defined by $\mathfrak{A}_{X}$; and $\left\|\overline{\mathcal{P}_{X}}\right\|=\mathcal{O}\left(|Q|^{3} \cdot\left|2^{\sigma}\right|+|\sigma| \cdot\left|2^{\sigma}\right|\right)=2^{\mathcal{O}}(\|\mathcal{P}\|)$.
After an appropriate renaming of IDB-predicates we can assume w.l.o.g. that $\operatorname{IDB}\left(\overline{\mathcal{P}_{X}}\right) \cap \operatorname{IDB}(\mathcal{P})=\emptyset$ and $\operatorname{IDB}\left(\overline{\mathcal{P}_{X}}\right) \cap \operatorname{IDB}\left(\overline{\mathcal{P}_{Y}}\right)=\emptyset$, for all distinct $X, Y \in \operatorname{IDB}(\mathcal{P})$.
Obviously, $\mathcal{P}^{\prime}:=\mathcal{P} \cup \bigcup_{X \in \operatorname{IDB}(\mathcal{P})} \overline{\mathcal{P}_{X}}$ is the desired MonDatalog-program of size $\left\|\mathcal{P}^{\prime}\right\| \leqslant\|\mathcal{P}\|+$ $|\operatorname{IDB}(\mathcal{P})| \cdot 2^{\mathcal{O}(\|\mathcal{P}\|)}=2^{\mathcal{O}(\|\mathcal{P}\|)}$.

## Proof of Theorem 6.9.

Let $\mathcal{P}$ be a S-MonDatalog-program of schema $\sigma$. Let $\left(\mathcal{P}_{1}, \ldots, \mathcal{P}_{n}\right)$ be a decomposition of $\mathcal{P}$ into strata. I.e., $\mathcal{P}=\mathcal{P}_{1} \cup \cdots \cup \mathcal{P}_{n}$, and each $\mathcal{P}_{i}$ can be viewed as a MonDatalog-program of schema $\sigma_{i}:=\sigma \cup \bigcup_{j<i} \operatorname{IDB}\left(\mathcal{P}_{j}\right)$. For each $i$, let $\tilde{\mathcal{P}}_{i}$ be the MonDatalog-program of schema $\tilde{\sigma}_{i}:=\sigma \cup\{X, \bar{X}:$ $\left.X \in \bigcup_{j<i} \operatorname{IDB}\left(\mathcal{P}_{j}\right)\right\}$, obtained from $\mathcal{P}_{i}$ by replacing every occurrence of a literal of the form $\neg X(y)$ by the literal $\bar{X}(y)$. Let $\tilde{\mathcal{P}}_{i}^{\prime}$ be the MonDatalog-program of size $2^{\mathcal{O}\left(\left\|\mathcal{P}_{i}\right\|\right)}$ that Proposition 6.8 provides
for $\tilde{\mathcal{P}}_{i}$. After an appropriate renaming we can assume w.l.o.g. that, viewed as programs of schema $\tilde{\sigma}_{i}$ and $\tilde{\sigma}_{j}$, respectively, $\tilde{\mathcal{P}}_{i}^{\prime}$ and $\tilde{\mathcal{P}}_{j}^{\prime}$ have no IDB-predicates in common (for distinct $i, j \leqslant n$ ).
It is straightforward to check that the program $\mathcal{P}^{\prime}:=\bigcup_{i=1}^{n} \tilde{\mathcal{P}}_{i}^{\prime}$, viewed as a MonDatalog-program of schema $\sigma$, has the desired property that $\operatorname{IDB}\left(\mathcal{P}^{\prime}\right) \supseteq \operatorname{IDB}(\mathcal{P})$ and $\left(\mathcal{P}^{\prime}, X\right)$ defines the same query as $(\mathcal{P}, X)$, for every $X \in \operatorname{IDB}(\mathcal{P})$.
Furthermore, $\left\|\mathcal{P}^{\prime}\right\| \leqslant \sum_{i=1}^{n} 2^{\mathcal{O}\left(\left\|\mathcal{P}_{i}\right\|\right)} \leqslant 2^{\mathcal{O}(\|\mathcal{P}\|)}$.
It is straightforward to see that $\mathcal{P}^{\prime}$ can be computed in time $2^{\mathcal{O}(\|\mathcal{P}\|)}$.

## Proof of Theorem 6.11.

The proof proceeds in 2 steps.
Step 1: Let us first consider the special case where $\mathcal{P}$ is a MonDatalog-program. For every $X \in \operatorname{IDB}(\mathcal{P})$ let $\mathcal{P}_{X}$ be a TMNF-program of size $\mathcal{O}(\|\mathcal{P}\|)$ with $X \in \operatorname{IDB}\left(\mathcal{P}_{X}\right)$ such that $\left(\mathcal{P}_{X}, X\right)$ defines the same unary query as $(\mathcal{P}, X)$ on the class of labelled trees (cf., Theorem 6.3). Let $n$ be the number of IDB-predicates of $\mathcal{P}_{X}$. Lemma 6.1 and Proposition 5.2 give us a MLFP $^{2}$-formula $\Psi_{X}(x)$ of size $\left\|\Psi_{X}\right\| \leqslant\left\|\mathcal{P}_{X}\right\|^{n} \leqslant\left\|\mathcal{P}_{X}\right\|^{\left\|\mathcal{P}_{X}\right\|} \leqslant 2^{\mathcal{O}(\|\mathcal{P}\| \cdot \lg | | \mathcal{P} \|)}$ which defines the same unary query as $(\mathcal{P}, X)$ on the class of labelled trees. Moreover, a close look at the proofs of Lemma 6.1 and Proposition 5.2 for the special case where $\mathcal{P}_{X}$ is in TMNF shows that $\Psi_{X}(x)$ is equivalent to a $\mathrm{FL}_{\mu}$-formula $\Phi_{X}$ of size $\mathcal{O}\left(\left\|\Psi_{X}\right\|\right)$.
Step 2: Now consider the case where $\mathcal{P}$ is an S-MonDatalog-program. Let $\sigma$ be the schema of $\mathcal{P}$, and let $\left(\mathcal{P}_{1}, \ldots, \mathcal{P}_{n}\right)$ be a decomposition of $\mathcal{P}$ into strata. I.e., $\mathcal{P}=\mathcal{P}_{1} \cup \cdots \cup \mathcal{P}_{n}$, and each $\mathcal{P}_{i}$ can be viewed as a MonDatalog-program of schema $\sigma_{i}:=\sigma \cup \bigcup_{j<i} \operatorname{IDB}\left(\mathcal{P}_{j}\right)$. Step 1 gives us, for every $i \leqslant n$ and every $X \in \operatorname{IDB}\left(\mathcal{P}_{i}\right)$, a $\mathrm{FL}_{\mu}$-formula $\Phi_{X}$ of size $\left\|\Phi_{X}\right\|=2^{\mathcal{O}\left(\left\|\mathcal{P}_{i}\right\| \cdot \mathrm{gg}\left\|\mathcal{P}_{i}\right\|\right)}$ that defines the same unary query on the class of $\sigma_{i}$-labelled trees as $\left(\mathcal{P}_{i}, X\right)$.
By induction on $i$ we define, for every $X \in \operatorname{IDB}\left(\mathcal{P}_{i}\right)$, the $\mathrm{FL}_{\mu}$-formula $\hat{\Phi}_{X}$ of size

$$
\left\|\hat{\Phi}_{X}\right\|=2^{\mathcal{O}\left(\left\|\mathcal{P}_{1}\right\| \cdot \lg \left\|\mathcal{P}_{1}\right\|+\cdots+\left\|\mathcal{P}_{i}\right\| \cdot \lg \left\|\mathcal{P}_{i}\right\|\right)}
$$

as follows: For every $X \in \operatorname{IDB}\left(\mathcal{P}_{i}\right)$ let $\hat{\Phi}_{X}$ be the formula obtained from $\Phi_{X}$ by replacing every occurrence of an $Y \in \bigcup_{j<i} \operatorname{IDB}\left(\mathcal{P}_{j}\right)$ by the formula $\hat{\Phi}_{Y}$. It is straightforward to verify for every $X \in \operatorname{IDB}(\mathcal{P})$ that $\hat{\Phi}_{X}$ defines the same unary query as $(\mathcal{P}, X)$ on the class of $\sigma$-labelled trees.

## Proof of Proposition 6.13.

We use the following notations and results of [FG03]:
For all $h \geqslant 1$ let $\left.\Sigma_{h}:=\{0,1,\langle 1\rangle,\langle/ 1\rangle, \ldots,\langle\mathrm{h}\rangle,</ \mathrm{h}\rangle\right\}$. The "tags" $\left.<\mathrm{i}\right\rangle$ and $</ \mathrm{i}>$ represent single letters of the alphabet and are just chosen to improve readability. For every $n \geqslant 1$ let $L(n)$ be the length of the binary representation of the number $n-1$, i.e., $L(0)=0, L(1)=1$, and $L(n)=\lfloor\log (n-1)\rfloor+1$, for all $n \geqslant 2$. By $\operatorname{bit}(i, n)$ we denote the $i$-th bit of the binary representation of $n$, i.e., bit $(i, n)$ is 1 if $\left\lfloor\frac{n}{2^{i}}\right\rfloor$ is odd, and $\operatorname{bit}(i, n)$ is 0 otherwise.
We encode every number $n \in \mathbb{N}$ by a string $\mu_{h}(n)$ over the alphabet $\Sigma_{h}$, where $\mu_{h}(n)$ is inductively defined as follows: $\left.\mu_{1}(0):=<1></ 1\right\rangle$, and

$$
\left.\mu_{1}(n):=<1>\operatorname{bit}(0, n-1) \operatorname{bit}(1, n-1) \cdots \operatorname{bit}(L(n)-1, n-1)</ 1\right\rangle
$$

for $n \geqslant 1$. For $h \geqslant 2$ we let $\mu_{h}(0):=<\mathrm{h}></ \mathrm{h}>$ and

$$
\begin{gathered}
\mu_{h}(n):=<\mathrm{h}> \\
\quad \begin{array}{c}
\mu_{h-1}(0) \operatorname{bit}(0, n-1) \\
\mu_{h-1}(1) \operatorname{bit}(1, n-1)
\end{array} \\
\vdots \\
\\
\quad \begin{array}{l}
\mu_{h-1}(L(n)-1) \operatorname{bit}(L(n)-1, n-1)
\end{array} \\
</ \mathrm{h}>,
\end{gathered}
$$

for $n \geqslant 1$. Here empty spaces and line breaks are just used to improve readability.

Lemma D. 1 ([FG03, Lemma 8]). For every $h \in \mathbb{N}$ there is a $\mathrm{FO}(<)$-formula $\chi_{h, 1}$ of size $e^{4}\left\|\chi_{h, 1}\right\|=$ $\mathcal{O}(h)$ such that the following is true for all strings $w$ over alphabet $\Sigma_{h}$, for all positions $a, b$ in $w$, and for all numbers $m, n \in\{0, \ldots, \operatorname{Tower}(h)\}$ : If $a$ is the first position of a substring $u$ of $w$ that is isomorphic to $\mu_{h}(m)$ and if b is the first position of a substring $v$ of $w$ that is isomorphic to $\mu_{h}(n)$, then $w \models \chi_{h, 1}(a, b)$ if and only if $m=n$.
For every $h \in \mathbb{N}$ let $H:=\operatorname{Tower}(h)$. For $a_{0} \cdots a_{H-1} \in\{0,1\}^{H}$ define the string

$$
w_{h, a_{0} \cdots a_{H-1}}:=<\mathrm{h}+1>\mu_{h}(0) a_{0} \mu_{h}(1) a_{1} \cdots \mu_{h}(H-1) a_{H-1}</ \mathrm{h}+1>
$$

We define the string-language

$$
L_{h}:=\left\{w_{h, a_{0} \cdots a_{H-1}}: a_{0} \cdots a_{H-1} \in\{0,1\}^{H}\right\}^{*}
$$

and we choose a designated string $w_{h} \in L_{h}$ via

$$
w_{h}:=w_{h, \operatorname{BIN}_{H}(0)} w_{h, \operatorname{BIN}_{H}(1)} \cdots w_{h, \operatorname{BIN}_{H}\left(2^{H}-1\right)}
$$

where $\operatorname{BIN}_{H}(n)$ denotes the reverse binary representation of length $H$ of $n$. For example, $\operatorname{BIN}_{4}(2)=$ 0100 and $\mathrm{BIN}_{4}(5)=1010$. It is straightforward to see the following:

Lemma D.2. For $h \in \mathbb{N}$ let $H:=\operatorname{Tower}(h)$. There is no nondeterministic finite automaton $\mathfrak{A}_{h}$ with less than $2^{H}=\operatorname{Tower}(h+1)$ states such that the following is true for every $w \in L_{h}: \mathfrak{A}_{h}$ accepts $w$ if and only if $w=w_{h}$.

Lemma D.3. For every $h \in \mathbb{N}$ there is $a \mathrm{FO}(<)$-sentence $\varphi_{h}$ of size $\mathcal{O}(h)$ such that the following is true for all $w \in L_{h}: w=\varphi_{h}$ if and only if $w=w_{h}$.

Proof. To simplify notation we, in the following, write $\operatorname{Succ}(x)$ to denote the successor of a position $x$ in a string $w$.
Using the formula $\chi_{h, 1}$ of Lemma D.1, it is straightforward to build a $\mathrm{FO}(<)$-sentence $\varphi_{h}$ stating for every input string $w \in L_{h}$ that

- the leftmost substring of $w$ of the form $<\mathrm{h}+1>\cdots</ \mathrm{h}+1>$ contains no position $x$ with $P_{</ \mathrm{h}>}(x) \wedge P_{1}(\operatorname{Succ}(x))$,
- the rightmost substring of $w$ of the form $<\mathrm{h}+1>\cdots</ \mathrm{h}+1>$ contains no position $y$ with $P_{</ \mathrm{h}>}(y) \wedge P_{0}(\operatorname{Succ}(x))$,
- for every two successive substrings $u$ and $v$ of $w$ of the form $<\mathrm{h}+1>\cdots</ \mathrm{h}+1>$, there is a position $x_{1}$ in $u$ with $P_{<\mathrm{h}>}\left(x_{1}\right)$ and a position $y_{1}$ in $v$ with $P_{<\mathrm{h}>}\left(y_{1}\right)$ such that $\chi_{h, 1}\left(x_{1}, y_{1}\right)$ is true and
- every position $x_{0}$ in $u$ to the left of $x_{1}$ satisfies $P_{</ \mathrm{h}>}\left(x_{0}\right) \rightarrow P_{1}\left(\operatorname{Succ}\left(x_{0}\right)\right)$,
- every position $y_{0}$ in $v$ to the left of $y_{1}$ satisfies $P_{</ \mathrm{h}>}\left(y_{0}\right) \rightarrow P_{0}\left(\operatorname{Succ}\left(y_{0}\right)\right)$,
- the first position $x_{1}^{\prime}$ to the right of $x_{1}$ in $u$ with $P_{</ \mathrm{h}>}\left(x_{1}^{\prime}\right)$ satisfies $P_{0}\left(\operatorname{Succ}\left(x_{1}^{\prime}\right)\right)$,
- the first position $y_{1}^{\prime}$ to the right of $y_{1}$ in $v$ with $P_{</ \mathrm{h}>}\left(y_{1}^{\prime}\right)$ satisfies $P_{1}\left(\operatorname{Succ}\left(y_{1}^{\prime}\right)\right)$,
- for all positions $x_{2}$ to the right of $x_{1}^{\prime}$ in $u$ and for all positions $y_{2}$ to the right of $y_{1}^{\prime}$ in $v$ such that $P_{<\mathrm{h}>}\left(x_{2}\right) \wedge P_{<\mathrm{h}>}\left(y_{2}\right) \wedge \chi_{h, 1}\left(x_{2}, y_{2}\right)$ is true, also $P_{1}\left(\operatorname{Succ}\left(x_{2}^{\prime}\right)\right) \leftrightarrow P_{1}\left(\operatorname{Succ}\left(y_{2}^{\prime}\right)\right)$ is true, where $x_{2}^{\prime}$ is the first position to the right of $x_{2}$ in $u$ with $P_{</ \mathrm{h}>}\left(x_{2}^{\prime}\right)$ and $y_{2}^{\prime}$ is the first position to the right of $y_{2}$ in $v$ with $P_{</ \mathrm{h}>}\left(y_{2}^{\prime}\right)$.

It is straightforward to check that the string $w_{h}$ is the unique string in $L_{h}$ that satisfies $\varphi_{h}$, and that $\varphi_{h}$ has size $\mathcal{O}\left(\left\|\chi_{h, 1}\right\|\right)$. Together with the bound $\left\|\chi_{h, 1}\right\|=\mathcal{O}(h)$ of Lemma D.1, this completes the proof of Lemma D.3.

[^4]From Lemma D. 2 and Lemma D. 3 we obtain, for every $h \in \mathbb{N}$, a FO $(<)$-sentence $\varphi_{h}$ of size $\mathcal{O}(h)$ that defines a string-language not definable by a nondeterministic finite automaton with less than $\operatorname{Tower}(h+1)$ states. This, in particular, completes the proof of Proposition 6.13.

## Proof of Theorem 6.14.

By contradiction. Assume that the translation from $\mathrm{FO}(<)$ to S-MonDatalog can be established by a function $f$ with bound $f(m) \leqslant \operatorname{Tower}(o(m))$. I.e., there is a $m_{0} \in \mathbb{N}$ and a function $g \in \omega(1)$ such that $f(m) \leqslant \operatorname{Tower}\left(\frac{m}{g(m)}\right)$, for all $m \geqslant m_{0}$; and for every $\mathrm{FO}(<)$-sentence $\varphi$ there is a S-MonDatalogprogram $\mathcal{P}$ of size $\|\mathcal{P}\| \leqslant f(\|\varphi\|) \leqslant \operatorname{Tower}\left(\frac{\|\varphi\|}{g(\|\varphi\|)}\right)$ and an $X \in \operatorname{IDB}(\mathcal{P})$ such that $(\mathcal{P}, X)$ defines the same Boolean query as $\varphi$ on the class of all finite strings. According to Theorem 6.9, $\mathcal{P}$ is equivalent to a MonDatalog-program $\mathcal{P}^{\prime}$ of size $\left\|\mathcal{P}^{\prime}\right\|=2^{\mathcal{O}(\|\mathcal{P}\|)} \leqslant \operatorname{Tower}\left(\frac{\|\varphi\|}{g(\|\varphi\|)}+c\right)$, for a suitable constant $c \in \mathbb{N}$. From Theorem 6.5 we obtain an STA $\mathfrak{A}$ that defines the same unary query on the class of labelled trees as $\left(\mathcal{P}^{\prime}, X\right)$ and that has $|Q| \leqslant 2^{\mathcal{O}\left(\left\|\mathcal{P}^{\prime}\right\|\right)} \leqslant \operatorname{Tower}\left(\frac{\|\varphi\|}{g(\|\varphi\|)}+d\right)$ states, for a suitable constant $d \in \mathbb{N}$. Restricting attention to strings again, it is straightforward to transform the STA $\mathfrak{A}$ into a nondeterministic finite automaton $\mathfrak{A}^{\prime}$ that accepts exactly those strings that belong to the Boolean query defined by $\mathfrak{A}$ and that has at most $2^{|Q|}$ different states. I.e., every $\mathrm{FO}(<)$-sentence $\varphi$ can be translated into a nondeterministic finite automaton that defines the same Boolean query as $\varphi$ on the class of all finite strings and that has at most $\operatorname{Tower}(h(\|\varphi\|))$ states, where the function $h$ is defined via $h(m):=\frac{m}{g(m)}+d+1$. Obviously, $h \in o(m)$, contradicting Proposition 6.13 and completing the proof of Theorem 6.14.


[^0]:    * supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD)

[^1]:    ${ }^{1}$ An operator $F: 2^{M} \rightarrow 2^{M}$ is monotone iff $F(A) \subseteq F(B)$ for all $A \subseteq B \subseteq M$.

[^2]:    ${ }^{2}$ In [FG03], an additional factor $\lg h$ occurs because there a logarithmic cost measure is used for the formula size, whereas here we use a uniform measure (cf., Section 2.4).

[^3]:    ${ }^{3}$ Recall that a pre fixed point of a monotone operator $F: 2^{M} \rightarrow 2^{M}$ is a set $A \subseteq M$ with $F(A) \subseteq A$, and that the least fixed point of $F$ is the intersection of all pre fixed points of $F$.

[^4]:    ${ }^{4}$ In [FG03], an additional factor $\lg h$ occurs because there a logarithmic cost measure is used for the formula size, whereas here we use a uniform measure (cf., Section 2.4).

