
THE SUCCINCTNESS OF FIRST-ORDER LOGIC ON LINEAR ORDERS

MARTIN GROHE AND NICOLE SCHWEIKARDT

Institut für Informatik, Humboldt-Universität, Unter den Linden 6, D-10099 Berlin, Germany
e-mail address: {grohe|schweika}@informatik.hu-berlin.de

ABSTRACT. Succinctness is a natural measure for comparing the strength of different logics. Intuitively, a logic L_1 is more succinct than another logic L_2 if all properties that can be expressed in L_2 can be expressed in L_1 by formulas of (approximately) the same size, but some properties can be expressed in L_1 by (significantly) smaller formulas.

We study the succinctness of logics on linear orders. Our first theorem is concerned with the finite variable fragments of first-order logic. We prove that:

- (i) Up to a polynomial factor, the 2- and the 3-variable fragments of first-order logic on linear orders have the same succinctness.
- (ii) The 4-variable fragment is exponentially more succinct than the 3-variable fragment.

Our second main result compares the succinctness of first-order logic on linear orders with that of monadic second-order logic. We prove that the fragment of monadic second-order logic that has the same expressiveness as first-order logic on linear orders is non-elementarily more succinct than first-order logic.

1. INTRODUCTION

It is one of the fundamental themes of logic in computer science to study and compare the *strength* of various logics. Maybe the most natural measure of strength is the *expressive power* of a logic. By now, researchers from finite model theory, but also from more application driven areas such as database theory and automated verification, have developed a rich toolkit that has led to a good understanding of the expressive power of the fundamental logics (e.g. [3, 10, 12]). It should also be said that there are clear limits to the understanding of expressive power, which are often linked to open problems in complexity theory.

In several interesting situations, however, one encounters different logics of the same expressive power. As an example, let us consider node selecting query languages for XML-documents. Here the natural deductive query language monadic datalog [7] and various automata based query “languages” [13, 14, 6] have the same expressive power as monadic second-order logic. XML-documents are usually modelled by labelled trees. Logics on trees and strings also play an important role in automated verification. Of the logics studied in the context of verification, the modal μ -calculus is another logic that has the same expressive power as monadic second-order logic on ranked trees and strings, and linear time temporal logic LTL has the same expressive power as first-order logic on strings [11].

2000 ACM Subject Classification: F.4.1.

Key words and phrases: finite model theory, first-order logic, succinctness.

Succinctness is a natural measure for comparing the strength of logics that have the same expressive power. Intuitively, a logic L_1 is more succinct than another logic L_2 if all properties that can be expressed in L_2 can be expressed in L_1 by formulas of (approximately) the same size, but some properties can be expressed in L_1 by (significantly) smaller formulas.

For both expressiveness and succinctness there is a trade-off between the strength of a logic and the complexity of evaluating formulas of the logic. The difference lies in the way the complexity is measured. Expressiveness is related to *data complexity*, which only takes into account the size of the structure in which the formula has to be evaluated, whereas succinctness is related to the *combined complexity*, which takes into account both the size of the formula and the structure [17].

Succinctness has received surprisingly little attention so far. A few scattered results are [18, 2, 1, 4, 15]; for example, it is known that first-order logic on strings is non-elementarily more succinct than LTL [11, 15]. In [8], we started a more systematic investigation. Specifically, we studied the succinctness of various logics on trees that all have the same expressive power as monadic second-order logic. While we were able to gain a reasonable picture of the succinctness of these logics, it also became clear that we are far from a thorough understanding of succinctness. In particular, very few techniques for proving lower bounds are available.

Most of the lower bound proofs use automata theoretic arguments, often combined with a clever encoding of large natural numbers that goes back to Stockmeyer [15]. In [8], these techniques were also combined with complexity theoretic reductions to prove lower bounds on succinctness under certain complexity theoretic assumptions. Wilke [18] used refined automata theoretic arguments to prove that CTL^+ is exponentially more succinct than CTL. Adler and Immerman [1] were able to improve Wilke's lower bound slightly, but what is more important is that they introduced games for establishing lower bounds on succinctness. These games vaguely resemble Ehrenfeucht-Fraïssé games, which are probably the most important tools for establishing inexpressibility results.

In this paper, we study the succinctness of logics on linear orders (without any additional structure). In particular, we consider finite variable fragments of first-order logic. It is known and easy to see that even the 2-variable fragment has the same expressive power as full first-order logic on linear orders (with respect to Boolean and unary queries). We prove the following theorem:

Theorem 1.1. (i) Up to a polynomial factor, the 2 and the 3-variable fragments of first-order logic on linear orders have the same succinctness.
(ii) The 4-variable fragment of first-order logic on linear orders is exponentially more succinct than the 3-variable fragment. \square

For the sake of completeness, let us also mention that full first-order logic is at most exponentially more succinct than the 3-variable fragment. It remains an open problem if there is also an exponential gap in succinctness between full first-order logic and the 4-variable fragment.

Of course the main result here is the exponential gap in Theorem 1.1 (ii), but it should be noted that (i) is also by no means obvious. The theorem may seem very technical and not very impressive at first sight, but we believe that to gain a deeper understanding of the issue of succinctness it is of fundamental importance to master basic problems such as those we consider here first (similar, maybe, to basic inexpressibility results such as the inability of first-order logic to express that a linear order has even length). The main technical

result behind both parts of the theorem is that a 3-variable first-order formula stating that a linear order has length m must have size at least $\frac{1}{2}\sqrt{m}$. Our technique for proving this result originated in the Adler-Immerman games, even though later it turned out that the proofs are clearer if the reference to the game is dropped.

There is another reason the gap in succinctness between the 3- and 4-variable fragments is interesting: It is a long standing open problem in finite model theory if, for $k \geq 3$, the k variable fragment of first-order logic is strictly less expressive than the $(k+1)$ -variable fragment on the class of all ordered finite structures. This question is still open for all $k \geq 3$. Our result (ii) at least shows that there are properties that require exponentially larger 3-variable than 4-variable formulas.

Succinctness as a measure for comparing the strength of logics is not restricted to logics of the same expressive power. Even if a logic L_1 is more expressive than a logic L_2 , it is interesting to know whether those properties that can be expressed in both L_1 and L_2 can be expressed more succinctly in one of the logics. Sometimes, this may even be more important than the fact that some esoteric property is expressible in L_1 , but not L_2 . We compare first-order logic with the more expressive monadic second-order logic and prove:

Theorem 1.2. The fragment of monadic second-order logic that has the same expressiveness as first-order logic on linear orders is non-elementarily more succinct than first-order logic. \square

The paper is organised as follows: After the Preliminaries, in Section 3 we prove the main technical result behind Theorem 1.1. In Sections 4 and 5 we formally state and prove the two parts of the theorem. Finally, Section 6 is devoted to Theorem 1.2.

The present paper is the full version of the conference contribution [9].

2. PRELIMINARIES

We write \mathbb{N} for the set of non-negative integers.

We assume that the reader is familiar with first-order logic FO (cf., e.g., the textbooks [3, 10]). For a natural number k we write FO^k to denote the k -variable fragment of FO. The three variables available in FO^3 will always be denoted x , y , and z . We write $\text{FO}^3(<, succ, min, max)$ (resp., $\text{FO}^3(<)$) to denote the class of all FO^3 -formulas of signature $\{<, succ, min, max\}$ (resp., of signature $\{<\}$), with binary relation symbols $<$ and $succ$ and constant symbols min and max . In the present paper, such formulas will always be interpreted in finite structures where $<$ is a linear ordering, $succ$ the successor relation associated with $<$, and min and max the minimum and maximum elements w.r.t. $<$.

For every $N \in \mathbb{N}$ let \mathcal{A}_N be the $\{<, succ, min, max\}$ -structure with universe $\{0, \dots, N\}$, $<$ the natural linear ordering, $min^{\mathcal{A}_N} = 0$, $max^{\mathcal{A}_N} = N$, and $succ$ the relation with $(a, b) \in succ$ iff $a+1 = b$. We identify the class of linear orders with the set $\{\mathcal{A}_N : N \in \mathbb{N}\}$.

For a structure \mathcal{A} we write $\mathcal{U}^{\mathcal{A}}$ to denote \mathcal{A} 's universe. When considering FO^3 , an interpretation is a tuple (\mathcal{A}, α) , where \mathcal{A} is one of the structures \mathcal{A}_N (for some $N \in \mathbb{N}$) and $\alpha : \{x, y, z\} \rightarrow \mathcal{U}^{\mathcal{A}}$ is a variable assignment in \mathcal{A} . To simplify notation, we will extend every assignment α to a mapping $\alpha : \{x, y, z, min, max\} \rightarrow \mathcal{U}^{\mathcal{A}}$, letting $\alpha(min) = min^{\mathcal{A}}$ and $\alpha(max) = max^{\mathcal{A}}$. For a variable $v \in \{x, y, z\}$ and an element $a \in \mathcal{U}^{\mathcal{A}}$ we write $\alpha[\frac{a}{v}]$ to denote the assignment that maps v to a and that coincides with α on all other variables. If A is a set of interpretations and φ is an $\text{FO}^3(<, succ, min, max)$ -formula, we write $A \models \varphi$ to indicate that φ is satisfied by every interpretation in A .

In a natural way, we view formulas as finite trees (precisely, as their *syntax trees*), where leaves correspond to the atoms of the formulas, and inner vertices correspond to Boolean connectives or quantifiers. We define the *size* $\|\varphi\|$ of φ to be the number of vertices of φ 's syntax tree.

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} iff there is a function $f \in F$ such that for every L_1 -sentence φ_1 there is an L_2 -sentence φ_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 on the size of L_2 -formulas needed to express *all* of L_1 . This definition may seem slightly at odds with the common use of the term ‘‘succinctness’’ in statements such as ‘‘ L_1 is exponentially more succinct than L_2 ’’ meaning that there is *some* L_1 -formula that is not equivalent to any L_2 -formula of sub-exponential size. In our terminology we would rephrase this last statement as ‘‘ L_1 is *not* $2^{o(m)}$ -succinct in L_2 ’’ (here we interpret sub-exponential as $2^{o(m)}$, 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_1 is exponentially more succinct than L_2 ’’ in informal discussions.

Example 2.2. $\text{FO}^3(<, succ, min, max)$ is $\mathcal{O}(m)$ -succinct in $\text{FO}^3(<)$ on the class of linear orders, because $succ(x, y)$ (respectively, $x=min$, respectively, $x=max$) can be expressed by the formula

$$(x < y) \wedge \neg \exists z ((x < z) \wedge (z < y))$$

(respectively, $\neg \exists y (y < x)$, respectively, $\neg \exists y (x < y)$). \square

3. LOWER BOUND FOR FO^3

3.1. Lower Bound Theorem. Before stating our main lower bound theorem, we need some more notation.

If S is a set we write $\mathcal{P}_2(S)$ for the set of all 2-element subsets of S . For a finite subset S of \mathbb{N} we write $\text{MAX } S$ (respectively, $\text{MIN } S$) to denote the maximum (respectively, minimum) element in S . For integers m, n we define

$$\text{diff}(m, n) := m - n$$

to be the difference between m and n . We define $<-type(m, n) \in \{<, =, >\}$ as follows:

$$\begin{aligned} \text{if } m < n & \text{ then } <-type(m, n) := '<', \\ \text{if } m = n & \text{ then } <-type(m, n) := '=', \\ \text{if } m > n & \text{ then } <-type(m, n) := '>'. \end{aligned}$$

We next fix the notion of a *separator*. Basically, if A and B are sets of interpretations and δ is a separator for $\langle A, B \rangle$, then δ contains information that allows to distinguish every interpretation $\mathcal{I} \in A$ from every interpretation $\mathcal{J} \in B$.

Definition 3.1. [separator]

Let A and B be sets of interpretations.

A *potential separator* is a mapping

$$\delta : \mathcal{P}_2(\{min, max, x, y, z\}) \longrightarrow \mathbb{N}.$$

δ is called *separator for* $\langle A, B \rangle$, if the following is satisfied: For every $\mathcal{I} := (\mathcal{A}, \alpha) \in A$ and $\mathcal{J} := (\mathcal{B}, \beta) \in B$ there are $u, u' \in \{\min, \max, x, y, z\}$ with $u \neq u'$, such that $\delta(\{u, u'\}) \geq 1$ and

1. $\text{<-type}(\alpha(u), \alpha(u')) \neq \text{<-type}(\beta(u), \beta(u'))$ or
2. $\delta(\{u, u'\}) \geq \text{MIN} \{|\text{diff}(\alpha(u), \alpha(u'))|, |\text{diff}(\beta(u), \beta(u'))|\}$ and $\text{diff}(\alpha(u), \alpha(u')) \neq \text{diff}(\beta(u), \beta(u'))$. □

Note that δ is a separator for $\langle A, B \rangle$ if, and only if, δ is a separator for $\langle \{\mathcal{I}\}, \{\mathcal{J}\} \rangle$, for all $\mathcal{I} \in A$ and $\mathcal{J} \in B$. For simplicity, we will often write $\langle \mathcal{I}, \mathcal{J} \rangle$ instead of $\langle \{\mathcal{I}\}, \{\mathcal{J}\} \rangle$.

Let us now state an easy lemma on the existence of separators.

Lemma 3.2. If A and B are sets of interpretations for which there exists an $\text{FO}^3(<)$ -formula ψ such that $A \models \psi$ and $B \models \neg\psi$, then there exists a *separator* δ for $\langle A, B \rangle$. □

Proof: We need the following notation: For a number $d \in \mathbb{N}$ and an interpretation (\mathcal{A}, α) choose

$$\text{Ord}(\mathcal{A}, \alpha) : \{0, 1, 2, 3, 4\} \rightarrow \{\min, x, y, z, \max\}$$

such that $\alpha(\text{Ord}(\mathcal{A}, \alpha)(i)) \leq \alpha(\text{Ord}(\mathcal{A}, \alpha)(i+1))$, for all $0 \leq i < 4$. Furthermore, choose

$$\text{Dist}_d(\mathcal{A}, \alpha) : \{(i, i+1) : 0 \leq i < 4\} \rightarrow \{0, \dots, 2^{d+1}\}$$

such that the following is true for all $0 \leq i < 4$:

$$\text{Dist}_d(\mathcal{A}, \alpha)(i, i+1) = \text{diff}(\alpha(\text{Ord}(\mathcal{A}, \alpha)(i+1)), \alpha(\text{Ord}(\mathcal{A}, \alpha)(i))), \text{ or}$$

$$\text{Dist}_d(\mathcal{A}, \alpha)(i, i+1) = 2^{d+1} \leq \text{diff}(\alpha(\text{Ord}(\mathcal{A}, \alpha)(i+1)), \alpha(\text{Ord}(\mathcal{A}, \alpha)(i))).$$

Finally, we define the d -*type* of (\mathcal{A}, α) as

$$\text{Type}_d(\mathcal{A}, \alpha) := (\text{Ord}(\mathcal{A}, \alpha), \text{Dist}_d(\mathcal{A}, \alpha)).$$

Using an Ehrenfeucht-Fraïssé game, it is an easy exercise to show the following (cf., e.g., [3]):

Lemma 3.3. Let $d \in \mathbb{N}$ and let (\mathcal{A}, α) and (\mathcal{B}, β) be interpretations. If $\text{Type}_d(\mathcal{A}, \alpha) = \text{Type}_d(\mathcal{B}, \beta)$, then (\mathcal{A}, α) and (\mathcal{B}, β) cannot be distinguished by $\text{FO}(<)$ -formulas of quantifier depth $\leq d$. □

Let d be the quantifier depth of the formula ψ . We define δ to be the potential separator with $\delta(p) := 2^{d+1}$, for all $p \in \mathcal{P}_2(\{\min, \max, x, y, z\})$.

To show that δ is, in fact, a *separator* for $\langle A, B \rangle$, let $(\mathcal{A}, \alpha) \in A$ and $(\mathcal{B}, \beta) \in B$. Since $(\mathcal{A}, \alpha) \models \psi$ and $(\mathcal{B}, \beta) \not\models \psi$, we obtain from Lemma 3.3 that $\text{Type}_d((\mathcal{A}, \alpha)) \neq \text{Type}_d((\mathcal{B}, \beta))$, i.e.,

1. $\text{Ord}(\mathcal{A}, \alpha) \neq \text{Ord}(\mathcal{B}, \beta)$, or
2. $\text{Dist}_d(\mathcal{A}, \alpha) \neq \text{Dist}_d(\mathcal{B}, \beta)$.

Therefore, there exist $u, u' \in \{\min, \max, x, y, z\}$ with $u \neq u'$, such that

1. $\text{<-type}(\alpha(u), \alpha(u')) \neq \text{<-type}(\beta(u), \beta(u'))$, or
2. $\text{diff}(\alpha(u), \alpha(u')) \neq \text{diff}(\beta(u), \beta(u'))$ and $\delta(\{u, u'\}) = 2^{d+1} \geq \text{MIN} \{|\text{diff}(\alpha(u), \alpha(u'))|, |\text{diff}(\beta(u), \beta(u'))|\}$.

Consequently, δ is a separator for $\langle A, B \rangle$, and the proof of Lemma 3.2 is complete. ■

Definition 3.4. [weight of δ]

Let δ be a potential separator. We define

(a) the *border-distance*

$$b(\delta) := \text{MAX} \{ \delta(\{\min, \max\}), \delta(\{\min, u\}) + \delta(\{u', \max\}) : u, u' \in \{x, y, z\} \}$$

(b) the *centre-distance*

$$c(\delta) := \text{MAX} \{ \delta(p) + \delta(q) : p, q \in \mathcal{P}_2(\{x, y, z\}), p \neq q \}$$

(c) the *weight*

$$w(\delta) := \sqrt{c(\delta)^2 + b(\delta)}. \quad \square$$

There is not much intuition we can give for this particular choice of weight function, except for the fact that it seems to be exactly what is needed for the proof of our main lower bound theorem (Theorem 3.6). At least it will later, in Remark 3.13, become clear why the $\sqrt{}$ -function is used for defining the weight function.

Definition 3.5. [minimal separator]

δ is called a *minimal separator* for $\langle A, B \rangle$ if δ is a separator for $\langle A, B \rangle$ and

$$w(\delta) = \text{MIN} \{ w(\delta') : \delta' \text{ is a separator for } \langle A, B \rangle \}. \quad \square$$

Now we are ready to formally state our main lower bound theorem on the size of $\text{FO}^3(<)$ -formulas:

Theorem 3.6. [main lower bound theorem]

If ψ is an $\text{FO}^3(<, \text{succ}, \min, \max)$ -formula, A and B are sets of interpretations such that $A \models \psi$ and $B \models \neg\psi$, and δ is a minimal separator for $\langle A, B \rangle$, then

$$\|\psi\| \geq \frac{1}{2} \cdot w(\delta). \quad \square$$

Before giving details on the proof of Theorem 3.6, let us first point out its following easy consequence:

Corollary 3.7.

Let $n > m \geq 0$. The two linear orders \mathcal{A}_m and \mathcal{A}_n (with universe $\{0, \dots, m\}$ and $\{0, \dots, n\}$, respectively) cannot be distinguished by an $\text{FO}^3(<, \text{succ}, \min, \max)$ -sentence of size $< \frac{1}{2}\sqrt{m}$. \square

Proof: Let ψ be an $\text{FO}^3(<, \text{succ}, \min, \max)$ -sentence with $\mathcal{A}_m \models \psi$ and $\mathcal{A}_n \models \neg\psi$. Let α be the assignment that maps each of the variables x, y , and z to the value 0. Consider the mapping

$$\begin{aligned} \delta_m &: \mathcal{P}_2(\{\min, \max, x, y, z\}) \rightarrow \mathbb{N} \quad \text{defined via} \\ \delta_m(p) &:= \begin{cases} m & , \text{ if } p = \{\min, \max\} \\ 0 & , \text{ otherwise.} \end{cases} \end{aligned}$$

It is straightforward to check that $w(\delta_m) = \sqrt{m}$ and that δ_m is a *minimal separator* for $\langle (\mathcal{A}_m, \alpha), (\mathcal{A}_n, \alpha) \rangle$.

From Theorem 3.6 we therefore obtain that

$$\|\psi\| \geq \frac{1}{2} \cdot w(\delta_m) = \frac{1}{2} \cdot \sqrt{m}.$$

This completes the proof of Corollary 3.7. \blacksquare

To prove Theorem 3.6 we need a series of intermediate results, as well as the notion of an *extended syntax tree* of a formula, which is a syntax tree where each node carries an additional label containing information about sets of interpretations satisfying, respectively, not satisfying, the associated subformula. More precisely, every node v of the extended syntax tree carries an *interpretation label* $il(v)$ which consists of a pair $\langle A, B \rangle$ of sets of interpretations such that every interpretation in A , but no interpretation in B , satisfies the subformula represented by the subtree rooted at node v . Basically, such an extended syntax tree corresponds to a game tree that is constructed by the two players of the *Adler-Immerman game* (cf., [1]).

For proving Theorem 3.6 we consider an extended syntax tree \mathcal{T} of the given formula ψ . We define a weight function on the nodes of \mathcal{T} by defining the weight $w(v)$ of each node v of \mathcal{T} to be the weight of a *minimal separator* for $il(v)$. Afterwards — and this is the main technical difficulty — we show that the weight of each node v is bounded (from above) by the weights of v 's children. This, in turn, enables us to prove a lower bound on the number of nodes in \mathcal{T} which depends on the weight of the root node.

3.2. Proof of Theorem 3.6. We start with the formal definition of extended syntax trees.

Definition 3.8. [extended syntax tree]

Let ψ be an $\text{FO}^3(<, succ, min, max)$ -formula, let A and B be sets of interpretations such that $A \models \psi$ and $B \models \neg\psi$. By induction on the construction of ψ we define an *extended syntax tree* $\mathcal{T}_\psi^{\langle A, B \rangle}$ as follows:

- If ψ is an atomic formula, then $\mathcal{T}_\psi^{\langle A, B \rangle}$ consists of a single node v that has a *syntax label* $sl(v) := \psi$ and an *interpretation label* $il(v) := \langle A, B \rangle$.
- If ψ is of the form $\neg\psi_1$, then $\mathcal{T}_\psi^{\langle A, B \rangle}$ has a root node v with $sl(v) := \neg$ and $il(v) := \langle A, B \rangle$.

The unique child of v is the root of $\mathcal{T}_{\psi_1}^{\langle B, A \rangle}$. Note that $B \models \psi_1$ and $A \models \neg\psi_1$.

- If ψ is of the form $\psi_1 \vee \psi_2$, then $\mathcal{T}_\psi^{\langle A, B \rangle}$ has a root node v with $sl(v) := \vee$ and $il(v) := \langle A, B \rangle$.

The first child of v is the root of $\mathcal{T}_{\psi_1}^{\langle A_1, B \rangle}$. The second child of v is the root of $\mathcal{T}_{\psi_2}^{\langle A_2, B \rangle}$, where, for $i \in \{1, 2\}$, $A_i = \{(\mathcal{A}, \alpha) \in A : (\mathcal{A}, \alpha) \models \psi_i\}$.

Note that $A = A_1 \cup A_2$, $A_i \models \psi_i$, and $B \models \neg\psi_i$.

- If ψ is of the form $\psi_1 \wedge \psi_2$, then $\mathcal{T}_\psi^{\langle A, B \rangle}$ has a root node v with $sl(v) := \wedge$ and $il(v) := \langle A, B \rangle$.

The first child of v is the root of $\mathcal{T}_{\psi_1}^{\langle A, B_1 \rangle}$. The second child of v is the root of $\mathcal{T}_{\psi_2}^{\langle A, B_2 \rangle}$, where, for $i \in \{1, 2\}$, $B_i = \{(\mathcal{B}, \beta) \in B : (\mathcal{B}, \beta) \not\models \psi_i\}$.

Note that $B = B_1 \cup B_2$, $A \models \psi_i$, and $B_i \models \neg\psi_i$.

- If ψ is of the form $\exists u \psi_1$, for a variable $u \in \{x, y, z\}$, then $\mathcal{T}_\psi^{\langle A, B \rangle}$ has a root node v with $sl(v) := \exists u$ and $il(v) := \langle A, B \rangle$. The unique child of v is the root of $\mathcal{T}_{\psi_1}^{\langle A_1, B_1 \rangle}$, where $B_1 := \{(\mathcal{B}, \beta[\frac{b}{u}]) : (\mathcal{B}, \beta) \in B, b \in \mathcal{U}^B\}$, and A_1 is chosen as follows: For every $(\mathcal{A}, \alpha) \in A$ fix an element $a \in \mathcal{U}^A$ such that $(\mathcal{A}, \alpha[\frac{a}{u}]) \models \psi_1$, and let $A_1 := \{(\mathcal{A}, \alpha[\frac{a}{u}]) : (\mathcal{A}, \alpha) \in A\}$. Note that $A_1 \models \psi_1$ and $B_1 \models \neg\psi_1$.

- If ψ is of the form $\forall u \psi_1$, for a variable $u \in \{x, y, z\}$, then $\mathcal{T}_\psi^{\langle A, B \rangle}$ has a root node v with $sl(v) := \forall u$ and $il(v) := \langle A, B \rangle$. The unique child of v is the root of $\mathcal{T}_{\psi_1}^{\langle A_1, B_1 \rangle}$, where $A_1 := \{(\mathcal{A}, \alpha[\frac{a}{u}]) : (\mathcal{A}, \alpha) \in A, a \in \mathcal{U}^A\}$, and B_1 is chosen as follows: For every $(\mathcal{B}, \beta) \in B$

fix an element $b \in \mathcal{U}^{\mathcal{B}}$ such that $(\mathcal{B}, \beta[\frac{b}{u}]) \models \neg\psi_1$, and let $B_1 := \{(\mathcal{B}, \alpha[\frac{b}{u}]) : (\mathcal{B}, \beta) \in B\}$. Note that $A_1 \models \psi_1$ and $B_1 \models \neg\psi_1$. \square

The following is the main technical result necessary for our proof of Theorem 3.6.

Lemma 3.9. Let ψ be an $\text{FO}^3(<, \text{succ}, \text{min}, \text{max})$ -formula, let A and B be sets of interpretations such that $A \models \psi$ and $B \models \neg\psi$, and let \mathcal{T} be an extended syntax tree $\mathcal{T}_\psi^{(A,B)}$.

For every node v of \mathcal{T} the following is true, where δ is a minimal separator for $il(v)$:

- (a) If v is a leaf, then $w(\delta) \leq 1$.
- (b) If v has 2 children v_1 and v_2 , and δ_i is a minimal separator for $il(v_i)$, for $i \in \{1, 2\}$, then $w(\delta) \leq w(\delta_1) + w(\delta_2)$.
- (c) If v has exactly one child v_1 , and δ_1 is a minimal separator for $il(v_1)$, then $w(\delta) \leq w(\delta_1) + 2$. \square

The proof of Lemma 3.9 is given in Section 3.3 below.

For a binary tree \mathcal{T} we write $\|\mathcal{T}\|$ to denote the number of nodes of \mathcal{T} . For the proof of Theorem 3.6 we also need the following easy observation.

Lemma 3.10. Let \mathcal{T} be a finite binary tree where each node v is equipped with a *weight* $w(v) \geq 0$ such that the following is true:

- (a) If v is a leaf, then $w(v) \leq 1$.
- (b) If v has 2 children v_1 and v_2 , then $w(v) \leq w(v_1) + w(v_2)$.
- (c) If v has exactly one child v_1 , then $w(v) \leq w(v_1) + 2$.

Then, $\|\mathcal{T}\| \geq \frac{1}{2} \cdot w(r)$, where r is the root of \mathcal{T} . \square

Proof: By induction on the size of \mathcal{T} .

If \mathcal{T} consists of a single node v , then $\|\mathcal{T}\| = 1 \geq \frac{1}{2} \cdot 1$; and $1 \geq w(v)$, since v is a leaf.

If \mathcal{T} consists of a root node v whose first child v_1 is the root of a tree \mathcal{T}_1 and whose second child v_2 is the root of a tree \mathcal{T}_2 , then $\|\mathcal{T}\| = 1 + \|\mathcal{T}_1\| + \|\mathcal{T}_2\|$. By induction we know for $i \in \{1, 2\}$ that $\|\mathcal{T}_i\| \geq \frac{1}{2}w(v_i)$. From the assumption we have that $w(v) \leq w(v_1) + w(v_2)$. Therefore,

$$\|\mathcal{T}\| \geq 1 + \frac{1}{2}w(v_1) + \frac{1}{2}w(v_2) \geq \frac{1}{2}w(v).$$

If \mathcal{T} consists of a root node v whose unique child v_1 is the root of a tree \mathcal{T}_1 , then $\|\mathcal{T}\| = 1 + \|\mathcal{T}_1\|$. By induction we know that $\|\mathcal{T}_1\| \geq \frac{1}{2}w(v_1)$. From the assumption we have that $w(v) \leq w(v_1) + 2$, i.e., $\frac{1}{2}w(v) \leq \frac{1}{2}w(v_1) + 1$. Therefore, $\|\mathcal{T}\| \geq 1 + \frac{1}{2}w(v_1) \geq \frac{1}{2}w(v)$.

This completes the proof of Lemma 3.10. \blacksquare

Using Lemma 3.9 and 3.10, we are ready for the

Proof of Theorem 3.6:

We are given an $\text{FO}^3(<, \text{succ}, \text{min}, \text{max})$ -formula ψ and sets A and B of interpretations such that $A \models \psi$ and $B \models \neg\psi$. Let \mathcal{T} be an extended syntax tree $\mathcal{T}_\psi^{(A,B)}$.

We equip each node v of \mathcal{T} with a *weight* $w(v) := w(\delta_v)$, where δ_v is a minimal separator for $il(v)$. From Lemma 3.9 we obtain that the preconditions of Lemma 3.10 are satisfied. Therefore, $\|\mathcal{T}\| \geq \frac{1}{2} \cdot w(r)$, where r is the root of \mathcal{T} , i.e., $w(r) = w(\delta)$, for a minimal separator δ for $il(r) = \langle A, B \rangle$.

From Definition 3.8 it should be obvious that $\|\psi\| = \|\mathcal{T}\|$. Therefore, the proof of Theorem 3.6 is complete. \blacksquare

3.3. Proof of Lemma 3.9. We partition the proof of Lemma 3.9 into proofs for the parts (a), (b), and (c), where part (c) turns out to be the most elaborate.

According to the assumptions of Lemma 3.9 we are given an $\text{FO}^3(<, succ, min, max)$ -formula ψ and sets A and B of interpretations such that $A \models \psi$ and $B \models \neg\psi$. Furthermore, we are given an extended syntax tree $\mathcal{T} = \mathcal{T}_\psi^{(A,B)}$. Throughout the remainder of this section, \mathcal{T} will always denote this particular syntax tree.

Proof of part (a) of Lemma 3.9:

Let v be a *leaf* of \mathcal{T} and let δ be a minimal separator for $\langle A_v, B_v \rangle := il(v)$. Our aim is to show that $w(\delta) \leq 1$.

By Definition 3.8 we know that $sl(v)$ is an *atomic* formula of the form $R(u, u')$ for $R \in \{<, =, succ\}$ and $u, u' \in \{min, max, x, y, z\}$. Furthermore, $A_v \models R(u, u')$ and $B_v \models \neg R(u, u')$. I.e., for all $(\mathcal{A}, \alpha) \in A_v$ and $(\mathcal{B}, \beta) \in B_v$,

$$\begin{aligned} <-type(\alpha(u), \alpha(u')) \neq <-type(\beta(u), \beta(u')) \\ \text{or} \end{aligned}$$

$$|diff(\alpha(u), \alpha(u'))| = 1 \neq |diff(\beta(u), \beta(u'))|.$$

In case that $u \neq u'$ we can define a *separator* $\tilde{\delta}$ for $\langle A_v, B_v \rangle$ via

$$\tilde{\delta}(p) := \begin{cases} 1 & , \text{ if } p = \{u, u'\} \\ 0 & , \text{ otherwise.} \end{cases}$$

Since δ is a *minimal* separator, we obtain that $w(\delta) \leq w(\tilde{\delta}) = 1$.

It remains to consider the case where $u = u'$. Here, $A_v \models R(u, u)$ and $B_v \models \neg R(u, u)$. Since $R \in \{<, =, succ\}$ this implies that $A_v = \emptyset$ or $B_v = \emptyset$. Therefore, according to Definition 3.1, the mapping $\tilde{\delta}$ with $\tilde{\delta}(p) = 0$, for all $p \in \mathcal{P}_2(\{min, max, x, y, z\})$, is a separator for $\langle A_v, B_v \rangle$. Hence, $w(\delta) \leq w(\tilde{\delta}) = 0$.

This completes the proof of part (a) of Lemma 3.9. ■

The essential step in the proof of part (b) of Lemma 3.9 is the following easy lemma.

Lemma 3.11. Let v be a node of \mathcal{T} that has two children v_1 and v_2 . Let δ_1 and δ_2 be separators for $il(v_1)$ and $il(v_2)$, respectively. Let $\tilde{\delta}$ be the potential separator defined on every $p \in \mathcal{P}_2(\{min, max, x, y, z\})$ via

$$\tilde{\delta}(p) := \delta_1(p) + \delta_2(p).$$

Then, $\tilde{\delta}$ is a *separator* for $il(v)$. □

Proof: Let $\langle A, B \rangle := il(v)$. We need to show that $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$, for all $\mathcal{I} \in A$ and $\mathcal{J} \in B$. Let therefore $\mathcal{I} := (\mathcal{A}, \alpha) \in A$ and $\mathcal{J} := (\mathcal{B}, \beta) \in B$ be fixed for the remainder of this proof.

Since v has 2 children, we know from Definition 3.8 that $sl(v) = \vee$ or $sl(v) = \wedge$. Let us first consider the case where $sl(v) = \vee$.

From Definition 3.8 we know that, for $i \in \{1, 2\}$, $il(v_i) = \langle A_i, B \rangle$, where $A_1 \cup A_2 = A$. Therefore, there is an $i \in \{1, 2\}$ such that $\mathcal{I} \in A_i$. From the assumption we know that δ_i is a separator for $\langle A_i, B \rangle$. Therefore, there are $u, u' \in \{min, max, x, y, z\}$ with $u \neq u'$, such that $\delta_i(\{u, u'\}) \geq 1$ and

$$1. \ <-type(\alpha(u), \alpha(u')) \neq <-type(\beta(u), \beta(u')) \quad \text{or}$$

2. $\delta(\{u, u'\}) \geq \text{MIN} \{ |\text{diff}(\alpha(u), \alpha(u'))|, |\text{diff}(\beta(u), \beta(u'))| \}$ and
 $\text{diff}(\alpha(u), \alpha(u')) \neq \text{diff}(\beta(u), \beta(u'))$.

Since $\tilde{\delta}(\{u, u'\}) = \delta_1(\{u, u'\}) + \delta_2(\{u, u'\})$, we know that $\tilde{\delta}(\{u, u'\}) \geq \delta_i(\{u, u'\})$. Therefore, $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$. This completes the proof of Lemma 3.11 for the case that $sl(v) = \vee$.

The case $sl(v) = \wedge$ follows by symmetry. \blacksquare

Using Lemma 3.11, the proof of part (b) of Lemma 3.9 is straightforward:

Proof of part (b) of Lemma 3.9:

Let v be a node of \mathcal{T} that has two children v_1 and v_2 . Let δ , δ_1 , and δ_2 , respectively, be minimal separators for $il(v)$, $il(v_1)$, and $il(v_2)$, respectively. Our aim is to show that $w(\delta) \leq w(\delta_1) + w(\delta_2)$.

Let $\tilde{\delta}$ be the separator for $il(v)$ obtained from Lemma 3.11. Since δ is a *minimal* separator for $il(v)$, it suffices to show that $w(\tilde{\delta}) \leq w(\delta_1) + w(\delta_2)$.

Using Definition 3.4, it is straightforward to check that $b(\tilde{\delta}) \leq b(\delta_1) + b(\delta_2)$ and $c(\tilde{\delta}) \leq c(\delta_1) + c(\delta_2)$. From this we obtain that

$$\begin{aligned} w(\tilde{\delta})^2 &= c(\tilde{\delta})^2 + b(\tilde{\delta}) \\ &\leq (c(\delta_1) + c(\delta_2))^2 + b(\delta_1) + b(\delta_2) \\ &= c(\delta_1)^2 + b(\delta_1) + c(\delta_2)^2 + b(\delta_2) + 2c(\delta_1)c(\delta_2) \\ &\leq w(\delta_1)^2 + w(\delta_2)^2 + 2w(\delta_1)w(\delta_2) \\ &= (w(\delta_1) + w(\delta_2))^2. \end{aligned}$$

I.e., we have shown that $w(\tilde{\delta}) \leq w(\delta_1) + w(\delta_2)$.

This completes the proof of part (b) of Lemma 3.9. \blacksquare

An essential step in the proof of part (c) of Lemma 3.9 is the following lemma.

Lemma 3.12. Let v be a node of \mathcal{T} that has syntax-label $sl(v) = \mathbf{Q}u$, for $\mathbf{Q} \in \{\exists, \forall\}$ and $u \in \{x, y, z\}$. Let δ_1 be a separator for $il(v_1)$, where v_1 is the unique child of v in \mathcal{T} . Let $\tilde{\delta}$ be the potential separator defined via

- $\tilde{\delta}(\{u, u'\}) := 0$, for all $u' \in \{\min, \max, x, y, z\} \setminus \{u\}$,
- $\tilde{\delta}(\{\min, \max\}) := \text{MAX} \{ \delta_1(\{\min, \max\}), \delta_1(\{\min, u\}) + \delta_1(\{u, \max\}) + 1 \}$,

and for all u', u'' such that $\{x, y, z\} = \{u, u', u''\}$ and all $m \in \{\min, \max\}$,

- $\tilde{\delta}(\{u', u''\}) := \text{MAX} \{ \delta_1(\{u', u''\}), \delta_1(\{u', u\}) + \delta_1(\{u, u''\}) + 1 \}$,
- $\tilde{\delta}(\{m, u'\}) := \text{MAX} \{ \delta_1(\{m, u'\}), \delta_1(\{m, u\}) + \delta_1(\{u, u'\}) + 1 \}$.

Then, $\tilde{\delta}$ is a separator for $il(v)$. \square

Proof: We only consider the case where $\mathbf{Q}u = \exists z$. All other cases $\mathbf{Q} \in \{\exists, \forall\}$ and $u \in \{x, y, z\}$ follow by symmetry.

Let $\langle A, B \rangle := il(v)$. We need to show that $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$, for all $\mathcal{I} \in A$ and $\mathcal{J} \in B$. Let therefore $\mathcal{I} := (\mathcal{A}, \alpha) \in A$ and $\mathcal{J} := (\mathcal{B}, \beta) \in B$ be fixed for the remainder of this proof. The aim is to show that $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.

Since $sl(v) = \exists z$, we know from Definition 3.8 that $il(v_1) = \langle A_1, B_1 \rangle$, where B_1 contains the interpretations $(\mathcal{B}, \beta[\frac{b}{z}])$, for all $b \in \mathcal{U}^{\mathcal{B}}$, and A_1 contains an interpretation $(\mathcal{A}, \alpha[\frac{a}{z}])$, for




Figure 1: Situation in Claim 3.

a particular $a \in \mathcal{U}^{\mathcal{A}}$. We define $\alpha_a := \alpha[\frac{a}{z}]$, $\mathcal{I}_a := (\mathcal{A}, \alpha_a)$, and for every $b \in \mathcal{U}^{\mathcal{B}}$, $\beta_b := \beta[\frac{b}{z}]$ and $\mathcal{J}_b := (\mathcal{B}, \beta_b)$.

From the fact that δ_1 is a separator for $sl(v_1)$, we in particular know, for every $b \in \mathcal{U}^{\mathcal{B}}$, that δ_1 is a separator for $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$. I.e., we know the following:

For every $b \in \mathcal{U}^{\mathcal{B}}$ there are $u_b, u'_b \in \{\min, \max, x, y, z\}$ with $u_b \neq u'_b$, such that $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is separated by $\delta_1(\{u_b, u'_b\})$, i.e., $\delta_1(\{u_b, u'_b\}) \geq 1$ and

- (1)_b: $<-type(\alpha_a(u_b), \alpha_a(u'_b)) \neq <-type(\beta_b(u_b), \beta_b(u'_b))$, or
(2)_b: $\delta_1(\{u_b, u'_b\}) \geq \text{MIN} \{ |diff(\alpha_a(u_b), \alpha_a(u'_b))|, |diff(\beta_b(u_b), \beta_b(u'_b))| \}$ and
 $diff(\alpha_a(u_b), \alpha_a(u'_b)) \neq diff(\beta_b(u_b), \beta_b(u'_b))$.

In what follows we will prove a series of claims which ensure that $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$. We start with

Claim 1. If there is a $b \in \mathcal{U}^{\mathcal{B}}$ such that $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is separated by $\delta_1(\{u_b, u'_b\})$ with $z \notin \{u_b, u'_b\}$, then $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$. \square

Proof: As $z \notin \{u_b, u'_b\}$, we have, by definition of $\tilde{\delta}$, that $\tilde{\delta}(\{u_b, u'_b\}) \geq \delta_1(\{u_b, u'_b\})$. Therefore, (1)_b and (2)_b imply that $\tilde{\delta}$ is a separator for $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ as well as for $\langle \mathcal{I}, \mathcal{J} \rangle$.

This completes the proof of Claim 1. \blacksquare

Due to Claim 1 it henceforth suffices to assume that for no $b \in \mathcal{U}^{\mathcal{B}}$, $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is separated by $\delta_1(\{u_b, u'_b\})$ with $z \notin \{u_b, u'_b\}$. I.e., we assume that for every $b \in \mathcal{U}^{\mathcal{B}}$, $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is separated by $\delta_1(\{\min, z\})$, $\delta_1(\{z, \max\})$, $\delta_1(\{x, z\})$, or $\delta_1(\{y, z\})$.

Claim 2. If $a = \alpha(u)$ for some $u \in \{\min, \max, x, y\}$, then $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$. \square

Proof: Choose $b := \beta(u)$. Therefore, $\alpha_a(z) = a = \alpha(u) = \alpha_a(u)$ and $\beta_b(z) = b = \beta(u) = \beta_b(u)$.

We know that $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is separated by $\delta_1(\{z, u'\})$, for some $u' \in \{\min, \max, x, y\}$. Furthermore, since $\alpha_a(z) = \alpha_a(u)$ and $\beta_b(z) = \beta_b(u)$, we have $u' \neq u$.

By definition of $\tilde{\delta}$ we know that $\tilde{\delta}(\{u, u'\}) \geq \delta_1(\{z, u'\})$. Therefore, $\tilde{\delta}$ is a separator for $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ as well as for $\langle \mathcal{I}, \mathcal{J} \rangle$.

This completes the proof of Claim 2. \blacksquare

Due to Claim 2 it henceforth suffices to assume that, $a \neq \alpha(u)$, for all $u \in \{\min, \max, x, y\}$.

Claim 3. If $\delta_1(\{\min, z\}) \geq diff(a, \min^{\mathcal{A}})$, then $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$. \square

Proof: We distinguish between two cases. An illustration is given in Figure 1.




Figure 2: Situation in Claim 5 (for the special case that $\alpha(x) \leq a$).

Case 1: $\text{diff}(\max^{\mathcal{B}}, \min^{\mathcal{B}}) \geq \text{diff}(a, \min^{\mathcal{A}})$.

In this case we can choose $b \in \mathcal{U}^{\mathcal{B}}$ with

$$\text{diff}(b, \min^{\mathcal{B}}) = \text{diff}(a, \min^{\mathcal{A}})$$

(simply via $b := a$). Obviously, $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is not separated by $\delta_1(\{\min, z\})$. However, we know that $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is separated by $\delta_1(\{z, u'\})$, for some $u' \in \{x, y, \max\}$.

Since

$$\text{diff}(a, \min^{\mathcal{A}}) + \delta_1(\{z, u'\}) \leq \delta_1(\{\min, z\}) + \delta_1(\{z, u'\}) \leq \tilde{\delta}(\{\min, u'\}),$$

it is straightforward to see that $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$, and also $\langle \mathcal{I}, \mathcal{J} \rangle$, is separated by $\tilde{\delta}(\{\min, u'\})$. I.e., $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.

Case 2: $\text{diff}(\max^{\mathcal{B}}, \min^{\mathcal{B}}) < \text{diff}(a, \min^{\mathcal{A}})$.

Since

$$\tilde{\delta}(\{\min, \max\}) \geq \delta_1(\{\min, z\}) \geq \text{diff}(a, \min^{\mathcal{A}}) > \text{diff}(\max^{\mathcal{B}}, \min^{\mathcal{B}}),$$

we know that $\langle \mathcal{I}, \mathcal{J} \rangle$ is separated by $\tilde{\delta}(\{\min, \max\})$. I.e., $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.

This completes the proof of Claim 3. ■

By symmetry we also obtain the following

Claim 4. If $\delta_1(\{z, \max\}) \geq \text{diff}(\max^{\mathcal{A}}, a)$, then $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$. □

In a similar way, we can also show the following

Claim 5. If $\delta_1(\{x, z\}) \geq |\text{diff}(\alpha(x), a)|$, then $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$. □

Proof: We distinguish between three cases. An illustration is given in Figure 2.

Case 1: There is a $b \in \mathcal{U}^{\mathcal{B}}$ such that $\text{diff}(\beta(x), b) = \text{diff}(\alpha(x), a)$.

Obviously, $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is not separated by $\delta_1(\{x, z\})$. However, we know that $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is separated by $\delta_1(\{z, u'\})$, for some $u' \in \{\min, y, \max\}$.

Since

$$|\text{diff}(\alpha(x), a)| + \delta_1(\{z, u'\}) \leq \delta_1(\{x, z\}) + \delta_1(\{z, u'\}) \leq \tilde{\delta}(\{x, u'\}),$$

it is straightforward to see that $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$, and also $\langle \mathcal{I}, \mathcal{J} \rangle$, is separated by $\tilde{\delta}(\{x, u'\})$. I.e., $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.




Figure 3: Situation at the beginning of Claim 7.

Case 2: $\alpha(x) \leq a$ and $\text{diff}(\max^B, \beta(x)) < \text{diff}(a, \alpha(x))$.

Since

$$\tilde{\delta}(\{x, \max\}) \geq \delta_1(\{x, z\}) \geq \text{diff}(a, \alpha(x)) > \text{diff}(\max^B, \beta(x)),$$

we know that $\langle \mathcal{I}, \mathcal{J} \rangle$ is separated by $\tilde{\delta}(\{x, \max\})$. I.e., $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.

Case 3: $a < \alpha(x)$ and $\text{diff}(\beta(x), \min^B) < \text{diff}(\alpha(x), a)$.

This case is analogous to Case 2.

Now the proof of Claim 5 is complete, because one of the three cases above must apply. ■

By symmetry we also obtain the following

Claim 6. If $\delta_1(\{y, z\}) \geq |\text{diff}(\alpha(y), a)|$, then $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$. □

Finally, we show the following

Claim 7. If none of the assumptions of the Claims 1–6 is satisfied, then $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$. □

Proof: We assume w.l.o.g. that $\min^A \leq \alpha(x) \leq \alpha(y) \leq \max^A$.

Since Claim 1 does not apply, we know that also $\min^B \leq \beta(x) \leq \beta(y) \leq \max^B$.

Since Claims 2–6 do not apply, we furthermore know that

1. $|\text{diff}(\alpha(u'), a)| > \delta_1(\{u', z\})$, for all $u' \in \{\min, \max, x, y\}$, and
2. $\min^A < a < \alpha(x)$ or $\alpha(x) < a < \alpha(y)$ or $\alpha(y) < a < \max^A$.

We distinguish between different cases, depending on the particular interval that a belongs to. An illustration is given in Figure 3.

Case 1: $\min^A < a < \alpha(x)$.

Case 1.1: $\text{diff}(\beta(x), \min^B) \leq \delta_1(\{\min, z\})$.

By definition of $\tilde{\delta}$ we have $\text{diff}(\beta(x), \min^B) \leq \tilde{\delta}(\{\min, x\})$. Since

$$\text{diff}(\alpha(x), \min^A) > \text{diff}(a, \min^A) > \delta_1(\{\min, z\}) \geq \text{diff}(\beta(x), \min^B),$$

we therefore know that $\langle \mathcal{I}, \mathcal{J} \rangle$ is separated by $\tilde{\delta}(\{\min, z\})$. I.e., $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.

Case 1.2: $\text{diff}(\beta(x), \min^B) > \delta_1(\{\min, z\})$.

In this case we can choose $b \leq \beta(x)$ such that $\text{diff}(b, \min^B) = \delta_1(\{\min, z\}) + 1$. An illustration is given in Figure 4. Then, $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is not separated by $\delta_1(\{\min, z\})$. However, we know that $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is separated by $\delta_1(\{z, u'\})$, for some $u' \in \{x, y, \max\}$. From the assumptions of




Figure 4: Situation in Case 1.2 of Claim 7.




Figure 5: Situation in Case 3.2 of Claim 7.

Claim 7 we know that $\text{diff}(\alpha(u'), a) > \delta_1(\{z, u'\})$. Hence we must have that $\text{diff}(\beta(u'), b) \leq \delta_1(\{z, u'\})$. Therefore,

$$\begin{aligned} \text{diff}(\beta(u'), \min^{\mathcal{B}}) &= \text{diff}(\beta(u'), b) + \text{diff}(b, \min^{\mathcal{B}}) \\ &\leq \delta_1(\{z, u'\}) + \delta_1(\{\min, z\}) + 1 \\ &\leq \tilde{\delta}(\{\min, u'\}). \end{aligned}$$

Since

$$\text{diff}(\alpha(u'), \min^{\mathcal{A}}) = \text{diff}(\alpha(u'), a) + \text{diff}(a, \min^{\mathcal{A}}) > \text{diff}(\beta(u'), \min^{\mathcal{B}}),$$

we hence obtain that $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.

Case 2: $\alpha(y) < a < \max^{\mathcal{A}}$.

This case is analogous to Case 1.

Case 3: $\alpha(x) < a < \alpha(y)$.

Case 3.1: $\text{diff}(\beta(y), \beta(x)) \leq \delta_1(\{x, z\})$.

By definition of $\tilde{\delta}$ we have $\text{diff}(\beta(y), \beta(x)) \leq \tilde{\delta}(\{x, y\})$. Since

$$\text{diff}(\alpha(y), \alpha(x)) > \text{diff}(a, \alpha(x)) > \delta_1(\{x, z\}) \geq \text{diff}(\beta(y), \beta(x)),$$

we therefore know that $\langle \mathcal{I}, \mathcal{J} \rangle$ is separated by $\tilde{\delta}(\{x, y\})$. I.e., $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.

Case 3.2: $\text{diff}(\beta(y), \beta(x)) > \delta_1(\{x, z\})$.

In this case we can choose b with $\beta(x) < b \leq \beta(y)$ such that $\text{diff}(b, \beta(x)) = \delta_1(\{x, z\}) + 1$. An illustration is given in Figure 5. Then, $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is not separated by $\delta_1(\{x, z\})$. However, we know that $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is separated by $\delta_1(\{z, u'\})$, for some $u' \in \{\min, y, \max\}$. From the assumptions of Claim 7 we know that $|\text{diff}(a, \alpha(u'))| > \delta_1(\{z, u'\})$. Hence we must have that $|\text{diff}(b, \beta(u'))| \leq \delta_1(\{z, u'\})$. We now distinguish between the cases where u' can be chosen from $\{y, \max\}$, on the one hand, and where u' must be chosen as \min , on the other hand.




Figure 6: Situation in Case 3.2.2 of Claim 7.

Case 3.2.1: $u' \in \{y, \max\}$.

In this case,

$$\begin{aligned} \text{diff}(\beta(u'), \beta(x)) &= \text{diff}(\beta(u'), b) + \text{diff}(b, \beta(x)) \\ &\leq \delta_1(\{z, u'\}) + \delta_1(\{x, z\}) + 1 \\ &\leq \tilde{\delta}(\{x, u'\}). \end{aligned}$$

Since

$$\text{diff}(\alpha(u'), \alpha(x)) = \text{diff}(\alpha(u'), a) + \text{diff}(a, \alpha(x)) > \text{diff}(\beta(u'), \beta(x)),$$

we hence obtain that $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.

Case 3.2.2: $u' \notin \{y, \max\}$.

In this case, $\langle \mathcal{I}_a, \mathcal{J}_b \rangle$ is separated by $\delta_1(\{\min, z\})$, and we may assume that it is neither separated by $\delta_1(\{z, y\})$ nor by $\delta_1(\{z, \max\})$ nor by $\delta_1(\{x, z\})$. In particular, we must have that

$$\text{diff}(\beta(y), b) \geq \delta_1(\{z, y\}) + 1.$$

Therefore, for every b' with

$$b \leq b' < \beta(y) - \delta_1(\{z, y\}),$$

the following is true: $\langle \mathcal{I}_a, \mathcal{J}_{b'} \rangle$ is neither separated by $\delta_1(\{x, z\})$ nor by $\delta_1(\{z, y\})$, but, consequently, by $\delta_1(\{\min, z\})$ or by $\delta_1(\{z, \max\})$. Let b_1 be the largest such b' for which $\langle \mathcal{I}_a, \mathcal{J}_{b'} \rangle$ is separated by $\delta_1(\{\min, z\})$. In particular,

$$\text{diff}(b_1, \min^B) \leq \delta_1(\{\min, z\}).$$

An illustration is given in Figure 6.

Case 3.2.2.1: $\text{diff}(\beta(y), b_1+1) \leq \delta_1(\{z, y\})$.

In this case we know that

$$\text{diff}(\beta(y), \min^B) \leq \delta_1(\{z, y\}) + 1 + \delta_1(\{\min, z\}) \leq \tilde{\delta}(\{\min, y\}).$$

Furthermore,

$$\begin{aligned} \text{diff}(\alpha(y), \min^A) &= \text{diff}(\alpha(y), a) + \text{diff}(a, \min^A) \\ &\geq \delta_1(\{z, y\}) + 1 + \delta_1(\{\min, z\}) + 1. \end{aligned}$$

Therefore,

$$\text{diff}(\alpha(y), \min^A) \neq \text{diff}(\beta(y), \min^B),$$

and $\langle \mathcal{I}, \mathcal{J} \rangle$ is separated by $\tilde{\delta}(\{\min, y\})$. I.e., $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.

Case 3.2.2.2: $\text{diff}(\beta(y), b_1+1) > \delta_1(\{z, y\})$.

In this case we know (by the maximal choice of b_1) that $\langle \mathcal{I}_a, \mathcal{J}_{b_1+1} \rangle$ must be separated by $\delta_1(\{z, \max\})$. In particular, $\text{diff}(\max^{\mathcal{B}}, b_1+1) \leq \delta_1(z, \max)$. Therefore,

$$\text{diff}(\max^{\mathcal{B}}, \min^{\mathcal{B}}) \leq \delta_1(\{z, \max\}) + 1 + \delta_1(\{\min, z\}) \leq \tilde{\delta}(\{\min, \max\}).$$

Furthermore,

$$\begin{aligned} \text{diff}(\max^{\mathcal{A}}, \min^{\mathcal{A}}) &\geq \text{diff}(\max^{\mathcal{A}}, a) + \text{diff}(a, \min^{\mathcal{A}}) \\ &\geq \delta_1(\{z, \max\}) + 1 + \delta_1(\{\min, z\}) + 1. \end{aligned}$$

Therefore,

$$\text{diff}(\max^{\mathcal{A}}, \min^{\mathcal{A}}) \neq \text{diff}(\max^{\mathcal{B}}, \min^{\mathcal{B}}),$$

and $\langle \mathcal{I}, \mathcal{J} \rangle$ is separated by $\tilde{\delta}(\{\min, \max\})$. I.e., $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$.

We now have shown that $\tilde{\delta}$ is a separator for $\langle \mathcal{I}, \mathcal{J} \rangle$, if Case 3 applies.

Together with the Cases 1 and 2 we therefore obtain that the proof of Claim 7 is complete. \blacksquare

Since at least one of the Claims 1–7 must apply, the proof of Lemma 3.12 finally is complete. \blacksquare

Proof of part (c) of Lemma 3.9:

Let v be a node of \mathcal{T} that has exactly one child v_1 . Let δ be a minimal separator for $il(v)$, and let δ_1 be a minimal separator for $\langle A_1, B_1 \rangle := il(v_1)$. Our aim is to show that $w(\delta) \leq w(\delta_1) + 2$.

From Definition 3.8 we know that either $sl(v) = \neg$ or $sl(v) = \mathbf{Q}u$, for some $\mathbf{Q} \in \{\exists, \forall\}$ and $u \in \{x, y, z\}$.

Case 1: $sl(v) = \neg$

In this case we know from Definition 3.8 that $il(v) = \langle B_1, A_1 \rangle$. Therefore, δ_1 also is a (minimal) separator for $il(v)$. In particular, $w(\delta) = w(\delta_1) \leq w(\delta_1) + 2$.

Case 2: $sl(v) = \mathbf{Q}u$

In this case let $\tilde{\delta}$ be the *separator for $il(v)$* defined in Lemma 3.12. Since δ is a *minimal* separator for $il(v)$, it suffices to show that $w(\tilde{\delta}) \leq w(\delta_1) + 2$.

Let u', u'' be chosen such that $\{x, y, z\} = \{u, u', u''\}$. Using Definition 3.4 and the particular choice of $\tilde{\delta}$, it is straightforward to see that

$$c(\tilde{\delta}) = \tilde{\delta}(\{u', u''\}) \leq c(\delta_1) + 1 \tag{3.1}$$

and that

$$\tilde{\delta}(\{\min, \max\}) \leq b(\delta_1) + 1. \tag{3.2}$$

Furthermore, for arbitrary $\tilde{u}, \tilde{u}' \in \{x, y, z\}$ we have

$$\tilde{\delta}(\{\min, \tilde{u}\}) + \tilde{\delta}(\{\tilde{u}', \max\}) \leq b(\delta_1) + 2c(\delta_1) + 2, \tag{3.3}$$

which can be seen as follows: If $\tilde{u} = u$ or $\tilde{u}' = u$, then $\tilde{\delta}(\{\min, \tilde{u}\}) = 0$ or $\tilde{\delta}(\{\tilde{u}', \max\}) = 0$.

Consequently,

$$\tilde{\delta}(\{\min, \tilde{u}\}) + \tilde{\delta}(\{\tilde{u}', \max\}) \leq \text{MAX} \{b(\delta_1), b(\delta_1) + c(\delta_1) + 1\}.$$

If \tilde{u} and \tilde{u}' both are different from u , then

$$\tilde{\delta}(\{\min, \tilde{u}\}) = \text{MAX} \{ \delta_1(\{\min, \tilde{u}\}), \delta_1(\{\min, u\}) + \delta_1(\{u, \tilde{u}\}) + 1 \}$$

and

$$\tilde{\delta}(\{\tilde{u}', \max\}) = \text{MAX} \{ \delta_1(\{\tilde{u}', \max\}), \delta_1(\{u, \max\}) + \delta_1(\{u, \tilde{u}'\}) + 1 \}.$$

Therefore, $\tilde{\delta}(\{\min, \tilde{u}\}) + \tilde{\delta}(\{\tilde{u}', \max\}) \leq$

$$\text{MAX} \left\{ \begin{array}{l} \delta_1(\{\min, \tilde{u}\}) + \delta_1(\{\tilde{u}', \max\}), \\ \delta_1(\{\min, \tilde{u}\}) + \delta_1(\{u, \max\}) + \delta_1(\{u, \tilde{u}'\}) + 1, \\ \delta_1(\{\min, u\}) + \delta_1(\{u, \tilde{u}\}) + 1 + \delta_1(\{\tilde{u}', \max\}), \\ \delta_1(\{\min, u\}) + \delta_1(\{u, \tilde{u}\}) + 1 + \delta_1(\{u, \max\}) + \delta_1(\{u, \tilde{u}'\}) + 1 \end{array} \right\}$$

which, in turn, is less than or equal to

$$\text{MAX} \{ b(\delta_1), b(\delta_1) + c(\delta_1) + 1, b(\delta_1) + 2c(\delta_1) + 2 \}.$$

I.e., we have shown that (3.3) is valid.

From (3.2) and (3.3) we obtain

$$b(\tilde{\delta}) \leq b(\delta_1) + 2c(\delta_1) + 2. \quad (3.4)$$

From (3.1) and (3.4) we conclude that

$$\begin{aligned} w(\tilde{\delta})^2 &= c(\tilde{\delta})^2 + b(\tilde{\delta}) \\ &\leq (c(\delta_1) + 1)^2 + b(\delta_1) + 2c(\delta_1) + 2 \\ &= c(\delta_1)^2 + b(\delta_1) + 4c(\delta_1) + 3 \\ &\leq w(\delta_1)^2 + 4w(\delta_1) + 4 \\ &= (w(\delta_1) + 2)^2. \end{aligned}$$

Therefore, $w(\tilde{\delta}) \leq w(\delta_1) + 2$.

This completes the proof of part (c) of Lemma 3.9. ■

Remark 3.13. From the above proof it becomes clear, why Definition 3.4 fixes the *weight* of a separator by using the $\sqrt{}$ -function. Let us consider the, at first glance, more straightforward weight function $\hat{w}(\delta) := \text{MAX} \{c(\delta), b(\delta)\}$. In the proof of part (c) of Lemma 3.9 we then obtain from the items (3.1) and (3.4) that $\hat{w}(\delta) \leq 2c(\delta_1) + b(\delta_1) + 2 \leq 3\hat{w}(\delta_1) + 2$. Therefore, a modified version of Lemma 3.10, where item (c) is replaced by the condition “If v has exactly one child v_1 , then $w(v) \leq 3w(v_1) + 2$ ”, leads to a (much weaker) bound of the form $\|\mathcal{T}\| \geq c \cdot \lg(w(v))$. This, in turn, leads to a weaker version of Theorem 3.6, stating that $\|\psi\| \geq c \cdot \lg(\hat{w}(\delta))$. However, this bound can already be proven by a conventional Ehrenfeucht-Fraïssé game and does not only apply for $\text{FO}^3(<)$ -formulas but for $\text{FO}(<)$ -formulas in general and therefore is of no use for comparing the succinctness of FO^3 and FO . □

4. FO² vs. FO³

As a first application, Theorem 3.6 allows us to translate every FO³-sentence ψ into an FO²-sentence χ that is equivalent to ψ on linear orders and that has size polynomial in the size of ψ . To show this, we use the following easy lemmas.

Lemma 4.1. Let φ be an FO($<$, *succ*, *min*, *max*)-sentence, and let d be the quantifier depth of φ . For all $N \geq 2^{d+1}$, $\mathcal{A}_N \models \varphi$ if, and only if, $\mathcal{A}_{2^{d+1}} \models \varphi$. \square

A proof can be found, e.g., in the textbook [3].

Lemma 4.2. For all $\ell \in \mathbb{N}$ there are FO²($<$)-sentences χ_ℓ and $\chi_{\geq \ell}$ of size $\mathcal{O}(\ell)$ such that, for all $N \in \mathbb{N}$, $\mathcal{A}_N \models \chi_\ell$ (respectively, $\chi_{\geq \ell}$) iff $N = \ell$ (respectively, $N \geq \ell$). \square

Proof: We choose

$$\chi'_{\geq 0}(x) := (x = x),$$

and, for all $\ell \geq 0$,

$$\chi'_{\geq \ell+1}(x) := \exists y (y < x) \wedge \chi_{\geq \ell}(y).$$

Obviously, for all $N \in \mathbb{N}$ and all $a \in \{0, \dots, N\}$, we have $\mathcal{A}_N \models \chi'_{\geq \ell}(a)$ iff $a \geq \ell$. Therefore, for every $\ell \in \mathbb{N}$, we can choose $\chi_{\geq \ell} := \exists x \chi'_{\geq \ell}(x)$ and $\chi_\ell := \chi_{\geq \ell} \wedge \neg \chi_{\geq \ell+1}$. \blacksquare

Theorem 4.3.

On linear orders, FO³($<$, *succ*, *min*, *max*)-sentences are $\mathcal{O}(m^4)$ -succinct in FO²($<$)-sentences. \square

Proof: Let ψ be an FO³($<$, *succ*, *min*, *max*)-sentence. Our aim is to find an FO²($<$)-sentence χ of size $\mathcal{O}(\|\psi\|^4)$ such that, for all $N \in \mathbb{N}$, $\mathcal{A}_N \models \chi$ iff $\mathcal{A}_N \models \psi$.

If ψ is satisfied by *all* linear orders or by *no* linear order, χ can be chosen in a straightforward way. In all other cases we know from Lemma 4.1 that there exists a $D \in \mathbb{N}$ such that either

- (1.) $\mathcal{A}_D \models \psi$ and, for all $N > D$, $\mathcal{A}_N \not\models \psi$, or
- (2.) $\mathcal{A}_D \not\models \psi$ and, for all $N > D$, $\mathcal{A}_N \models \psi$.

In particular, ψ is an FO³-sentence that distinguishes between the linear orders \mathcal{A}_D and \mathcal{A}_{D+1} . From Corollary 3.7 we therefore know that $\|\psi\| \geq \frac{1}{2}\sqrt{D}$.

We next construct an FO²-sentence χ equivalent to ψ : Let χ_ℓ and $\chi_{\geq \ell}$ be the FO²($<$)-sentences from Lemma 4.2. Let χ' be the disjunction of the sentences χ_ℓ for all those $\ell \leq D$ with $\mathcal{A}_\ell \models \psi$. Finally, if $\mathcal{A}_{D+1} \models \psi$, then choose $\chi := \chi' \vee \chi_{\geq D+1}$; otherwise choose $\chi := \chi'$. Obviously, χ is an FO²($<$)-sentence equivalent to ψ , and

$$\|\chi\| = \mathcal{O}\left(\sum_{\ell=0}^{D+1} \ell\right) = \mathcal{O}(D^2) = \mathcal{O}(\|\psi\|^4).$$

This completes the proof of Theorem 4.3. \blacksquare

5. FO³ vs. FO

Using Theorem 3.6, we will show in this section that there is an exponential succinctness gap between FO and FO³ on linear orders.

Lemma 5.1. For every FO($<, succ, min, max$)-sentence φ there is an FO²($<$)-sentence ψ of size $\|\psi\| \leq 2^{\mathcal{O}(\|\varphi\|)}$ which is equivalent to φ on the class of linear orders. \square

Proof: Let φ be an FO($<, succ, min, max$)-sentence, and let d be the quantifier depth of φ . In particular, $\|\varphi\| \geq d$. From Lemma 4.1 we know that, for all $N \geq 2^{d+1}$, $\mathcal{A}_N \models \varphi$ if, and only if, $\mathcal{A}_{2^{d+1}} \models \varphi$.

We use, for every $\ell \in \mathbb{N}$, the sentences χ_ℓ and $\chi_{\geq \ell}$ of Lemma 4.2. Let ψ' be the disjunction of the sentences χ_ℓ for all $\ell < 2^{d+1}$ such that $\mathcal{A}_\ell \models \varphi$. Finally, if $\mathcal{A}_{2^{d+1}} \models \varphi$, then choose $\psi := \psi' \vee \chi_{\geq 2^{d+1}}$; otherwise choose $\psi := \psi'$. Obviously, ψ is an FO²($<$)-sentence equivalent to φ , and

$$\|\psi\| = \mathcal{O}\left(\sum_{\ell=0}^{2^{d+1}} \ell\right) = \mathcal{O}(2^{2(d+1)}) = 2^{\mathcal{O}(\|\varphi\|)}.$$

■

Lemma 5.2. For all $m \in \mathbb{N}$ there are FO⁴($<$)-sentences φ_m and sets A_m and B_m of interpretations, such that $A_m \models \varphi_m$, $B_m \models \neg\varphi_m$, $\|\varphi_m\| = \mathcal{O}(m)$, and every FO³($<, succ, min, max$)-sentence ψ_m equivalent to φ_m has size $\|\psi_m\| \geq 2^{\frac{1}{2}m-1}$. \square

Proof: For every $N \in \mathbb{N}$ let $\alpha_N : \{\min, x, y, z, \max\} \rightarrow \{0, \dots, N\}$ be the assignment with $\alpha_N(x) = \alpha_N(\min) = 0$ and $\alpha_N(y) = \alpha_N(z) = \alpha(\max) = N$.

For every $m \in \mathbb{N}$ we choose $A_m := \{(\mathcal{A}_{2^m}, \alpha_{2^m})\}$ and $B_m := \{(\mathcal{A}_{2^{m+1}}, \alpha_{2^{m+1}})\}$.

Step 1: Choice of φ_m .

Inductively we define FO⁴($<$)-formulas $\varphi'_m(x, y)$ expressing that $|\text{diff}(x, y)| = 2^m$ via

$$\varphi'_m(x, y) := \exists z \forall u (u = x \vee u = y) \rightarrow \varphi'_{m-1}(z, u)$$

(and $\varphi'_0(x, y)$ chosen appropriately).

It is straightforward to see that $\|\varphi'_m\| = \mathcal{O}(m)$ and that $\varphi'_m(x, y)$ expresses that $|\text{diff}(x, y)| = 2^m$.

Therefore,

$$\varphi_m := \exists x \exists y \varphi'_m(x, y) \wedge \neg \exists z (z < x \vee y < z)$$

is an FO⁴($<$)-sentence with the desired properties.

Step 2: Size of equivalent FO³-sentences.

For every $m \in \mathbb{N}$ let ψ_m be an FO³($<, succ, min, max$)-sentence with $A_m \models \psi_m$ and $B_m \models \neg\psi_m$. From Corollary 3.7 we conclude that

$$\|\psi_m\| \geq \frac{1}{2}\sqrt{2^m} = 2^{\frac{1}{2}m-1}.$$

This completes the proof of Lemma 5.2. \square

From Lemma 5.1 and 5.2 we directly obtain

Theorem 5.3.

On the class of linear orders, FO($<$)-sentences are $2^{\mathcal{O}(m)}$ -succinct in FO³($<$)-sentences, but already FO⁴($<$)-sentences are not $2^{\mathcal{O}(m)}$ -succinct in FO³($<$)-sentences. \square

Note that the relation *succ* and the constants *min* and *max* are easily definable in FO³($<$) and could therefore be added in Theorem 5.3.

6. FO vs. MSO

In this section we compare the succinctness of FO and the FO-expressible fragment of monadic second-order logic (for short: MSO). This section's main result is a non-elementary succinctness gap between FO and the FO-expressible fragment of MSO on the class of linear orders (Theorem 6.5).

The main idea for proving this succinctness gap is to encode natural numbers by strings in such a way that there are extremely short MSO-formulas for “decoding” these strings back into numbers. The method goes back to Stockmeyer and Meyer [16, 15]; the particular encoding used in the present paper has been introduced in [5]. To formally state and prove this section's main result, we need some further notation:

We write $\text{Mon}\Sigma_1^1$ for the class of all MSO-formulas that consist of a prefix of existential set quantifiers, followed by a first-order formula. By $\exists X \text{FO}$ we denote the fragment of $\text{Mon}\Sigma_1^1$ with only a single existential set quantifier.

Let $\text{Tower} : \mathbb{N} \rightarrow \mathbb{N}$ be the function which maps every $h \in \mathbb{N}$ to the tower of 2s of height h . I.e., $\text{Tower}(0) = 1$ and, for every $h \in \mathbb{N}$, $\text{Tower}(h+1) = 2^{\text{Tower}(h)}$.

We use the following notations of [5]:

For $h \geq 1$ 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 $L(n) = \lfloor \log(n-1) \rfloor + 1$, 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:

$$\begin{aligned} \mu_1(0) &:= \langle 1 \rangle \langle /1 \rangle, \quad \text{and} \\ \mu_1(n) &:= \langle 1 \rangle \text{bit}(0, n-1) \text{bit}(1, n-1) \cdots \text{bit}(L(n)-1, n-1) \langle /1 \rangle, \end{aligned}$$

for $n \geq 1$. For $h \geq 2$ we let

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

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

For $h \in \mathbb{N}$ let $H := \text{Tower}(h)$. Let $\Sigma_h^\bullet := \Sigma_{h+1} \cup \{\bullet\}$, and let

$$v_h := \langle h+1 \rangle \mu_h(0) \bullet \mu_h(1) \bullet \cdots \mu_h(H-1) \bullet \langle /h+1 \rangle.$$

We consider the string-language $(v_h)^+$, containing all strings that are the concatenation of one or more copies of v_h . Let w_h be the (unique) string in $(v_h)^+$ that consists of exactly 2^H copies of v_h .

We write τ_h for the signature that consists of the symbol \langle and a unary relation symbol P_σ , for every $\sigma \in \Sigma_h^\bullet$. Non-empty strings over Σ_h^\bullet are represented by τ_h -structures in the

usual way (cf., e.g., [3]). We will shortly write $w \models \varphi$ to indicate that the τ_h -structure associated with a Σ_h^\bullet -string w satisfies a given τ_h -sentence φ .

The following lemma is our key tool for proving that there is a non-elementary succinctness gap between FO and the FO-expressible fragment of MSO.

Lemma 6.1. For every $h \in \mathbb{N}$ there is an $\exists X$ FO(τ_h)-sentence Φ_h of size $\mathcal{O}(h^2)$, such that the following is true for all strings w over the alphabet Σ_h^\bullet : $w \models \Phi_h$ iff $w = w_h$. \square

For proving Lemma 6.1 we need the following:

Lemma 6.2 ([5, Lemma 8]). For all $h \in \mathbb{N}$ there are FO(τ_h)-formulas $equal_h(x, y)$ of size¹ $\mathcal{O}(h)$ 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 $m, n \in \{0, \dots, Tower(h)\}$: 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 $w \models equal_h(a, b)$ if, and only if, $m = n$. \square

Using the above lemma, it is an easy exercise to show

Lemma 6.3. For every $h \in \mathbb{N}$ there is an FO(τ_h)-formula $inc_h(x, y)$ of size $\mathcal{O}(h)$ 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 $m, n \in \{0, \dots, Tower(h)\}$: 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 $w \models inc_h(a, b)$ if, and only if, $m+1 = n$. \square

We also need

Lemma 6.4. For every $h \in \mathbb{N}$, the language $(v_h)^+$ is definable by an FO(τ_h)-sentence $\varphi_{(v_h)^+}$ of size $\mathcal{O}(h^2)$. I.e., for all strings w over the alphabet Σ_h^\bullet we have $w \models \varphi_{(v_h)^+}$ if, and only if, $w \in (v_h)^+$. \square

Proof: The proof proceeds in 2 steps:

Step 1: Given $j \geq 1$, we say that a string w over Σ_h^\bullet satisfies the condition $C(j)$ if, and only if, for every position x (respectively, y) in w that carries the letter $\langle j \rangle$ (respectively, $\langle /j \rangle$) the following is true: There is a position y to the right of x that carries the letter $\langle /j \rangle$ (respectively, a position x to the left of y that carries the letter $\langle j \rangle$), such that the substring u of w that starts at position x and ends at position y is of the form $\mu_j(n)$ for some $n \in \{0, \dots, Tower(j)-1\}$.

We will construct, for all $j \in \{1, \dots, h\}$, FO(τ_h)-sentences ok_j of size $\mathcal{O}(j)$ such that the following is true for all $j \leq h$ and all strings w over Σ_h^\bullet that satisfy the conditions $C(j')$ for all $j' < j$:

$$w \models ok_j \iff w \text{ satisfies the condition } C(j).$$

Simultaneously we will construct, for all $j \in \{1, \dots, h\}$, FO(τ_h)-sentences $max_j(x)$ of size $\mathcal{O}(j)$ such that the following is true for all $j \leq h$, all strings w that satisfy the conditions $C(1), \dots, C(j)$, and all positions x in w :

$$w \models max_j(x) \iff \begin{array}{l} x \text{ is the starting position of} \\ \text{a substring of } w \text{ of the form} \\ \mu_j(Tower(j)-1). \end{array}$$

¹In [5], 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.

For the base case $j = 1$ note that $Tower(1) - 1 = 1$ and, by the definition of $\mu_1(n)$, $\mu_1(0) = \langle 1 \rangle \langle /1 \rangle$ and $\mu_1(1) = \langle 1 \rangle 0 \langle /1 \rangle$. It is straightforward to write down a formula ok_1 that expresses the condition $C(1)$. Furthermore, $max_1(x)$ states that the substring of length 3 starting at position x is of the form $\langle 1 \rangle 0 \langle /1 \rangle$.

For $j > 1$ assume that the formula max_{j-1} has already been constructed. For the construction of the formula ok_j we assume that the underlying string w satisfies the conditions $C(1), \dots, C(j-1)$.

The formula ok_j states that whenever x (respectively, y) is a position in w that carries the letter $\langle j \rangle$ (respectively, $\langle /j \rangle$) the following is true: There is a position y to the right of x that carries the letter $\langle /j \rangle$ (respectively, a position x to the left of y that carries the letter $\langle j \rangle$), such that the substring u of w that starts at position x and ends at position y is of the form $\mu_j(n)$ for some $n \in \{0, \dots, Tower(j) - 1\}$, i.e.,

1. the letters $\langle j \rangle$ and $\langle /j \rangle$ only occur at the first and the last position of u ,
2. whenever a position x' carries the letter $\langle /j-1 \rangle$, position $x'+1$ carries the letter 0 or 1, and position $x'+2$ carries the letter $\langle j-1 \rangle$ or the letter $\langle /j \rangle$,
3. either $u = \langle j \rangle \langle /j \rangle$, or the prefix of length 3 of u is of the form $\langle j \rangle \langle j-1 \rangle \langle /j-1 \rangle$,
4. whenever x' and y' are positions in u carrying the letter $\langle j-1 \rangle$ such that $x' < y'$ and no position between x' and y' carries the letter $\langle j-1 \rangle$, the formula $inc_{j-1}(x', y')$ from Lemma 6.3 is satisfied,
5. if the rightmost position x'' in u that carries the letter $\langle j-1 \rangle$ satisfies the formula $max_{j-1}(x'')$, then there must be at least one position x''' in u that carries the letter 0 such that $x''' - 1$ carries the letter $\langle /j-1 \rangle$.

Note that items 1.–4. ensure that u is indeed of the form $\mu_j(n)$, for some $n \in \mathbb{N}$. Item 5 guarantees that $n \in \{0, \dots, Tower(j) - 1\}$ because of the following: recall from the definition of the string $\mu_j(n)$ that $\mu_j(n)$ involves the (reverse) binary representation of the number $n-1$. In particular, for $n := Tower(j) - 1$, we need the (reverse) binary representation of the number $Tower(j) - 2$, which is of the form $011 \dots 11$ and of length $Tower(j) - 1$, i.e., its highest bit has the number $Tower(j) - 1$.

It is straightforward to see that the items 1.–5. and therefore also the formula ok_j can be formalised by an $FO(\tau_h)$ -formula of size $\mathcal{O}(j)$, and that this formula exactly expresses condition $C(j)$.

Furthermore, the formula $max_j(x)$ assumes that x is the starting position of a substring u of w of the form $\mu_j(n)$, for some $n \in \mathbb{N}$; and $max_j(x)$ states that

1. the (reverse) binary representation of n , i.e., the $\{0, 1\}$ -string built from the letters in u that occur directly to the right of letters $\langle /j-1 \rangle$, is of the form $011 \dots 11$, and
2. the highest bit of n has the number $Tower(j) - 1$, i.e., the rightmost position y in u that carries the letter $\langle j-1 \rangle$ satisfies the formula $max_{j-1}(y)$.

Obviously, $max_j(x)$ can be formalised in $FO(\tau_h)$ by a formula of size $\mathcal{O}(j)$. Finally, this completes Step 1.

Step 2: A string w over Σ_h^\bullet belongs to the language $(v_h)^+$ if, and only if, all the following conditions are satisfied:

1. w satisfies $ok_1 \wedge \dots \wedge ok_h$,
2. the first position in w carries the letter $\langle h+1 \rangle$, the last position in w carries the letter $\langle /h+1 \rangle$, the letter $\langle h+1 \rangle$ occurs at a position $x > 1$ iff position $x-1$ carries the letter $\langle /h+1 \rangle$, and the letter \bullet occurs at a position x iff position $x-1$ carries the letter $\langle /h \rangle$ and position $x+1$ carries the letter $\langle h \rangle$ or $\langle /h+1 \rangle$,

3. whenever x (respectively, y) is a position in w that carries the letter $\langle \mathbf{h+1} \rangle$ (respectively, $\langle / \mathbf{h+1} \rangle$) the following is true: There is a position y to the right of x that carries the letter $\langle / \mathbf{h+1} \rangle$ (respectively, a position x to the left of y that carries the letter $\langle \mathbf{h+1} \rangle$), such that the substring u of w that starts at position x and ends at position y is of the form v_h , i.e.,
- ★ the letters $\langle \mathbf{h+1} \rangle$ and $\langle / \mathbf{h+1} \rangle$ only occur at the first and the last position of u ,
 - ★ the prefix of length 3 of u is of the form $\langle \mathbf{h+1} \rangle \langle \mathbf{h} \rangle \langle / \mathbf{h} \rangle$, and the suffix of length 3 of u is of the form $\langle / \mathbf{h} \rangle \bullet \langle / \mathbf{h+1} \rangle$,
 - ★ whenever x' and y' are positions in u carrying the letter $\langle \mathbf{h} \rangle$ such that $x' < y'$ and no position between x' and y' carries the letter $\langle \mathbf{h} \rangle$, the formula $inc_h(x', y')$ from Lemma 6.3 is satisfied,
 - ★ the rightmost position x'' in u that carries the letter $\langle \mathbf{h} \rangle$ satisfies the formula $max_h(x'')$.

Using the formulas constructed in Step 1 and the preceding lemmas, it is straightforward to see that this can be formalised by an $\text{FO}(\tau_h)$ -formula $\varphi_{(v_h)^+}$ of size $\mathcal{O}(h^2)$. This finally completes the proof of Lemma 6.4. ■

Finally, we are ready for the

Proof of Lemma 6.1:

To determine whether an input string w is indeed the string w_h , one can proceed as follows: First, we make sure that the underlying string w belongs to $(v_h)^+$ via the $\text{FO}(\tau_h)$ -formula $\varphi_{(v_h)^+}$ of Lemma 6.4. Afterwards we, in particular, know that in each $\langle \mathbf{h+1} \rangle \cdots \langle / \mathbf{h+1} \rangle$ -block is of the form v_h and therefore contains exactly H positions that carry the letter \bullet . Now, to each \bullet -position in w we assign a letter from $\{0, 1\}$ in such a way that the $\{0, 1\}$ -string built from these assignments is an H -numbering, i.e., of one of the following forms:

1. $\text{BIN}_H(0) \text{BIN}_H(1) \text{BIN}_H(2) \cdots \text{BIN}_H(n)$, for some $n < 2^H$,
2. $\text{BIN}_H(0) \text{BIN}_H(1) \cdots \text{BIN}_H(2^H - 1) \left(\text{BIN}_H(0)^m \right)$, for some $m \geq 0$.

Here, $\text{BIN}_H(n)$ denotes the reverse binary representation of length H of the number $n < 2^H$. For example, $\text{BIN}_4(2) = 0100$ and $\text{BIN}_4(5) = 1010$. Of course, w is the string w_h , i.e., consists of exactly 2^H copies of v_h , if and only if the H -numbering's assignments in the rightmost copy of v_h form the string $\text{BIN}_H(2^H - 1)$, i.e., if and only if every \bullet -position in this copy of v_h was assigned the letter 1.

One way of assigning letters from $\{0, 1\}$ to the \bullet -positions in w is by choosing a set X of \bullet -positions with the intended meaning that a \bullet -position x is assigned the letter 1 if $x \in X$ and the letter 0 if $x \notin X$.

Using the $\text{FO}(\tau_h)$ -formulas $equal_h$ of Lemma 6.2 and $\varphi_{(v_h)^+}$ of Lemma 6.4, it is straightforward to construct the desired $\exists X \text{FO}(\tau_h)$ -formula Φ_h of size $\mathcal{O}(h)$.

This completes the proof of Lemma 6.1. ■

As a consequence of Lemma 6.1 and Lemma 4.1 one obtains a non-elementary succinctness gap between FO and MSO on the class of linear orders:

Theorem 6.5.

The $\text{FO}(<)$ -expressible fragment of $\text{Mon}\Sigma_1^1$ is not $Tower(o(\sqrt{m}))$ -succinct in $\text{FO}(<)$ on the class of linear orders. □

Proof: Recall that, for every $N \in \mathbb{N}$, \mathcal{A}_N denotes the linear order with universe $\{0, \dots, N\}$.

For every $h \in \mathbb{N}$ let $\ell(h) := |w_h| - 1$, where $|w_h|$ denotes the length of the string w_h . We say that a sentence χ *defines* the linear order $\mathcal{A}_{\ell(h)}$ if, and only if, $\mathcal{A}_{\ell(h)}$ is the unique structure in $\{\mathcal{A}_N : N \in \mathbb{N}\}$ that satisfies χ . For every $h \in \mathbb{N}$ we show the following:

- (a) Every $\text{FO}(<, \text{succ}, \text{min}, \text{max})$ -sentence ψ_h that defines $\mathcal{A}_{\ell(h)}$ has size $\|\psi_h\| \geq \text{Tower}(h)$.
- (a) There is a $\text{Mon}\Sigma_1^1(<)$ -sentence Ψ_h of size $\|\Psi_h\| = \mathcal{O}(h^2)$ that defines $\mathcal{A}_{\ell(h)}$.

Ad (a):

Since w_h consists of $2^{\text{Tower}(h)}$ copies of v_h , we know that $\ell(h) \geq 2^{\text{Tower}(h)}$. Therefore, every $\text{FO}(<)$ -sentence ψ_h that defines $\mathcal{A}_{\ell(h)}$ has quantifier depth, and therefore size, at least $\text{Tower}(h)$ (cf., Lemma 4.1).

Ad (b):

Let $\exists X \varphi_h$ be the $\exists X \text{FO}(\tau_h)$ -sentence obtained from Lemma 6.1. It is straightforward to formulate an $\text{FO}(\tau_h)$ -sentence ξ_h of size $\mathcal{O}(h^2)$ which expresses that every element in the underlying structure's universe belongs to exactly one of the sets P_σ , for $\sigma \in \Sigma_h^\bullet$.

The $\text{Mon}\Sigma_1^1(<)$ -sentence

$$\Psi_h := (\exists P_\sigma)_{\sigma \in \Sigma_h^\bullet} \exists X (\xi_h \wedge \varphi_h)$$

expresses that the nodes of the underlying linear order can be labelled with letters in Σ_h^\bullet in such a way that one obtains the string w_h . Such a labeling is possible if, and only if, the linear order has length $|w_h|$. I.e., Ψ_h *defines* $\mathcal{A}_{\ell(h)}$. Furthermore, $\|\Psi_h\| = \mathcal{O}(h^2)$, because $\|\xi_h\| = \mathcal{O}(h^2)$ and $\|\varphi_h\| = \mathcal{O}(h^2)$.

This completes the proof of Theorem 6.5 ■

Let us remark that by modifying the proof of the above result, one can also show that the $\text{FO}(<)$ -expressible fragment of *monadic least fixed point logic*, MLFP, is non-elementarily more succinct than $\text{FO}(<)$ on the class of linear orders.

7. CONCLUSION

Our main technical result is a lower bound on the size of a 3-variable formula defining a linear order of a given size. We introduced a new technique based on Adler-Immerman games that might be also useful in other situations. A lot of questions remain open, let us just mention a few here:

- Is first-order logic on linear orders $\text{poly}(m)$ -succinct in its 4-variable fragment, or is there an exponential gap?
- As a next step, it would be interesting to study the succinctness of the finite-variable fragments on strings, that is, linear orders with additional unary relation symbols. It is known that on finite strings, the 3-variable fragment of first-order logic has the same expressive power as full first-order logic. Our results show that there is an at least exponential succinctness gap between the 3-variable and the 4-variable fragment. We do not know, however, if this gap is only singly exponential or larger, and we also do not know what happens beyond 4 variables.
- Another interesting task is to study the succinctness of various extensions of (finite variable fragments of) first-order logic by transitive closure operators.

- It also remains to be investigated if our results can possibly help to settle the long standing open problem of whether the 3-variable and 4-variable fragments of first-order logic have the same expressive power on the class of all ordered finite structures.

Finally, let us express our hope that techniques for proving lower bounds on succinctness will further improve in the future so that simple results such as ours will have simple proofs!

REFERENCES

- [1] M. Adler and N. Immerman. An $n!$ lower bound on formula size. *ACM Transactions on Computational Logic*, 4(3):296–314, 2003.
- [2] N. Alechina and N. Immerman. Reachability logic: An efficient fragment of transitive closure logic. *Logic Journal of the IGPL*, 8(3):325–338, 2000.
- [3] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.
- [4] K. Etessami, M. Y. Vardi, and T. Wilke. First-order logic with two variables and unary temporal logic. *Information and Computation*, 179(2):279–295, 2002.
- [5] M. Frick and M. Grohe. The complexity of first-order and monadic second-order logic revisited. *Annals of Pure and Applied Logic*, 130(1–3):3–31, 2004.
- [6] M. Frick, M. Grohe, and C. Koch. Query evaluation on compressed trees. In *Proceedings of the 18th IEEE Symposium on Logic in Computer Science (LICS'03)*, pages 188–197, 2003.
- [7] G. Gottlob and C. Koch. Monadic datalog and the expressive power of web information extraction languages. *Journal of the ACM*, 51(1):74–113, 2004.
- [8] M. Grohe and N. Schweikardt. Comparing the succinctness of monadic query languages over finite trees. *RAIRO – Theoretical Informatics and Applications*, 38(4):343–374, 2004.
- [9] M. Grohe and N. Schweikardt. The succinctness of first-order logic on linear orders. In *Proceedings of the 19th IEEE Symposium on Logic in Computer Science (LICS'04)*, pages 438–447, 2004.
- [10] N. Immerman. *Descriptive Complexity*. Springer, 1999.
- [11] H. Kamp. *Tense Logic and the theory of linear order*. PhD thesis, University of California, Los Angeles, 1968.
- [12] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [13] F. Neven. *Design and Analysis of Query Languages for Structured Documents – A Formal and Logical Approach*. PhD thesis, Limburgs Universitair Centrum, 1999.
- [14] F. Neven and T. Schwentick. Query automata over finite trees. *Theoretical Computer Science*, 275(1-2):633–674, 2002.
- [15] L. J. Stockmeyer. *The Complexity of Decision Problems in Automata Theory*. PhD thesis, Department of Electrical Engineering, MIT, 1974.
- [16] L. J. Stockmeyer and A. R. Meyer. *Word problems requiring exponential time*. In *Proceedings of the 5th ACM Symposium on Theory of Computing (STOC'73)*, pages 1–9, 1973.
- [17] M. Y. Vardi. The complexity of relational query languages. In *Proceedings of the 14th ACM Symposium on Theory of Computing (STOC'82)*, pages 137–146, 1982.
- [18] T. Wilke. CTL^+ is exponentially more succinct than CTL. In *Proceedings of the 19th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'99)*, volume 1738 of *Lecture Notes in Computer Science*, pages 110–121. Springer, 1999.