

Comparing the succinctness of monadic query languages over finite trees

Martin Grohe and Nicole Schweikardt*

Laboratory for Foundations of Computer Science
University of Edinburgh, Scotland, U.K.
email: {grohe|vlnschwe}@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, μ -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 node-selecting 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 [NS02] 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 $\text{PTIME} \neq \text{NP}$, there is no algorithm evaluating a monadic second-order query in time $f(\text{size of query})p(\text{size of tree})$

* supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD)

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 MSO-formula 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 $f(\text{size of query})p(\text{size of tree})$ 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 $i \geq 1$, NP is not contained in $\text{DTIME}(n^{\log^{(i)}(n)})$, 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. MLFP^2 is exponentially more succinct than the full modal μ -calculus, that is, the modal μ -calculus with future and past modalities (see Theorem 4.5, Example 4.3, and Theorem 4.6).
4. The full modal μ -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 μ -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 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 known 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 second-order logic and the modal μ -calculus have the same expressive power on (ordered) trees, but again it 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 [AI01] 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 non-elementarily 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 μ -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 Σ we write Σ^* to denote the set of all finite strings over Σ , and we use ε to denote the empty string. We use \mathbb{N} to denote the set $\{0, 1, 2, \dots\}$ 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 $Tower(0) := 1$ and $Tower(h+1) = 2^{Tower(h)}$, for all $h \in \mathbb{N}$. I.e., $Tower(h)$ is a tower of 2s of height h .

We say that a function $f : \mathbb{N} \rightarrow \mathbb{N}$ has bound $f(m) \leq Tower(o(h(m)))$, 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 $m \geq m_0$ we have $f(m) \leq Tower(g(m))$.

Note that, in particular, every *elementary* function f has bound $f(m) \leq Tower(o(m))$. 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 2s of height h with an n on top.

2.2 Structures

A *signature* τ is a finite set of relation symbols and constant symbols. Each relation symbol $R \in \tau$ has a fixed arity $ar(R)$. A τ -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 $R^{\mathcal{A}} \subseteq (\mathcal{U}^{\mathcal{A}})^{ar(R)}$ 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

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

where *Root*, *Has-No-1stChild*, *Has-No-2ndChild* are unary relation symbols and *1stChild*, *2ndChild* are binary relation symbols. We define *Trees* to be the set of all τ_{Trees} -structures T that satisfy the following conditions:

1. $U^T \subseteq \{1, 2\}^*$ and for every string $si \in U^T$ with $i \in \{1, 2\}$ we also have $s \in U^T$.
2. $Root^T$ consists of the empty string ε .
3. $1^{st}Child^T$ consists of the pairs $(s, s1)$, for all $s1 \in U^T$.
4. $2^{nd}Child^T$ consists of the pairs $(s, s2)$, for all $s2 \in U^T$.
5. $Has-No-1^{st}Child^T$ consists of all strings $s \in U^T$ with $s1 \notin U^T$.
6. $Has-No-2^{nd}Child^T$ consists of all strings $s \in U^T$ with $s2 \notin U^T$.

For $T \in Trees$ and $t \in U^T$ we write T_t to denote the subtree of T with root t .

A *schema* σ is a set of unary relation symbols each of which is distinct from *Has-No-1stChild*, *Has-No-2ndChild*, *Root*. A σ -labelled tree is a $(\tau_{Trees} \cup \sigma)$ -structure consisting of some $T \in Trees$ and additional interpretations $P^T \subseteq U^T$ for all symbols $P \in \sigma$. We sometimes write $label(t)$ to denote the set $\{P \in \sigma : t \in P^T\}$ of labels at vertex t in T .

We identify a string $w = w_0 \dots w_{n-1}$ of length $|w| = n \geq 1$ over an alphabet Σ with a σ -labelled tree T^w in the following way: We choose σ 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 $U^{T^w} = \{\varepsilon, 1, 11, \dots, 1^{n-1}\}$, and we choose $P_a^{T^w} := \{1^i : w_i = a\}$, for each $a \in \Sigma$. This corresponds to the conventional representation of strings by structures in the sense that $\langle U^{T^w}, 1^{st}Child, (P_a^{T^w})_{a \in \Sigma} \rangle$ is isomorphic to the structure $\langle \{0, \dots, n-1\}, Succ, (P_a^w)_{a \in \Sigma} \rangle$ where *Succ* denotes the binary successor relation on $\{0, \dots, n-1\}$ 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 $\{0, \dots, n-1\}$ (respectively, the transitive closure of the relation *1stChild*). In these cases we explicitly write $FO(<)$ 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 *1stChild*, *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 second-order logic*, for short: MSO (cf., e.g., the textbooks [EF99, Imm99]). We use $\text{FO}(\tau)$ and $\text{MSO}(\tau)$, respectively, to denote the class of all first-order formulas and monadic second-order formulas, respectively, of signature τ . We write $\varphi(x_1, \dots, x_k, X_1, \dots, X_\ell)$ to indicate that the free first-order variables of the formula φ are x_1, \dots, x_k and the free set variables are X_1, \dots, X_ℓ . Sometimes we use \bar{x} and \bar{X} as abbreviations for sequences x_1, \dots, x_k and X_1, \dots, X_ℓ of variables.

A formula $\varphi(x)$ of signature τ defines the *unary query* which associates with every τ -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 φ when interpreting the free occurrences of the variable x by the element a . A *sentence* φ of signature τ (i.e., a formula that has no free variables) defines the *Boolean query* that associates the answer “yes” with all τ -structures that satisfy φ and the answer “no” with all other τ -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 τ be a signature and let $\varphi(x, X)$ be a formula of signature τ which is *positive* in the variable X , i.e., every atom of the form $X(z)$ occurs in φ within an *even* number of negation symbols. φ defines for every τ -structure \mathcal{A} a monotone operator¹ $F_{\mathcal{A}, \varphi} : 2^{\mathcal{U}^{\mathcal{A}}} \rightarrow 2^{\mathcal{U}^{\mathcal{A}}}$ via $F_{\mathcal{A}, \varphi}(A) := \{a \in \mathcal{U}^{\mathcal{A}} : \mathcal{A} \models \varphi(a, A)\}$, 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 φ in \mathcal{A} iff $F_{\mathcal{A}, \varphi}(A) = A$ (respectively, $F_{\mathcal{A}, \varphi}(A) \subseteq A$, respectively, $F_{\mathcal{A}, \varphi}(A) \supseteq A$).

For all $s \in \mathbb{N}$ we define $L_{\mathcal{A}, \varphi}^0 := \emptyset$, $L_{\mathcal{A}, \varphi}^{s+1} := F_{\mathcal{A}, \varphi}(L_{\mathcal{A}, \varphi}^s)$, and $G_{\mathcal{A}, \varphi}^0 := \mathcal{U}^{\mathcal{A}}$, $G_{\mathcal{A}, \varphi}^{s+1} := F_{\mathcal{A}, \varphi}(G_{\mathcal{A}, \varphi}^s)$. 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 φ in \mathcal{A} . It is straightforward to see that $L_{\mathcal{A}, \varphi}^\infty$ is indeed contained in every *pre fixed point* of φ in \mathcal{A} , and that every *post fixed point* of φ 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 φ in \mathcal{A} .

The logic $\text{MLFP}(\tau)$ is the extension of $\text{FO}(\tau)$ by least and greatest fixed point operators. I.e.: $\text{MLFP}(\tau)$ contains $\text{FO}(\tau)$ and is closed under Boolean connectives and first-order quantifications; and if $\varphi(x, X, \bar{y}, \bar{Y})$ is an $\text{MLFP}(\tau)$ -formula which is positive in the variable X then $[\text{LFP}_{x, X} \varphi](z)$ and $[\text{GFP}_{x, X} \varphi](z)$ are $\text{MLFP}(\tau)$ -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 $\mathcal{A} \models [\text{LFP}_{x, X} \varphi](a)$ iff $a \in L_{\mathcal{A}, \varphi}^\infty$, and $\mathcal{A} \models [\text{GFP}_{x, X} \varphi](a)$ 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 $[\text{GFP}_{x, X} \varphi](z)$ is equivalent to the formula $\neg[\text{LFP}_{x, X} \neg\varphi'](z)$, where φ' is obtained from φ by replacing every atom of the form $X(u)$ by the literal $\neg X(u)$. Therefore, every MLFP -formula Ψ can easily be transformed into an equivalent MLFP -formula Ψ' 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 φ to be the number of vertices of the tree that corresponds to φ .

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

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 $\|\varphi_2\| \leq f(\|\varphi_1\|)$ which is equivalent to φ_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 $[\text{LFP}_{x,X}(x, X, \bar{y}, \bar{Y})](z)$ is equivalent to $\forall X (Xz \vee \exists x \neg Xx \wedge \varphi(x, X, \bar{y}, \bar{Y}))$. \square

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. \square

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 Φ can be translated into a nondeterministic tree automaton with $\text{Tower}(\mathcal{O}(\|\Phi\|))$ states that accepts exactly those labelled trees that satisfy Φ . This leads to

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

To show that we cannot do essentially better, i.e., that there is no translation from MSO to MLFP of size $\text{Tower}(o(m))$ 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 worst-case complexity $2^{\Omega(n)}$ (cf., [DGH⁺02]). Although not answering the P vs. NP question, the exposition of a deterministic algorithm for SAT with worst-case complexity $\leq n^{\lg n}$ 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 $\lg^{(i+1)}(n) := \lg(\lg^{(i)}(n))$. Moreover, we write \lg^* to denote the “inverse” of the *Tower* function, that is, the (unique) integer valued function with $\text{Tower}(\lg^*(n)-1) < n \leq \text{Tower}(\lg^*(n))$.

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

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 γ ,

- constructs an MSO-formula $\Phi(z)$ of extremely small size that, when evaluated in w , specifies a canonical satisfying assignment for γ (if γ is satisfiable at all),
- tests, for all MLFP-formulas $\Psi(z)$ of size $\leq f(|\Phi|)$, whether Ψ specifies a satisfying assignment for γ .

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 $w \models \Psi(p)$. \square*

Let us now concentrate on the construction of a string w that represents a SAT-instance γ and of an MSO-formula $\Phi(z)$ that specifies a canonical satisfying assignment of γ (provided that γ is satisfiable at all). Since we want Φ to be extremely short, we cannot choose w to be the straightforward string-representation of γ . Instead, we use the following, more complicated, representation of [FG03]:

For all $h \geq 1$ let $\Sigma_h := \{0, 1, \langle 1 \rangle, \langle /1 \rangle, \dots, \langle h \rangle, \langle /h \rangle\}$. The “tags” $\langle i \rangle$ and $\langle /i \rangle$ represent single letters of the alphabet and are just chosen to improve readability. For every $n \geq 1$ 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 $n \geq 2$. By $\text{bit}(i, n)$ we denote the i -th bit of the binary representation of n , i.e., $\text{bit}(i, n)$ is 1 if $\lfloor \frac{n}{2^i} \rfloor$ is odd, and $\text{bit}(i, n)$ is 0 otherwise.

We encode every number $n \in \mathbb{N}$ by a string $\mu_h(n)$ over the alphabet Σ_h , where $\mu_h(n)$ is inductively defined as follows: $\mu_1(0) := \langle 1 \rangle \langle /1 \rangle$, and

$$\mu_1(n) := \langle 1 \rangle \text{bit}(0, n-1) \text{bit}(1, n-1) \cdots \text{bit}(L(n)-1, n-1) \langle /1 \rangle,$$

for $n \geq 1$. For $h \geq 2$ we let $\mu_h(0) := \langle h \rangle \langle /h \rangle$ and

$$\begin{aligned} \mu_h(n) := & \langle h \rangle \\ & \mu_{h-1}(0) \text{bit}(0, n-1) \\ & \mu_{h-1}(1) \text{bit}(1, n-1) \\ & \vdots \\ & \mu_{h-1}(L(n)-1) \text{bit}(L(n)-1, n-1) \\ & \langle /h \rangle, \end{aligned}$$

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

$$\begin{aligned} \Sigma'_h := & \Sigma_h \cup \\ & \{+, -, \langle \text{lit} \rangle, \langle / \text{lit} \rangle, \langle \text{clause} \rangle, \langle / \text{clause} \rangle, \langle \text{cnf} \rangle, \langle / \text{cnf} \rangle\} \cup \\ & \{\star, \langle \text{ass} \rangle, \langle / \text{ass} \rangle, \langle \text{val} \rangle, \langle / \text{val} \rangle\}. \end{aligned}$$

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

$$\mu_h(X_i) := \langle \text{lit} \rangle \mu_h(i) \langle / \text{lit} \rangle,$$

and the literal $\neg X_i$ is encoded by $\mu_h(\neg X_i) := \langle \text{lit} \rangle \mu_h(i) \langle / \text{lit} \rangle$.

A clause $\delta := \lambda_1 \vee \cdots \vee \lambda_r$ of literals is encoded by

$$\mu_h(\delta) := \langle \text{clause} \rangle \mu_h(\lambda_1) \cdots \mu_h(\lambda_r) \langle / \text{clause} \rangle.$$

A CNF-formula $\gamma := \delta_1 \wedge \cdots \wedge \delta_m$ is encoded by the string

$$\mu_h(\gamma) := \langle \text{cnf} \rangle \mu_h(\delta_1) \cdots \mu_h(\delta_m) \langle / \text{cnf} \rangle.$$

We write $\text{CNF}(n)$ to denote the class of all CNF-formulas the propositional variables of which are among X_0, \dots, X_{n-1} . To provide the “infrastructure” for specifying a truth assignment to the variables X_0, \dots, X_{n-1} , we use the string

$$\begin{aligned} \mu_h(X_0, \dots, X_{n-1}) := & \langle \text{ass} \rangle \\ & \langle \text{val} \rangle \mu_h(0) \star \langle / \text{val} \rangle \\ & \langle \text{val} \rangle \mu_h(1) \star \langle / \text{val} \rangle \\ & \vdots \\ & \langle \text{val} \rangle \mu_h(n-1) \star \langle / \text{val} \rangle \\ & \langle / \text{ass} \rangle. \end{aligned}$$

Remark 3.4. There is a 1–1-correspondence between assignments $\alpha : \{X_0, \dots, X_{n-1}\} \rightarrow \{true, false\}$, on the one hand, and sets P of positions of $\mu_h(X_0, \dots, X_{n-1})$ that carry the letter \star , on the other hand: Such a set P specifies the assignment α^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 $\mu_h(X_0, \dots, X_{n-1})$ belongs to P . Conversely, a given assignment α specifies the set P^α consisting of exactly those \star -positions of $\mu_h(X_0, \dots, X_{n-1})$ that occur directly after a substring $\mu_h(i)$ where $\alpha(X_i) = true$. \square

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

$$\mu_h(\gamma, \star) := \mu_h(\gamma) \mu_h(X_0, \dots, X_{n-1}).$$

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

Lemma 3.5.

- (a) *There is an algorithm that, given $h \in \mathbb{N}$ and $\gamma \in \text{CNF}(n)$, computes (a binary representation of) the string $\mu_h(\gamma, \star)$ in time $\mathcal{O}(h \cdot (\lg h) \cdot (\lg n)^2 \cdot (|\gamma| + n))$ (cf., [FG03, Lemma 9]).
The string $\mu_h(\gamma, \star)$ has length $|\mu_h(\gamma, \star)| = \mathcal{O}(h \cdot (\lg n)^2 \cdot (|\gamma| + n))$.*
- (b) *There is an algorithm that, given $h \in \mathbb{N}$, computes (the binary representation of) a FO($<$)-formula $\varphi_h(Z)$ in time $\mathcal{O}(h \cdot \lg h)$, such that for all $n \leq \text{Tower}(h)$, for all $\gamma \in \text{CNF}(n)$, 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 size² $|\varphi_h(Z)| = \mathcal{O}(h)$. \square

Given a $\text{CNF}(n)$ -formula γ and its representative $\mu_h(\gamma, \star)$, we now specify a *canonical* satisfying assignment of γ , provided that γ is satisfiable at all. As observed in Remark 3.4, every assignment $\alpha : \{X_0, \dots, X_{n-1}\} \rightarrow \{true, false\}$ corresponds to a set P^α of positions in $\mu_h(\gamma, \star)$ that carry the letter \star . P^α , again, can be identified with the 0-1-string of length $|\mu_h(\gamma, \star)|$ that carries the letter 1 exactly at those positions that belong to P^α . Now, the lexicographic ordering of these strings gives us a linear ordering on the set of all assignments $\alpha : \{X_0, \dots, X_{n-1}\} \rightarrow \{true, false\}$. As the canonical satisfying assignment of γ 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 $n \leq \text{Tower}(h)$, for all $\gamma \in \text{CNF}(n)$, and for all positions p of $\mu_h(\gamma, \star)$ that carry the letter \star , we have*

$$\mu_h(\gamma, \star) \models \Phi_h(p) \quad \text{iff} \quad \text{in the lexicographically smallest satisfying assignment for } \gamma, \text{ the propositional variable corresponding to position } p \text{ is assigned the value } true.$$

The formula $\Phi_h(z)$ has size $|\Phi_h| = \mathcal{O}(h)$. \square

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 $|\Psi| \leq f(|\Phi|)$ which defines the same query as Φ 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 γ 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 $\leq f(|\Phi_h|)$ defines the same query as $\Phi_h(z)$.

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).

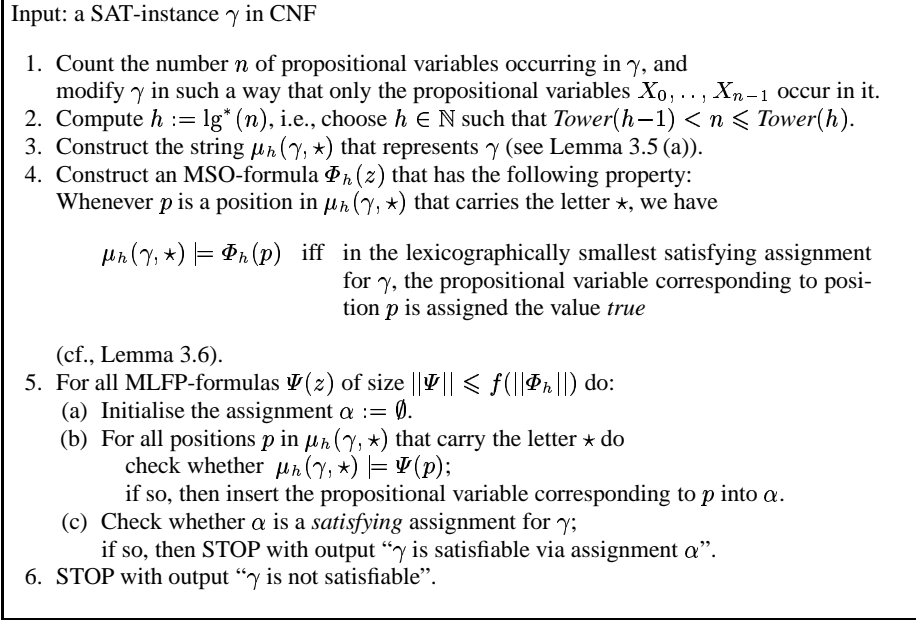


Fig. 1. A SAT-solving algorithm.

It remains to determine the worst-case running time of the algorithm. Let γ be an input CNF-formula for the algorithm, let n be the number of propositional variables of γ , 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 $\|\Phi_h\| \leq c \cdot h$, 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 $2^{c_1 \cdot f(\|\Phi_h\|) \cdot \lg(f(\|\Phi_h\|))}$ times, for a suitable constant $c_1 \in \mathbb{N}$. To see this, note that formulas of length $\leq f(\|\Phi_h\|)$ use at most $f(\|\Phi_h\|)$ different first-order variables and at most $f(\|\Phi_h\|)$ different set variables. I.e., these formulas can be viewed as strings of length $f(\|\Phi_h\|)$ over an alphabet of size $c_2 + 2 \cdot f(\|\Phi_h\|)$, for a suitable constant $c_2 \in \mathbb{N}$. Therefore, the number of such formulas is $\leq (c_2 + 2 \cdot f(\|\Phi_h\|))^{f(\|\Phi_h\|)} \leq 2^{c_1 \cdot f(\|\Phi_h\|) \cdot \lg(f(\|\Phi_h\|))}$.

Each performance of the loop in step 5 will take a number of steps polynomial in

$$|\mu_h(\gamma, \star)|^{\mathcal{O}(f(\|\Phi_h\|))} \leq (c_3 \cdot h \cdot (\lg n)^2 \cdot \|\gamma\|)^{c_4 \cdot f(c \cdot h)},$$

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 $f(m) \leq \text{Tower}(o(m))$. From Lemma 3.7 below we then obtain that our SAT-solving algorithm has, for every $i \in \mathbb{N}$, time bound $\|\gamma\|^{1 \lg^{(i)}(n)}$. This finally completes the proof of Theorem 3.2. ■

Lemma 3.7. *Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a function with bound $f(m) \leq \text{Tower}(o(m))$, 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 $n \geq n_0$ we have*

$$d \cdot f(c \cdot \lg^*(n)) \cdot \lg(f(c \cdot \lg^*(n))) \leq 1 \lg^{(i)}(n). \quad \square$$

4 The Two-Variable Fragment of MLFP and the Full Modal μ -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 first-order 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 τ be a signature that consists of relation symbols of arity ≤ 2 . Every MLFP(τ)-sentence φ is equivalent to an MLFP(τ)-sentence φ' of size $\|\varphi'\| \in \mathcal{O}(\|\varphi\|)$ with just two first-order variables. \square*

The following restriction MLFP² of MLFP is more interesting. MLFP² is the monadic version of the logic FP² considered by [GOR99,GO99], which is denoted by $\widehat{\text{FP}}^2$ in [Dzi96].

Definition 4.2 (MLFP²). *An MLFP-formula is strict if no parameters are allowed in the fixed points, that is, if for all subformulas of the form $[LFP_{x,x}\varphi](y)$, x is the only free first-order variable of φ .*

We define 2-variable monadic least fixed point logic, denoted by MLFP² to be the restriction of MLFP to strict formulas with at most 2 first-order variables. \square

We first note that MLFP², and actually FO², 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 $\sigma := \{L, R, P_1, \dots, P_n\}$ and

$$\varphi_n := \forall x \left(Lx \rightarrow \exists y \left(Ry \wedge \bigwedge_{i=1}^n (P_i x \leftrightarrow P_i y) \right) \right).$$

We claim that every nondeterministic finite automaton accepting precisely those strings over alphabet 2^σ that satisfy φ has at least 2^{2^n} states. To see this, for every $S \subseteq 2^{\{1, \dots, n\}}$, 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 $s = \{i \mid x \in P_i^{X_n(S)}\} = \{i \mid y \in P_i^{Y_n(S)}\}$.

Let $W_n(S, T) := X_n(S)Y_n(T)$ be the concatenation of $X_n(S)$ and $Y_n(T)$. Then $W_n(S, T) \models \varphi \iff S \subseteq T$. Clearly, a nondeterministic finite automaton accepting precisely those strings $W_n(S, T)$ with $S \subseteq T$ needs at least 2^{2^n} states. \square

Let us return to binary trees now. Following Vardi [Var98], we define the *full modal μ -calculus* FL _{μ} on binary trees as follows:

Definition 4.4 (FL _{μ}). *For each schema σ , an FL _{μ} -formula of schema σ is either:*

- *true, false, P , or $\neg P$, where $P \in \sigma \cup \{\text{Root}, \text{Has-No-1}^{\text{st}}\text{Child}, \text{Has-No-2}^{\text{nd}}\text{Child}\}$;*
- *$\Phi_1 \wedge \Phi_2$ or $\Phi_1 \vee \Phi_2$, where Φ_1 and Φ_2 are FL _{μ} -formulas of schema σ ;*
- *X , where X is a propositional variable;*
- *$\langle R \rangle \Phi$ or $[R] \Phi$, where $R \in \{1^{\text{st}}\text{Child}, 2^{\text{nd}}\text{Child}, 1^{\text{st}}\text{Child}^{-1}, 2^{\text{nd}}\text{Child}^{-1}\}$ and Φ is an FL _{μ} -formula of schema σ ;*
- *$\mu X. \Phi$ or $\nu X. \Phi$, where X is a propositional variable and Φ is an FL _{μ} -formula of schema σ . \square*

Instead of formally defining the semantics of FL _{μ} , we just give a translation of FL _{μ} into MLFP². For a σ -labelled tree T and a node $t \in \mathcal{U}^T$, we write $(T, t) \models \Phi$ to denote that the FL _{μ} -formula Φ holds at t in T . We identify propositional variables in FL _{μ} -formulas and set variables in MLFP²-formulas. For every FL _{μ} -formula Φ we define an MLFP²-formula $\varphi(x)$ in such a way that for all trees T and nodes $t \in \mathcal{U}^T$ we have $(T, t) \models \Phi \iff T \models \varphi(t)$.

- If $\Phi = \text{true}$ then $\varphi(x) := (x = x)$. If $\Phi = \text{false}$ then $\varphi(x) := \neg(x = x)$.
- If $\Phi = P$ then $\varphi(x) := P(x)$. 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 $R \in \{1^{st}Child, 2^{nd}Child\}$ 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 $\Phi := \langle R^{-1} \rangle . \Psi$ for $R \in \{1^{st}Child, 2^{nd}Child\}$ then $\varphi(x) := \exists y (R(y, x) \wedge \psi(y))$.
- If $\Phi := [R] . \Psi$ for $R \in \{1^{st}Child, 2^{nd}Child\}$ then $\varphi(x) := \forall y (R(x, y) \rightarrow \psi(y))$. Similarly, if $\Phi := [R^{-1}] . \Psi$ for $R \in \{1^{st}Child, 2^{nd}Child\}$ then $\varphi(x) := \forall y (R(y, x) \rightarrow \psi(y))$.
- If $\Phi := \mu X . \Psi$ then $\varphi(x) := [\text{LFP}_{x, X} \psi(x)](x)$.
- If $\Phi := \nu X . \Psi$ then $\varphi(x) := [\text{GFP}_{x, X} \psi(x)](x)$.

What we have seen above is that FL_μ is $\mathcal{O}(m)$ -succinct in MLFP^2 . Our next result is that there also is a reverse translation from MLFP^2 to FL_μ which only incurs an exponential blow-up in size:

Theorem 4.5. *MLFP^2 is $2^{\text{poly}(m)}$ -succinct in FL_μ on the class of labelled trees. More precisely: There is a number $c \in \mathbb{N}$ such that for every MLFP^2 -formula $\varphi(x)$ there is an FL_μ -formula Φ of size $\|\Phi\| \leq 2^{\|\varphi\|^c}$ such that for all labelled trees T and all nodes $t \in \mathcal{U}^T$, $T \models \varphi(t)$ iff $(T, t) \models \Phi$. \square*

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

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 MLFP^2 -formula $\varphi(x)$ there is a nondeterministic tree automaton of size $2^{2^{\text{poly}(\|\varphi\|)}}$ that accepts exactly those labelled trees in which φ holds at the root. \square*

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 τ be a signature and let $\varphi_1(x_1, X_1, \dots, X_n), \dots, \varphi_n(x_n, X_1, \dots, X_n)$ be formulas over the signature τ , each of which is positive in all the variables X_1, \dots, X_n . For every τ -structure \mathcal{A} every formula φ_i defines a monotone operator $F_{\mathcal{A}, \varphi_i} : (2^{\mathcal{U}^{\mathcal{A}}})^n \rightarrow 2^{\mathcal{U}^{\mathcal{A}}}$ via $F_{\mathcal{A}, \varphi_i}(A_1, \dots, A_n) := \{a \in \mathcal{U}^{\mathcal{A}} : \mathcal{A} \models \varphi_i(a, A_1, \dots, A_n)\}$. A tuple (A_1, \dots, A_n) is called a *simultaneous fixed point* of $(\varphi_1, \dots, \varphi_n)$ in \mathcal{A} iff, for all $i \leq n$, $F_{\mathcal{A}, \varphi_i}(A_1, \dots, A_n) = A_i$.

For all $i \leq n$ and $s \in \mathbb{N}$ we define $L_{\mathcal{A}, \varphi_i}^0 := \emptyset$ and $L_{\mathcal{A}, \varphi_i}^{s+1} := F_{\mathcal{A}, \varphi_i}(L_{\mathcal{A}, \varphi_1}^s, \dots, L_{\mathcal{A}, \varphi_n}^s)$. 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 $i \leq n$. We define $(L_{\mathcal{A}, \varphi_1}^\infty, \dots, L_{\mathcal{A}, \varphi_n}^\infty)$ to be the *simultaneous least fixed point* of $(\varphi_1, \dots, \varphi_n)$ in \mathcal{A} . It is straightforward to see that the simultaneous least fixed point is included in every simultaneous fixed point (A_1, \dots, A_n) of $(\varphi_1, \dots, \varphi_n)$ in \mathcal{A} , i.e., $L_{\mathcal{A}, \varphi_i}^\infty \subseteq A_i$, for all $i \leq n$ (cf., e.g., [EF99, Lemma 8.1.17]).

The logic $\text{Sim-MLFP}(\tau)$ is the extension of $\text{MLFP}(\tau)$ by simultaneous least fixed point operators. I.e.: $\text{Sim-MLFP}(\tau)$ contains $\text{MLFP}(\tau)$ and is closed under Boolean connectives and first-order quantifications; and if $n \in \mathbb{N}$ and $\varphi_1(x_1, X_1, \dots, X_n, \bar{y}, \bar{Y}), \dots, \varphi_n(x_n, X_1, \dots, X_n, \bar{y}, \bar{Y})$ are $\text{Sim-MLFP}(\tau)$ -formulas each of which is positive in the variables X_1, \dots, X_n then, for every $i \leq n$,

$$[\text{Sim-LFP}_{x_1, X_1, \dots, x_n, X_n} \varphi_1, \dots, \varphi_n]_{X_i}(x)$$

is a $\text{Sim-MLFP}(\tau)$ -formula such that for every $(\tau \cup \{\bar{y}, \bar{Y}\})$ -structure \mathcal{A} and every element $a \in \mathcal{U}^{\mathcal{A}}$ we have $\mathcal{A} \models [\text{Sim-LFP}_{x_1, X_1, \dots, x_n, X_n} \varphi_1, \dots, \varphi_n]_{X_i}(a)$ iff $a \in L_{\mathcal{A}, \varphi_i}^\infty$.

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

Every simultaneous least fixed point can be expressed by an MLFP-formula of exponential size:

Proposition 5.2. *Let τ be a signature, let $n, s \in \mathbb{N}$, and let, for each $i \leq n$, $\varphi_i(x_i, X_1, \dots, X_n)$ be an MLFP(τ)-formula of size $\|\varphi_i\| \leq s$ and positive in X_1, \dots, X_n . For every $i \leq n$ there is a MLFP(τ)-formula $\Phi_{X_i}(x)$ of size $\|\Phi_{X_i}\| \leq s^n$ such that, for all τ -structures \mathcal{A} , $\mathcal{A} \models \forall x \Phi_{X_i}(x) \leftrightarrow [\text{Sim-LFP}_{x_1, X_1, \dots, x_n, X_n} \varphi_1, \dots, \varphi_n]_{X_i}(x)$. Furthermore, if $\varphi_1, \dots, \varphi_n$ are in MLFP^2 , then also Φ_{X_i} is, for every $i \leq n$. \square*

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 $2^{2^{O(m)}}$ -succinct in MLFP on the class of all finite structures. \square*

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 $\|\gamma\|^{g^{(i)}(n)}$ (where γ is the input formula and n the number of propositional variables occurring in γ), MSO is not $\text{Tower}(o(m))$ -succinct in Sim-MLFP on the class of all finite strings. \square*

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 IDB-predicates (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 σ may use as EDB-predicates (i.e., extensional predicates which are determined by the structure the program is interpreted over) the predicates in τ_{Trees} , the predicates in σ , and a predicate $\neg P$ for every $P \in \sigma$ which is interpreted as the complement of P . We use $\text{IDB}(\mathcal{P})$ 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 σ is a finite set of rules of the form $X(x) \leftarrow \gamma(x, \bar{y})$, where γ is a conjunction of atomic formulas over the signature $\tau_{\text{Trees}} \cup \sigma \cup \{\neg P : P \in \sigma\} \cup \text{IDB}(\mathcal{P})$. 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 σ with $\text{IDB}(\mathcal{P}) = \{X_1, \dots, X_n\}$, 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 $i \leq n$. Let $\gamma_{i,1}(x, \bar{y}), \dots, \gamma_{i,\ell_i}(x, \bar{y})$ be a list of all bodies of rules with head $X_i(x)$ in \mathcal{P} . Of course, the FO-formula $\varphi_i(x, X_1, \dots, X_n) := \bigvee_{j=1}^{\ell_i} \exists \bar{y} \gamma_{i,j}(x, \bar{y})$ is positive in all the variables X_1, \dots, X_n . When evaluated in a σ -labelled tree T , the program \mathcal{P} defines the unary relations $(X_i)_{\mathcal{P}}^{\infty}(T)$, for $i \leq n$, to be the simultaneous least fixed point of $(\varphi_1, \dots, \varphi_n)$ in T .

A MonDatalog -program \mathcal{P} of schema σ , together with a designated goal predicate $X \in \text{IDB}(\mathcal{P})$, defines the *unary query* which yields, for every σ -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 $\text{IDB}(\mathcal{P}) = \{X_1, \dots, X_n\}$ there are FO-formulas $\varphi_i(x, X_1, \dots, X_n)$, for every $i \in \{1, \dots, n\}$, positive in X_1, \dots, X_n , such that $\|\varphi_{i1}\| + \dots + \|\varphi_n\| \leq \|\mathcal{P}\|$ and $[\text{Sim-LFP}_{x, X_1, \dots, x, X_n} \varphi_1, \dots, \varphi_n]_{X_i}(x)$ defines the same unary query as (\mathcal{P}, X_i) on the class of all labelled trees. \square*

Gottlob and Koch [GK02] proposed the following useful normal form for monadic datalog:

Definition 6.2 (TMNF). *A MonDatalog -program \mathcal{P} of schema σ is in tree marking normal form (TMNF, for short) iff each rule is of one of the four forms*

1. $X(x) \leftarrow S(x)$,

2. $X(x) \leftarrow X'(x) \wedge X''(x)$,
3. $X(x) \leftarrow R(x, y) \wedge X'(y)$,
4. $X(x) \leftarrow R(y, x) \wedge X'(y)$, where

$X, X', X'' \in \text{IDB}(\mathcal{P})$, $R \in \{1^{\text{st}}\text{Child}, 2^{\text{nd}}\text{Child}\}$, and $S \in \sigma \cup \{\neg P : P \in \sigma\} \cup \{\text{Root}, \text{Has-No-1}^{\text{st}}\text{Child}, \text{Has-No-2}^{\text{nd}}\text{Child}\}$. \square

Theorem 6.3 ([GK02, Theorem 4.11]). For every MonDatalog-program \mathcal{P} and every $X \in \text{IDB}(\mathcal{P})$ there is a TMNF-program \mathcal{P}_X of size $\mathcal{O}(\|\mathcal{P}\|)$ with $X \in \text{IDB}(\mathcal{P}_X)$ such that (\mathcal{P}_X, X) defines the same unary query as (\mathcal{P}, X) on the class of labelled trees. \square

One way of adding negation to datalog is to consider stratified datalog (cf., [AHV95]): A *stratified monadic datalog* program \mathcal{P} of schema σ is a finite set of rules of the form $X(x) \leftarrow \gamma(x, \vec{y})$, where γ is a conjunction of atomic formulas over the signature $\tau_{\text{Trees}} \cup \sigma \cup \{\neg P : P \in \sigma\} \cup \text{IDB}(\mathcal{P}) \cup \{\neg X : X \in \text{IDB}(\mathcal{P})\}$ that has the following property: There is a partition of \mathcal{P} into sets $\mathcal{P}_1, \dots, \mathcal{P}_n$, for some $n \in \mathbb{N}$, such that each \mathcal{P}_i is a MonDatalog-program of schema $\sigma \cup \bigcup_{j < i} \text{IDB}(\mathcal{P}_j)$. The programs $\mathcal{P}_1, \dots, \mathcal{P}_n$ are called the *strata* of \mathcal{P} , and $(\mathcal{P}_1, \dots, \mathcal{P}_n)$ 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, FL_μ , 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 σ be a schema. A *selecting σ -tree automaton* (σ -STA, for short) is a tuple $\mathfrak{A} = (Q, 2^\sigma, F, \delta, S)$, where $S \subseteq Q$ is the set of selecting states and $(Q, 2^\sigma, F, \delta)$ is a conventional nondeterministic bottom-up tree automaton (cf., e.g., [Tho96]) with finite state space Q , input alphabet 2^σ , accepting states $F \subseteq Q$, and transition function

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

A run of \mathfrak{A} on a σ -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 $\rho(t) = \delta(\text{label}(t))$; if $1^{\text{st}}\text{Child}(t, t_1) \wedge \text{Has-No-2}^{\text{nd}}\text{Child}(t)$ then $\rho(t) \in \delta(1, \rho(t_1), \text{label}(t))$; if $2^{\text{nd}}\text{Child}(t, t_2) \wedge \text{Has-No-1}^{\text{st}}\text{Child}(t)$ then $\rho(t) \in \delta(2, \rho(t_2), \text{label}(t))$; if $1^{\text{st}}\text{Child}(t, t_1) \wedge 2^{\text{nd}}\text{Child}(t, t_2)$ then $\rho(t) \in \delta(\rho(t_1), \rho(t_2), \text{label}(t))$.

A run ρ 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 σ -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 ρ of \mathfrak{A} on T . \square

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 σ , and let $X \in \text{IDB}(\mathcal{P})$. There is a σ -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 $2^{\mathcal{O}(\|\mathcal{P}\|)}$. \square

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 $w_1, w_2 \in \{0, 1\}^k$, $(X^{ok})_{\mathcal{P}_k}^\infty(w_1 2 w_2)$ 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:

$$\begin{aligned} X_0^{transfer}(x) &\leftarrow P_0(y_1) \wedge Succ(y_1, y_2) \wedge \cdots \wedge Succ(y_{k-1}, y_k) \wedge Succ(y_k, x) \\ X_1^{transfer}(x) &\leftarrow P_1(y_1) \wedge Succ(y_1, y_2) \wedge \cdots \wedge Succ(y_{k-1}, y_k) \wedge Succ(y_k, x) \\ X^{compare}(x) &\leftarrow P_0(x) \wedge X_0^{transfer}(x) \\ X^{compare}(x) &\leftarrow P_1(x) \wedge X_1^{transfer}(x) \\ X^{ok}(x) &\leftarrow P_2(x) \\ X^{ok}(x) &\leftarrow X^{ok}(y) \wedge Succ(y, x) \wedge X^{compare}(x). \end{aligned} \quad \square$$

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}' of size polynomial in $\|\mathcal{P}\|$.

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

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}' of size $2^{\mathcal{O}(\|\mathcal{P}\|)}$ with $IDB(\mathcal{P}') \supseteq \{X, \overline{X} : X \in IDB(\mathcal{P})\}$ such that, for every $X \in IDB(\mathcal{P})$, (\mathcal{P}', X) defines the same unary query as (\mathcal{P}, X) , and $(\mathcal{P}', \overline{X})$ defines the complement of the unary query defined by (\mathcal{P}, X) on the class of labelled trees.* \square

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}' of size $2^{\mathcal{O}(\|\mathcal{P}\|)}$ such that $IDB(\mathcal{P}') \supseteq IDB(\mathcal{P})$ and, for every $X \in IDB(\mathcal{P})$, (\mathcal{P}', X) defines the same unary query as (\mathcal{P}, X) on the class of all labelled trees. Given \mathcal{P} , the program \mathcal{P}' can be computed in time $2^{\mathcal{O}(\|\mathcal{P}\|)}$.* \square

6.3 S-MonDatalog vs FL_μ

From Theorem 4.6 and Lemma 6.7 one directly obtains

Theorem 6.10. *FL_μ is $2^{poly(m)}$ -succinct in S-MonDatalog on the class of labelled trees. More precisely: For every FL_μ -formula Φ there is an S-MonDatalog-program \mathcal{P} of size $2^{poly(\|\Phi\|)}$ and a predicate $X \in IDB(\mathcal{P})$, 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 Φ .* \square

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 FL_μ on the class of labelled trees. More precisely: For every S-MonDatalog-program \mathcal{P} and every $X \in IDB(\mathcal{P})$, there is a FL_μ -formula Φ_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.* \square

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 $Tower(\mathcal{O}(m))$ -succinct in S-MonDatalog on the class of labelled trees.* \square

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 $f(m) \leq Tower(o(m))$ such that for every $FO(<)$ -sentence φ there is a nondeterministic finite automaton \mathfrak{A} with at most $f(\|\varphi\|)$ states that accepts exactly those strings that satisfy φ .* \square

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 $f(m) \leq Tower(o(m))$ such that for every $FO(<)$ -sentence φ there is a S-MonDatalog-program \mathcal{P} of size $\|\mathcal{P}\| \leq f(\|\varphi\|)$ and a designated goal predicate $X \in IDB(\mathcal{P})$ such that (\mathcal{P}, X) defines the same Boolean query as φ on the class of all finite strings.* \square

Since $FO(<)$ 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.* \square

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 $FO(<)$ with $FO^k(<)$, 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.* \square

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 μ -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 $k \geq 3$? 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.
- [AI00] 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.
- [DGH⁺02] 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. Suciú 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 $\Psi(x_1, \dots, x_k, X_1, \dots, X_\ell)$, a string w , a sequence p_1, \dots, p_k of positions in w , and sets P_1, \dots, P_ℓ of positions in w , \mathcal{A} decides whether $w \models \Psi(p_1, \dots, p_k, P_1, \dots, P_\ell)$.

\mathcal{A} operates by recursion on the construction of Ψ . The only non-trivial case is when Ψ is of the form $[\text{LFP}_{y,Y} \Psi'(y, Y, x_1, \dots, x_k, X_1, \dots, X_\ell)](x_i)$. In this case, \mathcal{A} computes the stages of the least fixed point of Ψ' . 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 $w \models \Psi'(q, L^{s-1}, p_1, \dots, p_k, P_1, \dots, P_\ell)$;
 - 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 $\mathcal{O}(|w|^2 \cdot |w|^{\mathcal{O}(\|\Psi'\|)}) = |w|^{\mathcal{O}(\|\Psi\|)}$ steps. ■

Proof of Lemma 3.6.

First, use the algorithm of Lemma 3.5 (b) to construct the FO($<$)-formula $\varphi_h(Z)$.

It is straightforward to verify that the following formula has the desired properties:

$$\Phi_h(z) := \exists Z \left((\forall x Zx \rightarrow P_*x) \wedge Zz \wedge \varphi_h(Z) \wedge \forall Z' ((\forall x Z'x \rightarrow P_*x) \wedge \varphi_h(Z')) \rightarrow Z \leq_{\text{lex}} Z' \right),$$

where $Z \leq_{\text{lex}} Z'$ is an abbreviation for the FO($<$)-formula

$$(\forall x Zx \leftrightarrow Z'x) \vee \exists y (\neg Zy \wedge Z'y \wedge \forall x (x < y \rightarrow (Zx \leftrightarrow Z'x))).$$

Afterwards, we replace every occurrence of an atom of the form $x < y$ in $\Phi_h(z)$ by the MSO-formula $\exists Y (Yy \wedge \neg Yx \wedge \forall z_1 \forall z_2 (Succ(z_1, z_2) \wedge Yz_1) \rightarrow Yz_2)$. ■

Proof of Lemma 3.7.

Since f has bound $f(m) \leq \text{Tower}(o(m))$, 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 $f(m) \leq \text{Tower}\left(\frac{m}{g(m)}\right)$, for all $m \geq m_0$. Therefore, for $h := \lg^*(n)$ and $m := c \cdot h$ we have

$$d \cdot f(m) \cdot \lg(f(m)) \leq d \cdot \text{Tower}\left(\frac{m}{g(m)}\right) \cdot \text{Tower}\left(\frac{m}{g(m)} - 1\right) \leq \text{Tower}\left(\frac{m}{g(m)} + 1\right).$$

On the other hand, $n \geq \text{Tower}(h-1)$, and thus $\lg^{(i)}(n) \geq \text{Tower}(h-(i+1))$. It therefore suffices to show that $\frac{m}{g(m)} + 1 \leq h-(i+1)$, i.e., that $\frac{m}{g(m)} \leq h-(i+2)$. Since $g \in \omega(1)$, there is some $m_1 \geq m_0$ such that for all $m \geq m_1$ we have $g(m) \geq c \cdot (i+3)$. Hence,

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

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 φ is in negation normal form. Let x_1, \dots, x_k be the first-order variables occurring in φ , and let X_1, \dots, X_k be set variables *not* occurring in φ . We claim that for every subformula $\psi(x_{i_1}, \dots, x_{i_\ell}, \overline{Y})$ there is a formula $\psi'(X_{i_1}, \dots, X_{i_\ell}, \overline{Y})$ such that ψ' is positive in $X_{i_1}, \dots, X_{i_\ell}, \overline{Y}$ and for all structures \mathcal{A} and for all $a_1, \dots, a_\ell \in \mathcal{U}^{\mathcal{A}}, \overline{B} \subseteq \mathcal{U}^{\mathcal{A}}$ we have:

$$(*) \quad \mathcal{A} \models \psi'(\{a_1\}, \dots, \{a_\ell\}, \overline{B}) \iff \mathcal{A} \models \psi(a_1, \dots, a_\ell, \overline{B}).$$

ψ' is defined inductively as follows:

- If ψ is a positive atom, say, $\psi = R x_1 x_2$, then we let $\psi'(X_1, X_2) := \exists x \in X_1 \exists y \in X_2 R x y$.
- If ψ is a negated atom, say, $\psi = \neg R x_1 x_2$, then we let $\psi'(X_1, X_2) := \neg \forall x \in X_1 \forall y \in X_2 R x y$.
- If $\psi = \psi_1 \vee \psi_2$, then we let $\psi' := \psi'_1 \vee \psi'_2$. Conjunctions are handled analogously.
- If ψ starts with an existential quantifier, say, $\psi = \exists x_\ell \chi(x_1, \dots, x_\ell, \overline{Y})$, then we let

$$\psi' := \exists y \forall x [\text{LFP}_{x, X_\ell} x=y \vee \chi'(X_1, \dots, X_\ell, \overline{Y})](x).$$

Similarly, if ψ starts with a universal quantifier, say, $\psi = \forall x_\ell \chi(x_1, \dots, x_\ell, \overline{Y})$, then we let

$$\psi' := \forall y \forall x [\text{LFP}_{x, X_\ell} x=y \vee \chi'(X_1, \dots, X_\ell, \overline{Y})](x).$$

- If ψ starts with a least fixed point operator, say, $\psi = [\text{LFP}_{x_\ell, Y} \chi(x_1, \dots, x_\ell, Y, \overline{Y})](x_\ell)$ (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

$$\psi' := \exists x (X_\ell(x) \wedge [\text{LFP}_{y, Y} \forall x [\text{LFP}_{x, X_\ell} x=y \vee \chi'(X_1, \dots, X_\ell, Y, \overline{Y})](x)](x)).$$

The case that φ starts with a greatest fixed point operator can be treated similarly.

To see that the resulting sentence φ' is indeed equivalent to φ , the crucial observation is the following: Suppose that $\chi(x_1, \dots, x_\ell, \overline{Y})$ and $\chi'(X_1, \dots, X_\ell, \overline{Y})$ satisfy (*). Let

$$\chi''(X_1, \dots, X_{\ell-1}, y, \overline{Y}) := \forall x [\text{LFP}_{x, X_\ell} x=y \vee \chi'(X_1, \dots, X_\ell, \overline{Y})](x).$$

Then for all structures \mathcal{A} and for all $a_1, \dots, a_\ell \in \mathcal{U}^{\mathcal{A}}, \overline{B} \subseteq \mathcal{U}^{\mathcal{A}}$ we have:

$$(**) \quad \mathcal{A} \models \chi''(\{a_1\}, \dots, \{a_{\ell-1}\}, a_\ell, \overline{B}) \iff \mathcal{A} \models \chi(a_1, \dots, a_\ell, \overline{B}).$$

To see this, let $a_1, \dots, a_\ell \in \mathcal{U}^{\mathcal{A}}$ and $\overline{B} \subseteq \mathcal{U}^{\mathcal{A}}$. Suppose first that $\mathcal{A} \models \chi''(\{a_1\}, \dots, \{a_{\ell-1}\}, a_\ell, \overline{B})$. Then $\mathcal{A} \models \chi'(\{a_1\}, \dots, \{a_{\ell-1}\}, \emptyset, \overline{B})$ or $\mathcal{A} \models \chi'(\{a_1\}, \dots, \{a_{\ell-1}\}, \{a_\ell\}, \overline{B})$, because otherwise the fixed point process associated with χ'' remains empty. Since X_ℓ only occurs positively in χ' , $\mathcal{A} \models \chi'(\{a_1\}, \dots, \{a_{\ell-1}\}, \emptyset, \overline{B})$ implies that $\mathcal{A} \models \chi'(\{a_1\}, \dots, \{a_{\ell-1}\}, \{a_\ell\}, \overline{B})$, so we may assume the latter. Then by (*), $\mathcal{A} \models \chi(a_1, \dots, a_\ell, \overline{B})$.

Conversely, if $\mathcal{A} \models \chi(a_1, \dots, a_\ell, \overline{B})$ then, by (*), $\mathcal{A} \models \chi'(\{a_1\}, \dots, \{a_{\ell-1}\}, \{a_\ell\}, \overline{B})$. But this implies that all elements are in X in the second stage of the fixed point process associated with χ'' and hence $\mathcal{A} \models \chi''(\{a_1\}, \dots, \{a_{\ell-1}\}, a_\ell, \overline{B})$.

Now we are ready to prove (*) by induction on ψ . 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 $\psi = [\text{LFP}_{x_\ell, Y} \chi(x_1, \dots, x_\ell, Y, \overline{Y})](x_\ell)$. Then

$$\psi' = \exists x (X_\ell(x) \wedge [\text{LFP}_{y, Y} \chi''(X_1, \dots, X_{\ell-1}, y, Y, \overline{Y})](x)).$$

Let $a_1, \dots, a_{\ell-1} \in \mathcal{U}^A$ and $\overline{B} \subseteq \mathcal{U}^A$. By (**), for all $B \subseteq \mathcal{U}^A$ we have

$$\left\{ a \in \mathcal{U}^A \mid \mathcal{A} \models \chi''(\{a_1\}, \dots, \{a_{\ell-1}\}, a, B, \overline{B}) \right\} = \left\{ a \in \mathcal{U}^A \mid \mathcal{A} \models \chi(a_1, \dots, a_{\ell-1}, a, B, \overline{B}) \right\}.$$

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

$$\mathcal{A} \models [\text{LFP}_{y,Y} \chi''(\{a_1\}, \dots, \{a_{\ell-1}\}, y, \overline{B})](a) \iff \mathcal{A} \models [\text{LFP}_{y,Y} \chi(a_1, \dots, a_{\ell-1}, y, \overline{B})](a).$$

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

Proof of Theorem 4.5.

Let the *arity* of a formula be the number of free variables it has.

We say that an MLFP²-formula $\varphi(x, \overline{X})$ of the form $\exists y \psi(x, y, \overline{X})$ is in *normal form* if it is in negation normal form and $\psi(x, y, \overline{X})$ is a disjunction of formulas of the form $\alpha \wedge \theta$ satisfying the following conditions:

- (i) θ is a conjunction of at most unary formulas.
- (ii) α is one of the following five *basic binary formulas*:

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

Similarly, a formula of the form $\forall y \psi(x, y, \overline{X})$ is in normal form if it is a conjunction of formulas δ of the form $\alpha \rightarrow \theta$, where

- (i') θ is a disjunction of at most unary formulas.
- (ii') α is one of the basic binary formulas $\alpha^1, \dots, \alpha^5$.

Finally, an arbitrary MLFP²-formula is in normal form if all its subformulas of the form $\exists y \psi$ and $\forall y \psi$ are.

Step 1: For every MLFP²-formula φ we construct an equivalent MLFP²-formula φ' in normal form. The construction is by induction on φ , 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 $\varphi(x, \overline{X}) = \forall y \psi(x, y, \overline{X})$. We may view ψ as a Boolean combination of atomic formulas and at most unary formulas. We bring this Boolean combination into conjunctive normal form; let ψ' be the resulting formula. Now let us look at the clauses γ of ψ' . Let θ be the disjunction of all at most unary literals in γ , and let α' be the conjunction of the negations of the remaining literals. Then we may replace γ with $(\alpha' \rightarrow \theta)$.

α' is a conjunction of literals of the form $(\neg)R(x, y)$ and $(\neg)R(y, x)$, where $R \in \{1^{\text{st}} \text{Child}, 2^{\text{nd}} \text{Child}\}$. If α' is not satisfiable in any tree, then we may simply discard it from ψ' , because $(\alpha' \rightarrow \theta)$ always holds. If this makes ψ' empty, we replace the whole formula φ by $(x=x)$. So let us assume that α' is satisfiable in some tree. Then α' is equivalent to a disjunction α'' of some of the five basic binary formulas. Say, α' is equivalent to $\bigvee_{j=1}^{\ell} \alpha^{i_j}$, for some $\ell \leq 5$. Then we replace γ by the clauses $(\alpha^{i_1} \rightarrow \theta), \dots, (\alpha^{i_\ell} \rightarrow \theta)$.

We do this for all clauses of ψ' and obtain a formula ψ'' equivalent to ψ that has the desired form. This completes Step 1.

Before we translate MLFP²-formulas in normal form into FL _{μ} -formulas, we define a few auxiliary MLFP²-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 FL _{μ} -formula Φ we construct FL _{μ} -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. $(T, t) \models \exists\Phi$ if, and only if, there exists a $u \in \mathcal{U}^T$ such that $(T, u) \models \Phi$.
2. $(T, t) \models \exists\downarrow\Phi$ if, and only if, there exists a $u \in \mathcal{U}^{T_t}$ such that $(T, u) \models \Phi$.
3. $(T, t) \models \exists\uparrow\Phi$ if, and only if, there exists a $u \in \mathcal{U}^T \setminus \mathcal{U}^{T_t}$ 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 MLFP²-formula φ in normal form we construct an FL _{μ} -formula Φ such that for all labelled trees T and all nodes $t \in \mathcal{U}^T$, $T \models \varphi(t) \iff (T, t) \models \Phi$.

The construction is by induction on φ .

For all unary subformulas that do not start with an existential or universal quantifier, we can simply revert the translation from FL _{μ} to MLFP².

For 0-ary formulas $\varphi = \exists x\psi$, suppose Ψ is the FL _{μ} -formula already constructed for ψ . 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 $\psi = \bigvee_{i=1}^{\ell} (\alpha_i \wedge \theta_i)$, where θ_i and α_i satisfy the conditions (i) and (ii) for all i . For every i we define a formula Γ_i equivalent to $\exists y(\alpha_i \wedge \theta_i)$, and then we let $\Phi = \bigvee_{i=1}^{\ell} \Gamma_i$.

So let $1 \leq i \leq \ell$. Let θ_x be the conjunction of all at most unary formulas of $(\alpha_i \wedge \theta_i)$ that either do not have any free first-order variables or whose only free first-order variable is x , and let θ_y be the conjunction of all at most unary subformulas of $(\alpha_i \wedge \theta_i)$ whose only free first-order variable is y . If $\alpha_i = R(x, y)$ for $R \in \{1^{st}Child, 2^{nd}Child\}$, we let $\Gamma_i := \theta_x \wedge \langle R \rangle \theta_y$, where θ_x and θ_y are the FL _{μ} -formulas already constructed for θ_x and θ_y . Similarly, if $\alpha_i = R(y, x)$, we let $\Gamma_i := \theta_x \wedge \langle R^{-1} \rangle \theta_y$. If $\alpha_i = \neg 1^{st}Child(x, y) \wedge \neg 2^{nd}Child(x, y) \wedge \neg 1^{st}Child(y, x) \wedge \neg 2^{nd}Child(y, x)$, we let

$$\Gamma_i := \theta_x \wedge \left(\langle 1^{st}Child \rangle \langle 1^{st}Child \rangle \exists\downarrow\theta_y \vee \langle 1^{st}Child \rangle \langle 2^{nd}Child \rangle \exists\downarrow\theta_y \vee \langle 2^{nd}Child \rangle \langle 1^{st}Child \rangle \exists\downarrow\theta_y \vee \langle 2^{nd}Child \rangle \langle 2^{nd}Child \rangle \exists\downarrow\theta_y \vee \langle 1^{st}Child^{-1} \rangle \exists\uparrow\theta_y \vee \langle 2^{nd}Child^{-1} \rangle \exists\uparrow\theta_y \right).$$

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 Φ is at most exponential in the size of φ . 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 ψ of φ 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 ψ has either of these forms. Let ξ_1, \dots, ξ_m be the at most unary subformulas of χ that are not atoms such that χ can be written as a Boolean combination of ξ_1, \dots, ξ_m and atomic subformulas. Let Ξ_1, \dots, Ξ_m be the FL _{μ} -formulas corresponding to ξ_1, \dots, ξ_m . By induction hypothesis, for $1 \leq i \leq m$ we have

$$\|\Xi_i\| \leq 2^{\mathcal{O}((d-1) \cdot \|\xi_i\|)} \leq 2^{\mathcal{O}((d-1) \cdot \|\psi\|)}.$$

An inspection of our construction shows that

$$\|\Psi\| \leq 2^{\mathcal{O}(\|\psi\|)} \max_{1 \leq i \leq m} \|\Xi_i\| \leq 2^{\mathcal{O}(\|\psi\|)} 2^{\mathcal{O}((d-1) \cdot \|\psi\|)} = 2^{\mathcal{O}(d \cdot \|\psi\|)}.$$

■

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 $i \leq n$, the induction hypothesis

gives us for every $j \in \{1, \dots, n\} \setminus \{i\}$ a MLFP($\tau \cup \{X_i\}$)-formula $\tilde{\Phi}_{X_j}(x)$ of size $\leq s^{n-1}$ such that, for all $\tau \cup \{X_i\}$ -structures \mathcal{B} , $\mathcal{B} \models \forall x \tilde{\Phi}_{X_j}(x) \leftrightarrow$

$$[\text{Sim-LFP}_{x_1, X_1, \dots, x_{i-1}, X_{i-1}, x_{i+1}, X_{i+1}, \dots, x_n, X_n} \varphi_1, \dots, \varphi_{i-1}, \varphi_{i+1}, \dots, \varphi_n]_{X_j}(x).$$

Define $\tilde{\varphi}_i(x_i, X_i)$ to be the MLFP-formula obtained from $\varphi_i(x_i, X_1, \dots, X_n)$ 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 $\|\tilde{\varphi}_i\| < s^n$. To complete the proof of Proposition 5.2 it therefore suffices to show the following

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

$$\mathcal{A} \models \forall x [\text{LFP}_{x_i, X_i} \tilde{\varphi}_i](x) \leftrightarrow [\text{Sim-LFP}_{x_1, X_1, \dots, x_n, X_n} \varphi_1, \dots, \varphi_n]_{X_i}(x).$$

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

$$\begin{aligned} A_j(S) &:= \{a \in \mathcal{U}^{\mathcal{A}} : \langle \mathcal{A}, S \rangle \models \tilde{\Phi}_{X_j}(a)\}, \\ A_i^{(\ell)} &:= \{a \in \mathcal{U}^{\mathcal{A}} : \mathcal{A} \models [\text{LFP}_{x_i, X_i} \tilde{\varphi}_i](a)\}, \\ A_i^{(r)} &:= \{a \in \mathcal{U}^{\mathcal{A}} : \mathcal{A} \models [\text{Sim-LFP}_{x_1, X_1, \dots, x_n, X_n} \varphi_1, \dots, \varphi_n]_{X_i}(a)\}, \\ A_j^{(r)} &:= \{a \in \mathcal{U}^{\mathcal{A}} : \mathcal{A} \models [\text{Sim-LFP}_{x_1, X_1, \dots, x_n, X_n} \varphi_1, \dots, \varphi_n]_{X_j}(a)\}. \end{aligned}$$

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 $A_i^{(\ell)} =$

$$\{a \in \mathcal{U}^{\mathcal{A}} : \mathcal{A} \models \varphi_i(a, A_1(A_i^{(\ell)}), \dots, A_{i-1}(A_i^{(\ell)}), A_i^{(\ell)}, A_{i+1}(A_i^{(\ell)}), \dots, A_n(A_i^{(\ell)}))\}.$$

From the definition of $\tilde{\Phi}_{X_j}$ we furthermore know that $A_j(A_i^{(\ell)}) =$

$$\{a \in \mathcal{U}^{\mathcal{A}} : \mathcal{A} \models \varphi_j(a, A_1(A_i^{(\ell)}), \dots, A_{i-1}(A_i^{(\ell)}), A_i^{(\ell)}, A_{i+1}(A_i^{(\ell)}), \dots, A_n(A_i^{(\ell)}))\}.$$

I.e., $(A_1(A_i^{(\ell)}), \dots, A_{i-1}(A_i^{(\ell)}), A_i^{(\ell)}, A_{i+1}(A_i^{(\ell)}), \dots, A_n(A_i^{(\ell)}))$ is a simultaneous fixed point of the formulas $(\varphi_1, \dots, \varphi_n)$ in \mathcal{A} , the *least* fixed point of which is $(A_1^{(r)}, \dots, A_n^{(r)})$. In particular, this implies that $A_i^{(\ell)} \supseteq A_i^{(r)}$.

On the other hand, $(A_1^{(r)}, \dots, A_{i-1}^{(r)}, A_{i+1}^{(r)}, \dots, A_n^{(r)})$ is a simultaneous fixed point of $(\varphi_1, \dots, \varphi_{i-1}, \varphi_{i+1}, \dots, \varphi_n)$ in $\langle \mathcal{A}, A_i^{(r)} \rangle$, the *least* fixed point of which is $(A_1(A_i^{(r)}), \dots, A_{i-1}(A_i^{(r)}), A_{i+1}(A_i^{(r)}), \dots, A_n(A_i^{(r)}))$.

Therefore, $A_j(A_i^{(r)}) \subseteq A_j^{(r)}$, for all $j \neq i$, and hence $\{a \in \mathcal{U}^{\mathcal{A}} : \langle \mathcal{A}, A_i^{(r)} \rangle \models \tilde{\varphi}_i(a)\} =$

$$\begin{aligned} &\{a \in \mathcal{U}^{\mathcal{A}} : \mathcal{A} \models \varphi_i(a, A_1(A_i^{(r)}), \dots, A_{i-1}(A_i^{(r)}), A_i^{(r)}, A_{i+1}(A_i^{(r)}), \dots, A_n(A_i^{(r)}))\} \subseteq \\ &\{a \in \mathcal{U}^{\mathcal{A}} : \mathcal{A} \models \varphi_i(a, A_1^{(r)}, \dots, A_{i-1}^{(r)}, A_i^{(r)}, A_{i+1}^{(r)}, \dots, A_n^{(r)})\} = A_i^{(r)}. \end{aligned}$$

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}' of size $\mathcal{O}(\|\mathcal{P}\|)$ such that (\mathcal{P}', X) defines the same unary query as (\mathcal{P}, X) on the class of labelled trees. Let ℓ be the number of IDB-predicates of \mathcal{P}' . In [FGK03, Example 20], a σ -STA with 2^ℓ states was

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 .

constructed, defining the same unary query as (\mathcal{P}', X) on the class of labelled trees. \blacksquare

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 $Reach(t)$ is computed for every vertex t of T , where $Reach(t)$ consists of all states q such that there is a partial run ρ of \mathfrak{A} on T with $\rho(t) = q$. Afterwards, in a top-down pass the set $Acc(t)$ is computed for every $t \in \mathcal{U}^T$, where $Acc(t)$ consists of all states q such that there is an *accepting* run ρ 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 $Acc(t)$ contains a state in $Q \setminus S$.

To store the sets $Reach(\cdot)$ and $Acc(\cdot)$, respectively, the program \mathcal{P} will use predicates X_q^{Reach} and X_q^{Acc} , respectively, for every $q \in Q$. The intended meaning of these predicates is, that for all σ -labelled trees T , $t \in \mathcal{U}^T$, and $q \in Q$ we shall have $t \in (X_q^{Reach})_{\mathcal{P}}^{\infty}(T)$ iff $q \in Reach(t)$, and $t \in (X_q^{Acc})_{\mathcal{P}}^{\infty}(T)$ iff $q \in Acc(t)$.

\mathcal{P} has one further predicate, X^{no} 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 $Input_a(x) \leftarrow \bigwedge_{P \in a} P(x) \wedge \bigwedge_{P \in \sigma \setminus a} \neg P(x)$.
2. For every $q \in Q \setminus S$, \mathcal{P} contains the rule $X^{no}(x) \leftarrow X_q^{Acc}(x)$.
3. For all $q \in F$, \mathcal{P} contains the rule $X_q^{Acc}(x) \leftarrow X_q^{Reach}(x) \wedge Root(x)$.
4. For all $a \in 2^{\sigma}$ and $p \in \delta(a)$, \mathcal{P} contains the rule $X_p^{Reach}(x) \leftarrow Input_a(x) \wedge Has\text{-}No\text{-}1^{st}Child(x) \wedge Has\text{-}No\text{-}2^{nd}Child(x)$.
5. For all $q \in Q$, $a \in 2^{\sigma}$, and $p \in \delta(1, q, a)$, \mathcal{P} contains the rules $X_p^{Reach}(x) \leftarrow Input_a(x) \wedge 1^{st}Child(x, y) \wedge Has\text{-}No\text{-}2^{nd}Child(x) \wedge X_q^{Reach}(y)$, and $X_q^{Acc}(y) \leftarrow X_q^{Reach}(y) \wedge 1^{st}Child(x, y) \wedge Input_a(x) \wedge Has\text{-}No\text{-}2^{nd}Child(x) \wedge X_p^{Acc}(x)$.
6. For all $q \in Q$, $a \in 2^{\sigma}$, and $p \in \delta(2, q, a)$, \mathcal{P} contains the rules $X_p^{Reach}(x) \leftarrow Input_a(x) \wedge Has\text{-}No\text{-}1^{st}Child(x) \wedge 2^{nd}Child(x, y) \wedge X_q^{Reach}(y)$, and $X_q^{Acc}(y) \leftarrow X_q^{Reach}(y) \wedge 2^{nd}Child(x, y) \wedge Input_a(x) \wedge Has\text{-}No\text{-}1^{st}Child(x) \wedge X_p^{Acc}(x)$.
7. For all $q_1, q_2 \in Q$, $a \in 2^{\sigma}$ and $p \in \delta(q_1, q_2, a)$, \mathcal{P} contains the rules $X_p^{Reach}(x) \leftarrow Input_a(x) \wedge 1^{st}Child(x, y_1) \wedge X_{q_1}^{Reach}(y_1) \wedge 2^{nd}Child(x, y_2) \wedge X_{q_2}^{Reach}(y_2)$, $X_{q_1}^{Acc}(y_1) \leftarrow X_{q_1}^{Reach}(y_1) \wedge 1^{st}Child(x, y_1) \wedge 2^{nd}Child(x, y_2) \wedge Input_a(x) \wedge X_p^{Acc}(x) \wedge X_{q_2}^{Reach}(y_2)$, $X_{q_2}^{Acc}(y_2) \leftarrow X_{q_2}^{Reach}(y_2) \wedge 1^{st}Child(x, y_1) \wedge 2^{nd}Child(x, y_2) \wedge Input_a(x) \wedge X_p^{Acc}(x) \wedge X_{q_1}^{Reach}(y_1)$.

One can easily verify that (\mathcal{P}, X^{no}) defines the complement of the query defined by \mathfrak{A} on the class of σ -labelled trees. Furthermore, $\|\mathcal{P}\| = \mathcal{O}(|Q|^3 \cdot |2^{\sigma}| + |\sigma| \cdot |2^{\sigma}|)$. \blacksquare

Proof of Proposition 6.8.

Let σ be the schema of \mathcal{P} , and let $X \in IDB(\mathcal{P})$. From Theorem 6.5 we obtain a σ -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 $\overline{X} \in IDB(\overline{\mathcal{P}}_X)$ such that $(\overline{\mathcal{P}}_X, \overline{X})$ defines the complement of the query defined by \mathfrak{A}_X ; and $\|\overline{\mathcal{P}}_X\| = \mathcal{O}(|Q|^3 \cdot |2^{\sigma}| + |\sigma| \cdot |2^{\sigma}|) = 2^{\mathcal{O}(\|\mathcal{P}\|)}$.

After an appropriate renaming of IDB-predicates we can assume w.l.o.g. that $IDB(\overline{\mathcal{P}}_X) \cap IDB(\mathcal{P}) = \emptyset$ and $IDB(\overline{\mathcal{P}}_X) \cap IDB(\overline{\mathcal{P}}_Y) = \emptyset$, for all distinct $X, Y \in IDB(\mathcal{P})$.

Obviously, $\mathcal{P}' := \mathcal{P} \cup \bigcup_{X \in IDB(\mathcal{P})} \overline{\mathcal{P}}_X$ is the desired MonDatalog-program of size $\|\mathcal{P}'\| \leq \|\mathcal{P}\| + |IDB(\mathcal{P})| \cdot 2^{\mathcal{O}(\|\mathcal{P}\|)} = 2^{\mathcal{O}(\|\mathcal{P}\|)}$. \blacksquare

Proof of Theorem 6.9.

Let \mathcal{P} be a S-MonDatalog-program of schema σ . Let $(\mathcal{P}_1, \dots, \mathcal{P}_n)$ be a decomposition of \mathcal{P} into strata. I.e., $\mathcal{P} = \mathcal{P}_1 \cup \dots \cup \mathcal{P}_n$, and each \mathcal{P}_i can be viewed as a MonDatalog-program of schema $\sigma_i := \sigma \cup \bigcup_{j < i} IDB(\mathcal{P}_j)$. For each i , let $\tilde{\mathcal{P}}_i$ be the MonDatalog-program of schema $\tilde{\sigma}_i := \sigma \cup \{X, \overline{X} : X \in \bigcup_{j < i} IDB(\mathcal{P}_j)\}$, obtained from \mathcal{P}_i by replacing every occurrence of a literal of the form $\neg X(y)$ by the literal $\overline{X}(y)$. Let $\tilde{\mathcal{P}}'_i$ be the MonDatalog-program of size $2^{\mathcal{O}(\|\mathcal{P}_i\|)}$ 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, $\tilde{\mathcal{P}}_i'$ and $\tilde{\mathcal{P}}_j'$ have no IDB-predicates in common (for distinct $i, j \leq n$).

It is straightforward to check that the program $\mathcal{P}' := \bigcup_{i=1}^n \tilde{\mathcal{P}}_i'$, viewed as a MonDatalog-program of schema σ , has the desired property that $\text{IDB}(\mathcal{P}') \supseteq \text{IDB}(\mathcal{P})$ and (\mathcal{P}', X) defines the same query as (\mathcal{P}, X) , for every $X \in \text{IDB}(\mathcal{P})$.

Furthermore, $\|\mathcal{P}'\| \leq \sum_{i=1}^n 2^{\mathcal{O}(\|\mathcal{P}_i\|)} \leq 2^{\mathcal{O}(\|\mathcal{P}\|)}$.

It is straightforward to see that \mathcal{P}' can be computed in time $2^{\mathcal{O}(\|\mathcal{P}\|)}$. \blacksquare

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 $X \in \text{IDB}(\mathcal{P})$ let \mathcal{P}_X be a TMNF-program of size $\mathcal{O}(\|\mathcal{P}\|)$ with $X \in \text{IDB}(\mathcal{P}_X)$ such that (\mathcal{P}_X, X) 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 MLFP²-formula $\Psi_X(x)$ of size $\|\Psi_X\| \leq \|\mathcal{P}_X\|^n \leq \|\mathcal{P}_X\|^{\|\mathcal{P}_X\|} \leq 2^{\mathcal{O}(\|\mathcal{P}\| \cdot \lg \|\mathcal{P}\|)}$ 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 FL _{μ} -formula Φ_X of size $\mathcal{O}(\|\Psi_X\|)$.

Step 2: Now consider the case where \mathcal{P} is an S-MonDatalog-program. Let σ be the schema of \mathcal{P} , and let $(\mathcal{P}_1, \dots, \mathcal{P}_n)$ be a decomposition of \mathcal{P} into strata. I.e., $\mathcal{P} = \mathcal{P}_1 \cup \dots \cup \mathcal{P}_n$, and each \mathcal{P}_i can be viewed as a MonDatalog-program of schema $\sigma_i := \sigma \cup \bigcup_{j < i} \text{IDB}(\mathcal{P}_j)$. Step 1 gives us, for every $i \leq n$ and every $X \in \text{IDB}(\mathcal{P}_i)$, a FL _{μ} -formula Φ_X of size $\|\Phi_X\| = 2^{\mathcal{O}(\|\mathcal{P}_i\| \cdot \lg \|\mathcal{P}_i\|)}$ that defines the same unary query on the class of σ_i -labelled trees as (\mathcal{P}_i, X) .

By induction on i we define, for every $X \in \text{IDB}(\mathcal{P}_i)$, the FL _{μ} -formula $\hat{\Phi}_X$ of size

$$\|\hat{\Phi}_X\| = 2^{\mathcal{O}(\|\mathcal{P}_1\| \cdot \lg \|\mathcal{P}_1\| + \dots + \|\mathcal{P}_i\| \cdot \lg \|\mathcal{P}_i\|)}$$

as follows: For every $X \in \text{IDB}(\mathcal{P}_i)$ let $\hat{\Phi}_X$ be the formula obtained from Φ_X by replacing every occurrence of an $Y \in \bigcup_{j < i} \text{IDB}(\mathcal{P}_j)$ by the formula $\hat{\Phi}_Y$. It is straightforward to verify for every $X \in \text{IDB}(\mathcal{P})$ that $\hat{\Phi}_X$ defines the same unary query as (\mathcal{P}, X) on the class of σ -labelled trees. \blacksquare

Proof of Proposition 6.13.

We use the following notations and results of [FG03]:

For all $h \geq 1$ let $\Sigma_h := \{0, 1, <1>, </1>, \dots, <h>, </h>\}$. The “tags” $<i>$ and $</i>$ represent single letters of the alphabet and are just chosen to improve readability. For every $n \geq 1$ 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 $n \geq 2$. By $\text{bit}(i, n)$ we denote the i -th bit of the binary representation of n , i.e., $\text{bit}(i, n)$ is 1 if $\lfloor \frac{n}{2^i} \rfloor$ is odd, and $\text{bit}(i, n)$ is 0 otherwise.

We encode every number $n \in \mathbb{N}$ by a string $\mu_h(n)$ over the alphabet Σ_h , where $\mu_h(n)$ is inductively defined as follows: $\mu_1(0) := <1></1>$, and

$$\mu_h(n) := <1> \text{bit}(0, n-1) \text{bit}(1, n-1) \dots \text{bit}(L(n)-1, n-1) </1>,$$

for $n \geq 1$. For $h \geq 2$ we let $\mu_h(0) := <h></h>$ and

$$\begin{aligned} \mu_h(n) := & <h> \\ & \mu_{h-1}(0) \text{bit}(0, n-1) \\ & \mu_{h-1}(1) \text{bit}(1, n-1) \\ & \vdots \\ & \mu_{h-1}(L(n)-1) \text{bit}(L(n)-1, n-1) \\ & </h>, \end{aligned}$$

for $n \geq 1$. 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 FO($<$)-formula $\chi_{h,1}$ of size⁴ $\|\chi_{h,1}\| = \mathcal{O}(h)$ such that the following is true for all strings w over alphabet Σ_h , for all positions a, b in w , and for all numbers $m, n \in \{0, \dots, \text{Tower}(h)\}$: 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 $m = n$. \square

For every $h \in \mathbb{N}$ let $H := \text{Tower}(h)$. For $a_0 \cdots a_{H-1} \in \{0, 1\}^H$ define the string

$$w_{h, a_0 \cdots a_{H-1}} := \langle h+1 \rangle \mu_h(0) a_0 \mu_h(1) a_1 \cdots \mu_h(H-1) a_{H-1} \langle /h+1 \rangle.$$

We define the string-language

$$L_h := \{ w_{h, a_0 \cdots a_{H-1}} : a_0 \cdots a_{H-1} \in \{0, 1\}^H \}^*,$$

and we choose a designated string $w_h \in L_h$ via

$$w_h := w_{h, \text{BIN}_H(0)} w_{h, \text{BIN}_H(1)} \cdots w_{h, \text{BIN}_H(2^H-1)},$$

where $\text{BIN}_H(n)$ denotes the reverse binary representation of length H of n . For example, $\text{BIN}_4(2) = 0100$ and $\text{BIN}_4(5) = 1010$. It is straightforward to see the following:

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

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

Proof. To simplify notation we, in the following, write $\text{Succ}(x)$ 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 FO($<$)-sentence φ_h stating for every input string $w \in L_h$ that

- the leftmost substring of w of the form $\langle h+1 \rangle \cdots \langle /h+1 \rangle$ contains no position x with $P_{\langle /h \rangle}(x) \wedge P_1(\text{Succ}(x))$,
- the rightmost substring of w of the form $\langle h+1 \rangle \cdots \langle /h+1 \rangle$ contains no position y with $P_{\langle /h \rangle}(y) \wedge P_0(\text{Succ}(x))$,
- for every two successive substrings u and v of w of the form $\langle h+1 \rangle \cdots \langle /h+1 \rangle$, there is a position x_1 in u with $P_{\langle h \rangle}(x_1)$ and a position y_1 in v with $P_{\langle h \rangle}(y_1)$ such that $\chi_{h,1}(x_1, y_1)$ is true and
 - every position x_0 in u to the left of x_1 satisfies $P_{\langle /h \rangle}(x_0) \rightarrow P_1(\text{Succ}(x_0))$,
 - every position y_0 in v to the left of y_1 satisfies $P_{\langle /h \rangle}(y_0) \rightarrow P_0(\text{Succ}(y_0))$,
 - the first position x'_1 to the right of x_1 in u with $P_{\langle /h \rangle}(x'_1)$ satisfies $P_0(\text{Succ}(x'_1))$,
 - the first position y'_1 to the right of y_1 in v with $P_{\langle /h \rangle}(y'_1)$ satisfies $P_1(\text{Succ}(y'_1))$,
 - for all positions x_2 to the right of x'_1 in u and for all positions y_2 to the right of y'_1 in v such that $P_{\langle h \rangle}(x_2) \wedge P_{\langle h \rangle}(y_2) \wedge \chi_{h,1}(x_2, y_2)$ is true, also $P_1(\text{Succ}(x'_2)) \leftrightarrow P_1(\text{Succ}(y'_2))$ is true, where x'_2 is the first position to the right of x_2 in u with $P_{\langle /h \rangle}(x'_2)$ and y'_2 is the first position to the right of y_2 in v with $P_{\langle /h \rangle}(y'_2)$.

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

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).

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

Proof of Theorem 6.14.

By contradiction. Assume that the translation from FO($<$) to S-MonDatalog can be established by a function f with bound $f(m) \leq Tower(o(m))$. I.e., there is a $m_0 \in \mathbb{N}$ and a function $g \in \omega(1)$ such that $f(m) \leq Tower(\frac{m}{g(m)})$, for all $m \geq m_0$; and for every FO($<$)-sentence φ there is a S-MonDatalog-program \mathcal{P} of size $|\mathcal{P}| \leq f(|\varphi|) \leq Tower(\frac{|\varphi|}{g(|\varphi|)})$ and an $X \in \text{IDB}(\mathcal{P})$ such that (\mathcal{P}, X) defines the same Boolean query as φ on the class of all finite strings. According to Theorem 6.9, \mathcal{P} is equivalent to a MonDatalog-program \mathcal{P}' of size $|\mathcal{P}'| = 2^{\mathcal{O}(|\mathcal{P}|)} \leq Tower(\frac{|\varphi|}{g(|\varphi|)} + c)$, 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 (\mathcal{P}', X) and that has $|Q| \leq 2^{\mathcal{O}(|\mathcal{P}'|)} \leq Tower(\frac{|\varphi|}{g(|\varphi|)} + d)$ 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}' 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 FO($<$)-sentence φ can be translated into a nondeterministic finite automaton that defines the same Boolean query as φ on the class of all finite strings and that has at most $Tower(h(|\varphi|))$ 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. ■