On the ordered conjecture

Yijia Chen
Department of Computer Science
Shanghai Jiaotong University, China
yijia.chen@cs.sjtu.edu.cn

Jörg Flum
Mathematisches Institut
Albert-Ludwigs-Universität Freiburg, Germany
joerg.flum@math.uni-freiburg.de

***Abstract*—It is well-known that least fixed-point logic LFP captures the complexity class PTIME on ordered structures. The ordered conjecture claims that LFP is more expressive than first-order logic FO (in short, LFP $>$ FO) on every infinite class O of finite ordered structures. We present two methods which yield that LFP $>$ FO on various types of classes of ordered structures. The first method, the *model-checking method*, among others, can be applied for all such classes O of bounded cliquewidth. By the second method, *the padding method*, we show that for classes O of "bounded treewidth," more precisely, for classes O such that there is a bound for the treewidth of the successor structures associated with the members of O, even DTC $>$ FO on O, where DTC denotes the deterministic transitive closure logic, a logic that captures the complexity class L on ordered structures. Furthermore, with the padding method we get that for every infinite class of ordered structures O we have DTC $>$ FO on the class of all ordered sums $\mathcal{A} \oplus \mathcal{B}$ with $\mathcal{A}, \mathcal{B} \in$ O.**

Under some complexity theoretic assumption, we prove the existence of a class O of ordered structures such that on O not only LFP $>$ FO but even LFP has the expressive power of existential second-order logic Σ^1_1. Furthermore, we characterize those classes of structures whose corresponding class of all ordered versions has bounded treewidth.

I. Introduction

Least fixed-point logic LFP is the extension of first-order logic FO obtained by adding a least fixed-point operator for positive first-order formulas. On the class of all *ordered* structures[1] LFP captures the complexity class P (= PTIME) and is hence more expressive than FO. Thus, LFP $>$ FO on the class of *all* structures. However, there are classes of unordered structures on which LFP collapses to first-order logic FO. In his paper [22] McColm studies the collapse of LFP to FO (in short, LFP $\equiv$ FO); in particular, he states a certain conjecture concerning necessary and sufficient conditions for such a collapse. The following special case of McColm's conjecture, which Kolaitis and Vardi [19] call the *ordered conjecture* (OC), still remains open:

(OC) *If* O *is an infinite class of ordered structures, then* LFP $>$ FO *on* O.

The ordered conjecture is linked with open questions in complexity theory. On the one hand, Dawar and Hella [10] showed that on every infinite class of ordered structures partial fixed-point logic PFP, the logic capturing polynomial space (PSPACE) on the class of all ordered structures, is more expressive than FO. Hence, a refutation of (OC) would yield PTIME $\neq$ PSPACE. On the other hand, Dawar et al. [11] proved that (OC) implies Lin-H $\neq$ E. Here, Lin-H is the class of problems that can be solved by a linear time, bounded depth, alternating machine and E $=$ DTIME$(2^{O(n)})$.

The following quest was one of the starting points of our analysis of (OC):

To what extent is the notion of treewidth helpful in the study of (OC)?

Clearly, every infinite class of ordered structures has unbounded treewidth (as any two elements of such a structure are comparable in the ordering). As in LFP we can define the ordering from its successor relation, instead of an ordered structure $(\mathcal{A}, <^{\mathcal{A}})$ we consider the associated successor structure $(\mathcal{A}, S(<^{\mathcal{A}}))$, where $S(<^{\mathcal{A}})$ denotes the binary successor relation associated with $<^{\mathcal{A}}$. We divide the ordered conjecture into two parts, the *bounded ordered conjecture* (BOC) and the *unbounded ordered conjecture* (UOC):

(BOC) [(UOC)] *If* O *is an infinite class of ordered structures whose class* S(O) *of associated successor structures has* bounded [unbounded] *treewidth, then* LFP $>$ FO *on* O.

As somehow classes with unbounded treewidth have a richer structure than bounded ones, at first glance one could think that it should be easier to prove (UOC) than (BOC). However, probably it will not be easy to verify (UOC) as the class of ordered structures used by Dawar et al. in [11] to show the implication "(OC) $\Rightarrow$ Lin-H $\neq$ E" is

[1] If not stated otherwise explicitly, all structures in this paper are assumed to be finite.

of "unbounded treewidth," hence the implication "(UOC) $\Rightarrow$ Lin-H $\neq$ E" is true.

Our quest concerning the role of treewidth in the study of (OC) was motivated by results of [11] and [24]. In [11], Dawar and Hella showed that LFP is more expressive than FO on every infinite class O of ordered structures in a vocabulary which besides $<$ only contains unary relation symbols; of course, S(O) has bounded treewdth for such an O. A result of Podewski and Ziegler [24] implies that for every infinite class O of ordered structures with $\mathrm{S(O)}$ of bounded treewidth (it even suffices to assume that $\mathrm{S(O)}$ only is of bounded local treewidth) no ordering at all is FO-definable. Hence, DTC $>$ FO on S(O) (as the ordering is DTC-definable from its successor relation). Is it possible to show that DTC $>$ FO or, at least, that LFP $>$ FO even on O, that is, does (BOC) hold?

In fact, we could prove (BOC). There are two distinct methods yielding this result. The first one, the model-cheking method, even yields that LFP $>$ FO on classes of bounded cliquewidth. By the second method, the padding method, (under the hypotheses of (BOC)) we even get that DTC $>$ FO. By the same method we show that for *every* infinite class of ordered structures O the logic DTC does not collapse to FO on the class of ordered sums $\mathcal{A} \oplus \mathcal{B}$ with $\mathcal{A}, \mathcal{B} \in$ O.

We were also encouraged to study the validity of statements similar to the ordered conjecture as in previous papers [5], [6] we introduced a complexity-theoretic assumption (namely $\mathrm{NP[TC]} \neq \mathrm{P[TC]}$), whose failure implies:

there exists an infinite set $I \subseteq \{0,1\}^*$ such that $Q \cap I \in$ P for all $Q \subseteq \{0,1\}^*$ in NP. (1)

At first glance, this seems to be a complexity-theoretic version of $\Sigma^1_1 \equiv$ LFP on (a class of structures related to) I. Here, Σ^1_1 denotes existential second-order logic, the fragment of second-order logic that captures NP by Fagin's Theorem. However, it turned out that to draw this conclusion one needs a version of (1) for functions. Under this hypothesis, we even get a class of ordered structures O such that $\Sigma^1_1 \equiv$ LFP $>$ FO on O.

Finally, we characterize those classes of structures whose corresponding class of *all* ordered versions have bounded treewidth.

II. Some preliminaries

We start by fixing notation and recalling some definitions and some results. For $n \in \mathbb{N}$ we denote by $[n]$ the set $\{1, \dots, n\}$.

Structures. A *vocabulary* τ is a finite set of relation symbols. Each relation symbol has an *arity*. A *structure* $\mathcal{A}$ of vocabulary τ, or τ-*structure*, consists of a nonempty set A called the *universe*, and an interpretation $R^{\mathcal{A}} \subseteq A^r$ of each r-ary relation symbol $R \in \tau$. As we already mentioned, *if not stated otherwise explicitly, structures are assumed to have a finite universe.* If τ contains a binary relation symbol $<$ and in the structure $\mathcal{A}$ the relation $<^{\mathcal{A}}$ is a (total) ordering of the universe, then $\mathcal{A}$ is an *ordered* structure. To avoid technicalities, where necessary, we assume that structures have as universe the set $[n]$ for some $n \geq 1$ and that in the ordered case $<$ is interpreted by the natural ordering on $[n]$. Then, in a canonical way, we identify structures with strings over the alphabet $\{0,1\}$; in particular, the size $|\mathcal{A}|$ of a structure $\mathcal{A}$ is the length of the string $\mathcal{A}$.

Often we deal with classes of structures. Thereby we always assume that all structures of a fixed class have the same vocabulary. But distinct vocabularies may correspond to distinct classes. By convention, we say that *a class is finite if, up to isomorphism, it only contains finitely many structures; otherwise, it is infinite.* Furthermore, O *will always denote an infinite class of ordered structures.*

For a vocabulary τ and a binary relation symbol R not in τ we denote by τ_R the vocabulary $\tau \cup \{R\}$. In particular, we deal with the binary relation symbols $<$ (for an ordering) and S (for a successor relation). Sometimes (but not always) it will be convenient to write τ_R-structures in the form $(\mathcal{A}, R^{\mathcal{A}})$.

For an ordered $\tau_<$-structure $(\mathcal{A}, <^{\mathcal{A}})$ we denote by $S(<^{\mathcal{A}})$ the successor relation of $<^{\mathcal{A}}$,

$$S(<^{\mathcal{A}}) := \big\{(a,a') \bigm| a, a' \in A,\ a <^{\mathcal{A}} a', \text{ and } a' \leq^{\mathcal{A}} b \text{ for all } b \text{ with } a <^{\mathcal{A}} b\big\}.$$

The τ_S-structure $(\mathcal{A}, S(<^{\mathcal{A}}))$ is the *successor structure associated with* $(\mathcal{A}, <^{\mathcal{A}})$. If O is a class of ordered $\tau_<$-structures, we let

$$\mathrm{S(O)} := \big\{(\mathcal{A}, S(<^{\mathcal{A}})) \bigm| (\mathcal{A}, <^{\mathcal{A}}) \in \mathrm{O}\big\}$$

be the class of its associated successor structures.

We assume that the reader is familiar with the notions of tree decomposition and treewidth of structures. We write tree decompositions in the form $(\mathcal{T}, (B_t)_{t \in T})$, where $\mathcal{T}$ is the corresponding tree and B_t the *bag* at t. For a structure $\mathcal{A}$ we denote by $\mathrm{tw}(\mathcal{A})$ the treewidth of $\mathcal{A}$. Similarly, tw(C) denotes the treewidth of the class C of structures.

Logics. We assume familiarity with first-order logic FO, monadic second-order logic MSO, least fixed-point logic LFP, deterministic transitive closure logic DTC, and with partial fixed-point logic PFP. For every of these logics L we denote by $L[\tau]$ the set of L-formulas of vocabulary τ. Concerning LFP, DTC, and PFP essentially we only need the following property.

Theorem II.1 ([16], [17], [25], [1])**.** *On ordered structures,* LFP. DTC*, and* PFP *capture the complexity classes* P, L *(logarithmic space), and* PSPACE*, respectively.*

Let L and L' be logics and C a class of structures. An L-formula $\varphi(x_1, \dots, x_n)$ with free individual variables among $x_1, \dots, x_n$ and an L'-formula $\psi(x_1, \dots, x_n)$ *are equivalent on* C if for all $\mathcal{A} \in$ C and all $a_1, \dots, a_n \in A$,

$$\mathcal{A} \models \varphi(a_1, \dots, a_n) \iff \mathcal{A} \models \psi(a_1, \dots, a_n).$$

The logic L' *is at least as expressive as* L, written $L' \geq L$ on C, if every L-formula is equivalent to an L'-formula on C. If, on C, $L' \geq L$ but not $L \geq L'$, then L' *is more expressive than* L on C, written $L' > L$ on C, The logics L and L' *have the same expressive power on* C, written $L \equiv L'$ on C, if $L \geq L'$ on C and $L' \geq L$ on C.

Clearly, $L \equiv L'$ on C for every finite class C of structures and any logics L and L' with $L \geq$ FO and $L' \geq$ FO on C. We already mentioned in the Introduction:

Theorem II.2 ([10])**.** *If* C *is an infinite class of ordered structures, then* PFP $>$ FO *on* C*.*

By taking an LFP-formula $\varphi(x)$ of vocabulary $\tau := \{<\}$ expressing that x is the nth element of the ordering $<$ for some even n we see:

Proposition II.3. *If* C *is an infinite class of ordered* $\{<\}$*-structures, then* LFP $>$ FO *on* C*.*

The following result shows that we get a quite different notion if we restrict ourselves to sentences in the definition of $L \equiv L'$ on C. We present a proof in Appendix A.

Proposition II.4 ([10])**.** *Let* L *be a logic with* $L \geq$ FO*. Assume that for every vocabulary* τ *the set of* L*-sentences of vocabulary* τ *is countable. Then every infinite class* C *of structures contains an infinite subclass* C$'$ *such that every* L*-sentence is equivalent to an* FO*-sentence on* C$'$ *("*$L \equiv$ FO *for sentences on* C$'$*").*

III. The model-checking method

Central to the model-checking method is the observation that for every infinite class O of ordered structures and every natural number d there is an LFP-formula whose validity in structures $\mathcal{A} \in$ O may not be checked in time $O(|\mathcal{A}|^d)$. Thus, for suitable d, such a formula can't be equivalent to any first-order formula if on O the FO-model-checking problem parameterized by the length of the formula is fixed-parameter tractable.

A standard diagonalization argument yields the observation on LFP just mentioned:

Proposition III.1. *Let* τ *be a vocabulary containing the relation symbol* $<$ *and let* $d \geq 1$ *be a natural number. There is an* LFP*-formula* $\varphi(x)$ *such that for every infinite class* O *of ordered* τ*-structures the property*

$$\mathcal{A} \models \varphi(a)$$

cannot be decided in time $O(|\mathcal{A}|^d)$ *for every* $\mathcal{A} \in$ O *and* $a \in A$*.*

Proof: We fix an effective enumeration $\mathbb{M}_1, \mathbb{M}_2, \dots$ such that

(E1) each $\mathbb{M}_i$ is a Turing machine which on inputs of the form $(\mathcal{A}, a)$, where $\mathcal{A}$ is an *ordered* τ-structure and $a \in A$, runs in time $|(\mathcal{A}, a)|^{d+1}$;
(E2) there is an algorithm that on input $(i, \mathcal{A}, a)$ simulates $\mathbb{M}_i$ on $(\mathcal{A}, a)$ in time $i \cdot |(\mathcal{A}, a)|^{d+1}$;
(E3) for every Turing machine $\mathbb{M}$ which on inputs of the form $(\mathcal{A}, a)$ runs in time $|(\mathcal{A}, a)|^{d+1}$ there is an $i \geq 1$ such that $\mathbb{M}$ and $\mathbb{M}_i$ accept the same pairs $(\mathcal{A}, a)$.

Let the property $\mathscr{P}$ hold for an ordered τ-structure $\mathcal{A}$ and $a \in A$, written $\mathscr{P}(\mathcal{A}, a)$, if the a-th machine $\mathbb{M}_a$ does not accept $(\mathcal{A}, a)$ (as $\mathcal{A}$ is ordered, each element $a \in A$ can be identified with a natural number). By (E2) it is easy to see that $\mathscr{P}(\mathcal{A}, a)$ can be decided in time $a \cdot |(\mathcal{A}, a)|^{d+1}$, and thus, in polynomial time. Hence, by Theorem II.1, there is an LFP-formula $\varphi(x)$ such that

$$\mathcal{A} \models \varphi(a) \iff \mathscr{P}(\mathcal{A}, a).$$

We show that $\varphi(x)$ satisfies the claim of the proposition. Assume otherwise, then there is an infinite class O, a $c \in \mathbb{N}$, and a Turing machine $\mathbb{M}$ such that $\mathbb{M}$ decides $\mathscr{P}(\mathcal{A}, a)$ in time $c \cdot |\mathcal{A}|^d$ for every $\mathcal{A} \in$ O and $a \in A$. Clearly, for sufficiently large $\mathcal{A}$, this term is dominated by $|\mathcal{A}|^{d+1}$ and hence, by $|(\mathcal{A}, a)|^{d+1}$. Thus, by (E3) there exists an $i \in \mathbb{N}$ such that

$$\mathbb{M}_i \text{ accepts } (\mathcal{A}, a) \iff \mathbb{M} \text{ accepts } (\mathcal{A}, a)$$

for all sufficiently large $\mathcal{A} \in \mathrm{O}$ and $a \in A$.

Choose such a sufficiently large $\mathcal{A} \in \mathrm{O}$ and let a be the ith element in $\mathcal{A}$. Then we conclude

$$\begin{aligned} \mathscr{P}(\mathcal{A}, a) &\iff \mathbb{M}_i \text{ does not accept } (\mathcal{A}, a) \\ &\iff \mathbb{M} \text{ does not accept } (\mathcal{A}, a) \\ &\iff \mathscr{P}(\mathcal{A}, a) \text{ does not hold,} \end{aligned}$$

a contradiction. □

For a class C of structures and a logic L the *(parameterized) L-model-checking problem is fixed-parameter tractable*, written $p\text{-MC}(\mathrm{C}, L) \in$ FPT, if there is a $d \geq 1$, a computable function $f : \mathbb{N} \to \mathbb{N}$, and an algorithm which decides whether $\mathcal{A} \models \varphi(a_1, \ldots, a_n)$ in time $f(\varphi) \cdot |\mathcal{A}|^d$ for every L-formula $\varphi(x_1, \ldots, x_n)$, all $\mathcal{A} \in \mathrm{O}$, and $a_1, \ldots, a_n \in A$.

Choosing for such a $d \geq 1$ the formula φ according to the previous proposition, we see:

Corollary III.2. *If* $p\text{-MC}(\mathrm{O}, \mathrm{FO}) \in$ FPT, *then* LFP $>$ FO *on* O.

Corollary III.3. *If* $p\text{-MC}(\mathrm{S(O)}, \mathrm{MSO}) \in$ FPT, *then* LFP $>$ FO *on* O.

Proof: As in MSO the ordering may be recovered from the successor relation, we see that $p\text{-MC}(\mathrm{S(O)}, \mathrm{MSO}) \in$ FPT implies $p\text{-MC}(\mathrm{O}, \mathrm{MSO}) \in$ FPT. Thus, the previous corollary yields the claim. □

By Courcelle's Theorem, $p\text{-MC}(\mathrm{C}, \mathrm{MSO}) \in$ FPT holds for every class C of bounded treewidth. Therefore, we get (BOC), the bounded ordered conjecture, from the last corollary.

A further application of Proposition III.1 yields:

Theorem III.4. *If* O *or* S(O) *is of bounded cliquewidth, then* LFP $>$ FO *on* O.

Hereby, we only consider vocabularies τ of arity ≤ 2, that is, every relation symbol in τ has arity ≤ 2. The notion of cliquewidth introduced for graphs in [8] has been extended in [15] to arbitrary structures. The idea of the proof of Theorem III.4 is the following (in Appendix B we present the relevant definitions and the details of the proof): Assume that the class O of structures in a vocabulary τ of arity ≤ 2 is of bounded cliquewidth. In polynomial time we assign to every τ-structure $\mathcal{A}$ a graph $\mathcal{G}_\mathcal{A}$ such that the class $\mathrm{C_O} := \{\mathcal{G}_\mathcal{A} \mid \mathcal{A} \in \mathrm{O}\}$ is of bounded cliquewidth, too, and $p\text{-MC}(\mathrm{O}, \mathrm{MSO})$ is fpt-reducible to $p\text{-MC}(\mathrm{C_O}, \mathrm{MSO})$. By results of [9] and [23], we know that $p\text{-MC}(\mathrm{C_O}, \mathrm{MSO}) \in$ FPT, thus $p\text{-MC}(\mathrm{O}, \mathrm{MSO}) \in$ FPT. The claim follows by Corollary III.2.

IV. The padding method

If $\mathcal{A}$ and $\mathcal{B}$ are τ-structures (where τ does not contain constants), the *union* $\mathcal{A} \cup \mathcal{B}$ is the structure with universe $A \cup B$ and with $R^{\mathcal{A} \cup \mathcal{B}} := R^\mathcal{A} \cup R^\mathcal{B}$. If $< \in \tau$ and $\mathcal{A}$ and $\mathcal{B}$ are ordered τ-structures with $A \cap B = \emptyset$, we obtain the *ordered sum* $\mathcal{A} \oplus \mathcal{B}$ from $\mathcal{A} \cup \mathcal{B}$ by changing the interpretation of $<$ to

$$<^{\mathcal{A} \oplus \mathcal{B}} := <^{\mathcal{A} \cup \mathcal{B}} \cup \big\{(a, b) \bigm| a \in A \text{ and } b \in B\big\}.$$

If $A \cap B \neq \emptyset$, one defines $\mathcal{A} \oplus \mathcal{B}$ by passing, in a canonical way, to isomorphic structures with disjoint universes and taking their ordered sum.

Theorem IV.1. *Let* O *and* O′ *be infinite classes of ordered structures of vocabulary* τ. *Then,* DTC $>$ FO *on* $\mathrm{O} \oplus \mathrm{O}' := \big\{\mathcal{A} \oplus \mathcal{A}' \bigm| \mathcal{A} \in \mathrm{O} \text{ and } \mathcal{A}' \in \mathrm{O}'\big\}$.

We prove this result by the padding method. We first explain its basic idea. The validity of a PFP-formula $\varphi(\bar{x})$ in a structure $\mathcal{A}$ of O may be checked in the "padded" structure $\mathcal{A} \oplus \mathcal{A}'$ in logarithmic space if $\mathcal{A}' \in \mathrm{O}'$ is sufficiently large (and we have a constant for the last element of $\mathcal{A}$). As DTC captures L on ordered structures, the validity of $\varphi(\bar{x})$ in $\mathcal{A}$ may be expressed in $\mathcal{A} \oplus \mathcal{A}'$ by a DTC-formula $\psi(\bar{x})$. If, by contradiction, we assume that DTC $\equiv$ FO on $\mathrm{O} \oplus \mathrm{O}'$, we may replace $\psi(\bar{x})$ by an FO-formula $\xi(\bar{x})$. There is a suitable infinite subclass of $\mathrm{O} \oplus \mathrm{O}'$ such that for all (but finitely many) structures $\mathcal{A} \oplus \mathcal{A}'$ and $\mathcal{B} \oplus \mathcal{B}'$ in this subclass the structures $\mathcal{A}'$ and $\mathcal{B}'$ satisfy the same first-order sentences of sufficiently large quantifier rank (namely, all those of quantifier rank at most that of $\xi(\bar{x})$). Then, we may replace the FO-formula $\xi(\bar{x})$ (evaluated in structures of $\mathrm{O} \oplus \mathrm{O}'$) by an FO-formula $\rho(\bar{x})$ evaluated in structures of O. Altogether, then $\varphi(\bar{x})$ is equivalent to $\rho(\bar{x})$ on a suitable infinite subclass of O, which contradicts Theorem II.2.

Before we turn to the details we need some tools related to the Ehrenfeucht-Fraïssé method. If $\mathcal{A}$ and $\mathcal{B}$ are τ-structures and $i \in \mathbb{N}$, we write $\mathcal{A} \equiv^i \mathcal{B}$ if $\mathcal{A}$ and $\mathcal{B}$ satisfy the same FO-sentences of quantifier rank $\leq i$. Given $\mathcal{A}$ and a sequence $\bar{a} = a_1, \ldots, a_\ell$ of elements of A and $i \in \mathbb{N}$ there is a formula, the *i-type of $\bar{a}$ in $\mathcal{A}$*, we denote by $\mathrm{tp}^i_{\bar{a}}(\mathcal{A})(x_1, \ldots, x_\ell)$ (and by $\mathrm{tp}^i(\mathcal{A})$ if $\ell = 0$), such that for all $\mathcal{B}$ and $\bar{b} = b_1, \ldots, b_\ell$ in B:

$$\begin{aligned} &\mathcal{B} \models \mathrm{tp}^i_{\bar{a}}(\mathcal{A})(\bar{b}) \\ &\iff \bar{a} \text{ in } \mathcal{A} \text{ and } \bar{b} \text{ in } \mathcal{B} \text{ satisfy the same formulas of quantifier rank} \leq i \\ &\iff (\mathcal{A}, \bar{a}) \equiv^i (\mathcal{B}, \bar{b}) \\ &\iff \mathrm{tp}^i_{\bar{a}}(\mathcal{A})(x_1, \ldots, x_\ell) = \mathrm{tp}^i_{\bar{b}}(\mathcal{B})(x_1, \ldots, x_\ell). \end{aligned}$$

Here $(\mathcal{A},\bar{a})$ and $(\mathcal{B},\bar{b})$ are structures in an enlarged vocabulary containing constants. Occasionally (and then we mention it explicitly), we consider such vocabularies containing finitely many constants (more precisely, constant symbols).

The following well-known result may be obtained by the Ehrenfeucht-Fraïssé method:

Proposition IV.2. *Let $\mathcal{A}$, $\mathcal{A}'$, $\mathcal{B}$, and $\mathcal{B}'$ be ordered τ-structures and $\bar{a} = a_1,\dots,a_\ell$ and $\bar{a}' = a'_1,\dots,a'_\ell$ be sequences of elements of A and A', respectively. Let $i \in \mathbb{N}$. If $(\mathcal{A},\bar{a}) \equiv^i (\mathcal{A}',\bar{a}')$ and $\mathcal{B} \equiv^i \mathcal{B}'$, then $(\mathcal{A} \oplus \mathcal{B},\bar{a}) \equiv^i (\mathcal{A}' \oplus \mathcal{B}',\bar{a}')$.*

Proof of Theorem IV.1: By contradiction, assume that DTC $\equiv$ FO on $\mathrm{O} \oplus \mathrm{O}'$. We show that there is an infinite class $\mathrm{C} \subseteq \mathrm{O}$ such that PFP $\equiv$ FO on C. However, this contradicts Theorem II.2.

For this purpose we choose, for every structure $\mathcal{A} \in \mathrm{O}$, a structure $\mathcal{A}^+ \in \mathrm{O}'$ such that

$$2^{2^{|A|}} \le |A| + |A^+|. \tag{2}$$

We define a sequence $(\mathrm{C}_i)_{i\ge 0}$ of infinite classes C_i with

$$\mathrm{O} = \mathrm{C}_0 \supseteq \mathrm{C}_1 \supseteq \cdots \supseteq \mathrm{C}_i \supseteq \cdots \tag{3}$$

inductively. Assume we already have defined the infinite subclass C_i of O. Then, by (2), the class $\{\mathcal{A}^+ \in \mathrm{O}' \mid \mathcal{A} \in \mathrm{C}_i\}$ is infinite. So there is at least one first-order i-type θ_i such that infinitely many structures in this class are of i-type θ_i. We let C_{i+1} be the class

$$\{\mathcal{A} \in \mathrm{C}_i \mid \mathrm{tp}^i(\mathcal{A}^+) = \theta_i\} \cup \{\mathcal{B} \mid \mathcal{B} \text{ has a } (\le i)\text{th smallest size in } \mathrm{C}_i\}.^2$$

Then C_{i+1} is infinite and $\mathrm{C}_i \supseteq \mathrm{C}_{i+1}$. Finally, we set $\mathrm{C} := \bigcap_{i\ge 0} \mathrm{C}_i$.

One easily verifies Claim 1.

Claim 1. The class C is infinite and for every $i \ge 0$, $\mathrm{tp}^i(\mathcal{A}^+) = \theta_i$ for all but finitely many structures $\mathcal{A} \in \mathrm{C}$.

Claim 2. For every PFP-formula $\varphi(\bar{x})$ there is an $i \in \mathbb{N}$ and an FO-formula $\rho(\bar{x})$ such that for every $\mathcal{A} \in \mathrm{C}$ with $\mathrm{tp}^i(\mathcal{A}^+) = \theta_i$ and $\bar{a} \in A^{|\bar{x}|}$ we have

$$\mathcal{A} \models \varphi(\bar{a}) \iff \mathcal{A} \models \rho(\bar{a}).$$

Proof of Claim 2. Let $\varphi(\bar{x})$ be a PFP-formula and, for notational simplicity, assume that $\bar{x} = x_1, x_2$. We consider the following property of arbitrary ordered τ-structures $\mathcal{A}$ and $a, a_1, a_2 \in A$:

$$a_1, a_2 \le^{\mathcal{A}} a, \quad 2^{2^{n(a)}} \le |A|, \quad \text{and } \mathcal{A}{\restriction}a \models \varphi(a_1,a_2). \tag{4}$$

Here, $\mathcal{A}{\restriction}a$ denotes the substructure of $\mathcal{A}$ induced on the set $\{b \in A \mid b \le^{\mathcal{A}} a\}$ and $n(a) := \big|\{b \in A \mid b \le^{\mathcal{A}} a\}\big|$, that is, a is the $n(a)$th element of the ordering in $\mathcal{A}$.

Clearly, the property (4) of ordered τ-structures $\mathcal{A}$ and of $a, a_1, a_2 \in A$ can be checked in logarithmic space. As DTC captures logarithmic space on ordered structures (cf. Theorem II.1), there is a DTC-formula $\psi(x, x_1, x_2)$ such that for all ordered τ-structures $\mathcal{A}$ and $a, a_1, a_2 \in A$:

$$\mathcal{A} \models \psi(a, a_1, a_2) \iff \mathcal{A} \text{ and } a, a_1, a_2 \text{ satisfy (4)}. \tag{5}$$

By our assumption DTC $\equiv$ FO on $\mathrm{O} \oplus \mathrm{O}'$, there is an FO-formula $\xi(x, x_1, x_2)$ such that for every $\mathcal{A} \in \mathrm{O}$, $\mathcal{A}' \in \mathrm{O}'$, and $a, a_1, a_2 \in A \cup A'$ we have

$$\mathcal{A} \oplus \mathcal{A}' \models \psi(a, a_1, a_2) \iff \mathcal{A} \oplus \mathcal{A}' \models \xi(a, a_1, a_2). \tag{6}$$

Let i be the quantifier rank of ξ. We let $\chi(x, x_1, x_2)$ be the FO-formula (recall that θ_i is the i-type chosen to define C_{i+1})

$$\bigvee \{\mathrm{tp}^i_{b,b_1,b_2}(\mathcal{B})(x,x_1,x_2) \mid \mathcal{B} \in \mathrm{C},\ b, b_1, b_2 \in B, \mathrm{tp}^i(\mathcal{B}^+) = \theta_i \text{ and } \mathcal{B} \oplus \mathcal{B}^+ \models \xi(b, b_1, b_2)\}.$$

Then, for every $\mathcal{A} \in \mathrm{C}$ with $\mathrm{tp}^i(\mathcal{A}^+) = \theta_i$ and all $a, a_1, a_2 \in A$, we have

$$\mathcal{A} \models \chi(a, a_1, a_2) \iff \mathcal{A} \oplus \mathcal{A}^+ \models \xi(a, a_1, a_2). \tag{7}$$

The implication from right to left is easy: If $\mathcal{A} \oplus \mathcal{A}^+ \models \xi(a, a_1, a_2)$, then $\mathrm{tp}^i_{a,a_1,a_2}(\mathcal{A})(x, x_1, x_2)$ is a disjunct of χ; thus, $\mathcal{A} \models \chi(a, a_1, a_2)$. Conversely, assume that $\mathcal{A} \models \chi(a, a_1, a_2)$. Then, for some $\mathcal{B} \in \mathrm{C}$ and $b, b_1, b_2 \in B$ with $\mathrm{tp}^i(\mathcal{B}^+) = \theta_i$ and $\mathcal{B} \oplus \mathcal{B}^+ \models \xi(b, b_1, b_2)$, we have $\mathcal{A} \models \mathrm{tp}^i_{b,b_1,b_2}(\mathcal{B})(a, a_1, a_2)$. Hence,

$$(\mathcal{A}, a, a_1, a_2) \equiv^i (\mathcal{B}, b, b_1, b_2) \quad \text{and} \quad \mathcal{A}^+ \equiv^i \mathcal{B}^+.$$

Thus, $(\mathcal{A} \oplus \mathcal{A}^+, a, a_1, a_2) \equiv^i (\mathcal{B} \oplus \mathcal{B}^+, b, b_1, b_2)$ by Proposition IV.2; in particular, $\mathcal{A} \oplus \mathcal{A}^+ \models \xi(a, a_1, a_2)$.

Finally, we let $\rho(x_1, x_2)$ be the FO-formula

$$\exists x\big(\chi(x, x_1, x_2) \land \text{“}x \text{ is the last element of } <\text{”}\big).$$

Let $\mathcal{A} \in \mathrm{C}$ with last element a and with $\mathrm{tp}^i(\mathcal{A}^+) = \theta_i$. Then, by (2),

$$2^{2^{n(a)}} = 2^{2^{|A|}} \le |A| + |A^+| \quad \text{and } (\mathcal{A} \oplus \mathcal{A}^+){\restriction}a = \mathcal{A}. \tag{8}$$

[2] The structure $\mathcal{B}$ has a $(\le i)$th smallest size in a class D of structures if $\mathcal{B} \in \mathrm{D}$ and there is a $j \le i$ and a sequence $\mathcal{B}_1, \dots, \mathcal{B}_{j-1}$ of structures in D such that $|B_1| < |B_2| < \dots |B_{j-1}| < |B|$, and there is no such sequence for $j = i+1$.

Therefore, for all $a_1, a_2 \in A$,

$$\begin{aligned} & \mathcal{A} \models \varphi(a_1, a_2) \\ \iff & \mathcal{A} \oplus \mathcal{A}^+ \models \psi(a, a_1 a_2) \text{ (by (4), (5), (8))} \\ \iff & \mathcal{A} \oplus \mathcal{A}^+ \models \xi(a, a_1 a_2) \text{ (by (6))} \\ \iff & \mathcal{A} \models \chi(a, a_1 a_2) \text{ (by (7))} \\ \iff & \mathcal{A} \models \rho(a_1, a_2) \text{ (by definition of } \rho). \end{aligned}$$

This finishes the proof of Claim 2. ⊣

Now it is easy to show that PFP $\equiv$ FO on C. Given a PFP-formula $\varphi(\bar{x})$ we determine $i \geq 0$ and an FO-formula $\rho(\bar{x})$ satisfying Claim 2. By Claim 1, we know that $\varphi(\bar{x})$ and $\rho(\bar{x})$ are equivalent in all but finitely many structures of C. Taking care of these finitely many structures separately, we get an FO-formula equivalent to $\varphi(\bar{x})$ in C. □

By the padding method we obtain the following strengthening of (BOC).

Theorem IV.3. *If* $\mathrm{tw}(\mathrm{S}(\mathrm{O})) < \infty$, *then* DTC $>$ FO *on* O.

We give the technical details of a proof in Appendix C; here we only mention the main idea. To apply the padding method we have to write the structures $\mathcal{A}$ in O as an "ordered union" of a small and a large substructure. This is achieved along a suitable bag of a tree-decomposition of $\mathcal{A}$. Using the algorithm of Elberfeld et al. [12] we get such a tree-decomposition in logarithmic space.

In view of [12] the question arises whether Theorem IV.3 can be obtained by the simpler model-checking method. Of course, there is a DTC-analogue of Proposition III.1 showing that under the hypotheses of this proposition there is a DTC-formula $\varphi(x)$ such that $\mathcal{A} \models \varphi(a)$ cannot be solved in space $O(1) + d \cdot \log|\mathcal{A}|$ for structures $\mathcal{A}$ ranging over a given O. The space version of Courcelle's Theorem obtained in [12] shows p-MC(C, MSO) $\in$ XL. The authors of [12] communicated to us that apparently their proof technique seems not to yield p-MC(C, MSO) $\in$ para-L, which would be needed to obtain the result by the model-checking method.

V. Further (non-) collapse results

Recently [5], [6], when analyzing the complexity of some parameterized problems, we introduced a complexity-theoretic assumption, namely NP[TC] $\neq$ P[TC]. Its failure implies that

there exists an infinite set $I \subseteq \{0,1\}^*$ *such that* $Q \cap I \in$ P *for all* $Q \subseteq \{0,1\}^*$ *in* NP. (9)

Hence, "NP collapses to P on I." As, by Fagin's Theorem, the fragment Σ_1^1 of second-order logic captures the complexity class NP, we thus hoped to show that $\Sigma_1^1 \equiv$ LFP on the class of structures representing strings of such an $I \subseteq \{0,1\}^*$. However, we realized that it only yields the collapse for *sentences* which, by Proposition II.4, is true without any assumption. In this section we relate the "function version" of (9) with a collapse result.

Let Σ^* denote the set of strings over the alphabet $\Sigma := \{0,1\}$.

Definition V.1. Let $h : \mathbb{N} \to \mathbb{N}$ and $f : \Sigma^* \to \Sigma^*$ be functions.

(i) The function f *is computable in time polynomial in* h or, in short, $f \in$ FP(h), if there is a deterministic Turing machine $\mathbb{M}$ such that for every input $w \in \Sigma^*$ the machine $\mathbb{M}$ computes $f(w)$ in time $h(|w|)^{O(1)}$.

(ii) The function f *is nondeterministically computable in time polynomial in* h or, in short, $f \in$ FNP(h) if there is a nondeterministic Turing machine $\mathbb{M}$ such that for every input $w \in \Sigma^*$ and $i \in \big[|f(w)|\big]$ the ith bit of $f(w)$ is 1 if and only if $\mathbb{M}$ accepts (w, i) in time $h(|w|)^{O(1)}$, and for all $i > |f(w)|$ the machine $\mathbb{M}$ always outputs a special symbol $\perp$ in time $h(|w|)^{O(1)}$.

For time constructible h one easily verifies that $f \in$ FP(h) if and only if for f there is a *deterministic* machine $\mathbb{M}$ satisfying the conditions in (ii).

For a function $f : \Sigma^* \to \Sigma^*$ and a set $I \subseteq \{0,1\}^*$ let $f\restriction_I : \Sigma^* \to \Sigma^*$ be defined by

$$f\restriction_I(w) = \begin{cases} f(w), & \text{if } w \in I \\ \lambda \ (= \text{the empty string}), & \text{otherwise.} \end{cases}$$

By $id_{\mathbb{N}}$ we denote the identity function on $\mathbb{N}$. The function version (FV) of (9) is the statement:

There is an infinite set $I \subseteq \Sigma^*$ *such that* $f\restriction_I \in$ FP[$id_{\mathbb{N}}$] *for every function* $f \in$ FNP[$id_{\mathbb{N}}$].

Clearly, (FV) implies (9). By taking as f the identity function $id_{\mathbb{N}}$ we see that any I satisfying (FV) is decidable in polynomial time.

Let τ be the vocabulary $\{O, <\}$ with unary O and binary $<$. Every string $w = a_1 \cdots a_m \in \Sigma^*$ may be represented by an ordered τ-structure $\mathcal{S}(w)$, where the universe is $[m]$, where $<^{\mathcal{S}(w)}$ is the natural ordering on $[m]$, and where $O^{\mathcal{S}(w)} := \{i \in [m] \mid a_i = 1\}$ is the set of positions in w carrying a one. If $I \subseteq \Sigma^*$, then we set $\mathrm{O}(I) := \{\mathcal{S}(w) \mid w \in I\}$.

Proposition V.2. *If* (FV) *holds and* I *is according to* (FV), *then* $\Sigma_1^1 \equiv$ LFP *on* O(I).

Proof: Let I be as stated. Let $\varphi(\bar{x})$ be any Σ_1^1-formula. For simplicity, assume that $\bar{x} = x, y$.

For every $w \in \Sigma^*$, say, of length m and every $i, j \in [m]$ we let

$$bit\big(\mathcal{S}(w) \models \varphi(i,j)\big) := \begin{cases} 1, & \text{if } \mathcal{S}(w) \models \varphi(i,j) \\ 0, & \text{if } \mathcal{S}(w) \not\models \varphi(i,j) \end{cases}$$

be the "truth value" of $\mathcal{S}(w) \models \varphi(i,j)$. Furthermore, we define $f_\varphi : \Sigma^* \to \Sigma^*$ by letting $f_\varphi(w)$ be the string consisting of the bits corresponding to the truth values of $\mathcal{S}(w) \models \varphi(i,j)$, where the pairs $(i,j) \in [m] \times [m]$ are taken in the lexicographic order.

As the evaluation of φ can be performed by a nondeterministic machine in polynomial time, we have $f_\varphi \in$ FNP$[id_{\mathbb{N}}]$). By the assumption (FV) there is a polynomial time deterministic Turing machine that computes f_φ on I. As LFP captures PTIME on ordered structures, one easily gets an LFP-formula $\rho(x,y)$ such that for all $w \in I$ and $i, j \in [|w|]$ we have

$$\begin{aligned} \mathcal{S}(w) \models \rho(i,j) \quad &\iff \quad bit\big(\mathcal{S}(w) \models \varphi(i,j)\big) = 1 \\ &\iff \quad \mathcal{S}(w) \models \varphi(i,j), \end{aligned}$$

and thus, $\rho(x,y)$ and $\varphi(x,y)$ are equivalent on O(I). □

As all structures in S(O(I)) have treewidth 1, by the previous proposition together with (BOC) we get:

Corollary V.3. *If* (FV) *holds and I is according to* (FV)*, then* $\Sigma^1_1 \equiv$ LFP $>$ FO *on* O(I).

We present an assumption which yields the validity of (FV).

Lemma V.4. *Let $h : \mathbb{N} \to \mathbb{N}$ be a time constructible and increasing function. If* FP$[h]$ = FNP$[h]$*, then* (FV) *holds.*

Proof: For $f \in$ FNP$[id_{\mathbb{N}}]$ let $f' : \Sigma^* \to \Sigma^*$ be defined by

$$f'(x) = \begin{cases} f(1^{h(m)}), & \text{if } x = 1^m \\ \lambda, & \text{otherwise.} \end{cases}$$

As h is time constructible and increasing, f' is well-defined and $f' \in$ FNP$[h]$. Hence, by assumption, $f' \in$ FP$[h]$. Thus we see that we can set $I := \{1^{h(m)} \mid m \in \mathbb{N}\}$. □

In analogy to NP[TC] $\neq$ P[TC] we let FP[TC] $\neq$ FNP[TC] mean that FP$[h] \neq$ FNP$[h]$ for every time constructible and increasing function $h : \mathbb{N} \to \mathbb{N}$. Putting Proposition V.2 and Lemma V.4 together we obtain:

Theorem V.5. *If* $\Sigma^1_1 \neq$ LFP *on every infinite class of structures, then* FP[TC] $\neq$ FNP[TC].

VI. On bounded treewidth preserving successor relations

First we characterize those classes C such that for the class O(C) := $\{(\mathcal{A}, <^{\mathcal{A}}) \mid \mathcal{A} \in$ C and $<^{\mathcal{A}}$ ordering of $A\}$ of all ordered versions of structures in C, we have tw(S(O(C))) $< \infty$ (and thus DTC $>$ FO on O(C) by Theorem IV.3).

Example VI.1. Let τ be a vocabulary containing only unary relation symbols and C be the class of all τ-structures. Then O(C) is the "class of paths coloured by τ." Hence, tw(S(O(C))) $= 1$.

Recall that the *Gaifman graph* $\mathcal{G}(\mathcal{A})$ of a τ-structure $\mathcal{A}$ is the graph $\mathcal{G}(\mathcal{A}) = (G(A), E^{\mathcal{G}(\mathcal{A})})$ with $G(A) := A$ and where $E^{\mathcal{G}(\mathcal{A})}$ is the set

$$\big\{(a,b) \bigm| a, b \in A,\ a \neq b, \text{there exists } R \in \tau \text{ and } (a_1, \ldots, a_r) \in R^{\mathcal{G}} \text{ with } a, b \in \{a_1, \ldots, a_r\}\big\}.$$

Example VI.2. Let $\tau = \{E\}$ with binary E and let C be the class of paths (that is, the class of graphs that are paths). Then tw(S(O(C))) $= \infty$. In fact, it is easy to see that on a path $\mathcal{P} = (P, E^{\mathcal{P}})$ of length n^2 an ordering $<^{\mathcal{P}}$ may be chosen in such a way that the Gaifman graph of $(P, E^{\mathcal{P}}, <^{\mathcal{P}})$ is an $n \times n$-grid. Hence, tw$((P, E^{\mathcal{P}}, <^{\mathcal{P}})) = n$.

The previous examples are contained in the main result of this section we address now.

Recall that a *matching* in a graph is a set of pairwise disjoint edges (that is, edges having no vertices in common). The *matching number* $m(\mathcal{G})$ of the graph $\mathcal{G}$ is the maximum of the cardinalities of matchings in $\mathcal{G}$,

$$m(\mathcal{G}) := \max\big\{|M| \bigm| M \text{ is a matching in } \mathcal{G}\big\}.$$

For a structure $\mathcal{A}$ we let $m(\mathcal{A})$, the *matching number of* $\mathcal{A}$, be the matching number of the Gaifman graph $\mathcal{G}(\mathcal{A})$ of $\mathcal{A}$. Finally, for a class C the *matching number* m(C) *of* C is given by

$$m(\mathrm{C}) := \sup\big\{m(\mathcal{A}) \bigm| \mathcal{A} \in \mathrm{C}\big\}.$$

Theorem VI.3. *For every class* C *of structures*

$$\text{tw(S(O(C)))} < \infty \iff m(\mathrm{C}) < \infty,$$

that is, the class of successor structures associated with the class of all ordered versions in C *has bounded treewidth if and only if the matching number of* C *is finite.*

Proof: First we show that for a graph $\mathcal{G}$ we get an ordering $<^{\mathcal{G}}$ with tw$\big((\mathcal{G}, S(<^{\mathcal{G}}))\big) \geq n$ for a given $n \in \mathbb{N}$ if the matching number $m(\mathcal{G})$ is sufficiently large. This is exemplified in Figure 1,

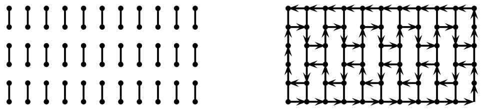

Fig. 1.

where the left hand side shows a matching in a graph $\mathcal{G}$, and in the right hand side the arrows represent the successor relation of an ordering. In this way we get, as minor of the Gaifman graph of $\big(\mathcal{G}, S(<^{\mathcal{G}})\big)$, a large grid. Hence, we see that $m(\mathrm{C}) = \infty$ implies $\mathrm{tw}(\mathrm{S}(\mathrm{O}(\mathrm{C}))) = \infty$.

To show that $\mathrm{tw}(\mathrm{S}(\mathrm{O}(\mathrm{C}))) < \infty$ if $m(\mathrm{C}) < \infty$ it suffices to prove the following claim.

Claim: For every graph $\mathcal{G}$ and any ordering $<^{\mathcal{G}}$ on $\mathcal{G}$ we have $\mathrm{tw}\big((\mathcal{G}, S(<^{\mathcal{G}}))\big) \leq 2 \cdot m(\mathcal{G}) + 1$.

Proof of the Claim: Let M be a maximum matching of $\mathcal{G}$. Then:

– For the set $V(M)$ of vertices of edges in M we have $|V(M)| = 2 \cdot m(\mathcal{G})$.
– Every edge of $\mathcal{G}$ has at least one endpoint in M.

For an arbitrary ordering $<^{\mathcal{G}}$ on $\mathcal{G}$ let $a_1, \ldots, a_s$ be the enumeration of the elements of $G \setminus V(M)$ according to this ordering. Then for $i, j \in [s]$ we have

$$\text{if } (a_i, a_j) \in S(<^{\mathcal{G}}), \text{ then } j = i + 1 \qquad (10)$$

(but, in general, the converse will not be true). We present a tree decomposition $\big(T, (B_t)_{t \in T}\big)$ of $(\mathcal{G}, S(<^{\mathcal{G}}))$ of width $\leq 2 \cdot m(\mathcal{G}) + 1$. The tree $\mathcal{T}$ is a path with nodes $r, t_1, \ldots, t_s$, where r is the root and t_{i+1} the child of t_i for $i \in [s-1]$. The bags B_t are given by $B_r := V(M)$, $B_{t_1} := V(M) \cup \{a_1\}$, and for i with $2 \leq i \leq s$ by $B_{t_i} := V(M) \cup \{a_{i-1}, a_i\}$. □

As the tree of the last tree decomposition is a path, we see that in the previous theorem we may replace the treewidth by the pathwidth.

So far we dealt with *all* ordered versions of a class of structures of bounded treewidth. Now we remark that for *every* class C of structures of bounded treewidth there is a class O obtained by adding to every structure in C an ordering such that S(O) is of bounded treewidth, too, (and hence, DTC > FO on O by Theorem IV.3). This is an immediate consequence of the following result stated in [21] with an even better bound. However, we were unable to verify the proof of that paper and therefore we show the following result in Appendix D.

Theorem VI.4. *For every structure $\mathcal{A}$ there is an ordering $<^{\mathcal{A}}$ such that*

$$\mathrm{tw}((\mathcal{A}, S(<^{\mathcal{A}})) \leq \mathrm{tw}(\mathcal{A}) + 5.$$

The proof of Theorem VI.4 yields the effective version used in the following observation. An $\mathrm{MSO}[\tau_<]$-sentence is *order-invariant* if for every τ-structure $\mathcal{A}$ and any two orderings $<_1$ and $<_2$ on A we have

$$(\mathcal{A}, <_1) \models_{\mathrm{MSO}} \varphi \iff (\mathcal{A}, <_2) \models_{\mathrm{MSO}} \varphi$$

(we write $\models_{\mathrm{MSO}}$ to distinguish the classical semantics from the one we are going to introduce next). Let the *invariant monadic second-order logic* inv-MSO be the logic defined as follows. Its sentences of vocabulary τ are the order-invariant sentences in $\mathrm{MSO}[\tau_<]$.[3] A τ-structure $\mathcal{A}$ is a model of an inv-MSO$[\tau]$-sentence φ, written $\mathcal{A} \models_{\text{inv-MSO}} \varphi$, if $(\mathcal{A}, <^{\mathcal{A}}) \models_{\mathrm{MSO}} \varphi$ for some ordering $<^{\mathcal{A}}$ on A (or, equivalently, for all orderings $<^{\mathcal{A}}$ on A).

Let the *(modulo)-counting monadic second-order logic* CMSO be obtained from MSO by adding for every $m \geq 1$ a (first-order) quantifier $\exists^m_{\equiv 0}$ where $\exists^m_{\equiv 0} x \varphi(x)$ expresses that the number of x satisfying φ is congruent 0 modulo m.

It is easy to see that every CMSO-sentence is equivalent to an inv-MSO-sentence. On the other hand, Courcelle [7] showed that their expressive power coincides on the class of forests. For some time [4], [14] it was believed that it had been proven that inv-MSO and CMSO have the same expressive power on all classes of bounded treewidth. However, in their forthcoming book [8] Courcelle and Engelfried write concerning the validity of inv-MSO $\equiv$ CMSO:

The cases of graphs of path-width at most k and of all graphs (of tree-width at most k) have been announced in [18] and [20] respectively, but the full proofs have not been published. The conjecture is still open.

Using Theorem VI.4 one can generalize Courcelle's Theorem from MSO to inv-MSO. On the class of all structures it is known [14] that inv-MSO > CMSO. Therefore it is worthwhile to mention that, by the same reasons, the version of Courcelle's Theorem in terms of fixed-parameter tractability (e,g, see [13, Theorem 11.37]) generalizes from MSO to inv-MSO.

[3] It is well-known that inv-MSO$[\tau]$ is not a decidable set for every τ containing an at least binary relation symbol.

VII. Structures of bounded local treewidth

Recall that a class C of structures is of *bounded local treewidth* if there is a function $h : \mathbb{N} \to \mathbb{N}$ such that ltw$(\mathcal{A}, r) \leq h(r)$ for all $\mathcal{A} \in$ C and $r \in \mathbb{N}$, where

$$\text{ltw}(\mathcal{A}, r) := \max\{\text{tw}(\mathcal{N}_r^{\mathcal{A}}(a)) \mid a \in A\}$$

and where $\mathcal{N}_r^{\mathcal{A}}(a)$ denotes the subgraph of the Gaifman graph $\mathcal{G}(\mathcal{A})$ induced on the r-neighborhood

$$N_r^{\mathcal{A}}(a) := \{b \in E \mid d^{\mathcal{G}(\mathcal{A})}(a, b) \leq r\}$$

of a. Here, $d^{\mathcal{G}(\mathcal{A})}(a, b)$ denotes the distance between a and b in the graph $\mathcal{G}(\mathcal{A})$. Of course, every class of bounded treewidth is of bounded local treewidth; furthermore, classes of graphs of bounded degree or of planar graphs are of bounded local treewidth.

Let the *locally bounded ordered conjecture* (LBOC) be the statement:

(LBOC) *If* O *is an infinite class of ordered structures whose class* S(O) *of associated successor structures is of bounded local treewidth, then* LFP $>$ FO *on* O.

We believe that (LBOC) holds, even though so far we were unable to prove it. One reason for our belief is the fact that the results of Podewski and Ziegler [24] (see also [3]) mentioned in the Introduction imply that the ordering is not FO-definable on any class of structures of bounded local treewidth.

To help the interested reader to find the relevant parts in [24], [3] we mention here the main steps to obtain this result.

Throughout this section structures maybe finite or infinite, however vocabularies are further assumed to be finite sets of relation symbols.

Let C be a class of structures. We say that *an ordering is* FO*-definable* on C if there is an FO-formula $\varphi(x, y)$ such that for all $\mathcal{A} \in$ C

$$\varphi(\mathcal{A}) := \{(a, b) \mid a, b \in A \text{ and } \mathcal{A} \models \varphi(a, b)\}$$

is an ordering of A. We aim to show:

Theorem VII.1. *If* O *is an infinite class of finite ordered structures with* S(O) *of bounded local treewidth, then no ordering is* FO*-definable on* S(O). *In particular,* DTC $>$ FO *on* S(O).

Clearly, DTC $>$ FO on S(O) if DTC $>$ FO on O; thus, the last conclusion of the theorem already follows from Theorem IV.3 as S(O) is of bounded treewidth.

For $n \geq 1$ denote by $\mathcal{C}_n$ a clique with n elements and for $n, m \geq 1$ by $\mathcal{C}_n^m$ the graph obtained from $\mathcal{C}_n$ by inserting m new vertices on each edge, that is, by replacing every original edge by a path of length $m + 1$. For single structures the following notion was introduced in [24], for classes in [3].

Definition VII.2. A class C of structures is *superflat* if for all $m \geq 1$ there is a $n \geq 1$ such that $\mathcal{C}_n^m \not\subseteq \mathcal{G}(\mathcal{A})$ for all $\mathcal{A} \in$ C.

Here $\mathcal{C}_n^m \subseteq \mathcal{G}$ means that no copy of $\mathcal{C}_n^m$ is a subgraph (in the graph-theoretic sense) of the graph $\mathcal{G}$. A structure $\mathcal{A}$ is *superflat* if the class $\{\mathcal{A}\}$ is.

Clearly, every finite structure is superflat. More important for us is the fact:

Lemma VII.3. *Every class* C *of finite structures of bounded local treewidth is superflat.*

Proof: By definition of bounded local treewidth there is a function $h : \mathbb{N} \to \mathbb{N}$ such that ltw$(\mathcal{A}, r) \leq h(r)$ for all $\mathcal{A} \in$ C and $r \in \mathbb{N}$. To show that C is superflat, let $m \geq 1$. Choose $n > h(m) + 1$. Then, for all $\mathcal{A} \in$ C we have $\mathcal{C}_n^m \not\subseteq \mathcal{G}(\mathcal{A})$ as otherwise for every a (in the copy) of $\mathcal{C}_n^m$ (in $\mathcal{G}(\mathcal{A})$), we have $C_n^m \subseteq N_{2(m+1)}^{\mathcal{A}}(a)$ and hence, tw$(\mathcal{N}_{2(m+1)}^{\mathcal{A}}(a)) \geq n > h(m)$, a contradiction. □

The result of [24] we need reads as follows:

Theorem VII.4. *If* $\mathcal{A}$ *is an infinite and superflat structure, then no ordering is* FO*-definable on* $\mathcal{A}$.

In [24] this result corresponds to Corollary 9 and part a) of the remark at the end of that paper. Note on the one hand that there the authors show a more general result than the version we need. On the other hand, they only consider structures in vocabularies with at most binary relation symbols; however, their proofs easily generalize to arbitrary finite vocabularies.

Let C be a class of structures. Let Th(C) be the first-order theory of C, that is, the set of FO-sentences holding in *all* structures of C. Denote by C* the class of models of Th(C),

$$\text{C}^* := \{\mathcal{A} \mid \mathcal{A} \models \text{Th(C)}\}.$$

Clearly, C $\subseteq$ C*.

Lemma VII.5. *If* C *is a superflat class of structures, then* C* *is superflat.*

Proof: For $m, n \geq 1$ there is a FO-sentence expressing "no subgraph in the Gaifman graph is isomorphic to $\mathcal{C}_n^m$." Here, we use that our vocabulary is finite. □

Proof of Theorem VII.1: Let O be an infinite class of finite ordered structures with $\mathrm{S(O)}$ of bounded local treewidth. By contradiction, assume that $\varphi(x, y) \in$ FO defines an ordering on every structure in C $:= S(\mathrm{O})$ and hence, on every structure in C*. As C is infinite, by the compactness theorem the class C* contains an infinite structure $\mathcal{A}$, which is superflat by the previous lemma. As $\varphi(x, y)$ defines an ordering on $\mathcal{A}$, this contradicts Theorem VII.4. □

References

[1] S. Abiteboul and V. Vianu. Fixpoint extensions of first-order logic and datalog-like languages. In *Proc. of the Fourth Annual Symposium on Logic in Computer Science, LICS'89*, pages 71–79, 1989.

[2] H. Adler and I. Adler. A note on clique-width and tree-width for structures. *CoRR*, abs/0806.0103, 2008.

[3] H. Adler and I. Adler. Nowhere dense graph classes, stability, and the independence property. *CoRR*, abs/1011.4016, 2010.

[4] M. Benedikt and L. Segoufin. Towards a characterization of order-invariant queries over tame graphs. *The Journal Symbolic Logic*, 74(1):168–186, 2009.

[5] Y. Chen and J. Flum. A logic for PTIME and a parameterized halting problem. In *Fileds of Logic and Computation*, pages 251–276. Springer, 2010.

[6] Y. Chen and J. Flum. On the complexity of Gödel's proof predicate. *The Journal of Symbolic Logic*, 75(1):239–254, 2010.

[7] B. Courcelle. The monadic second-order logic of graphs X: Linear orderings. *Theoretical Computer Science*, 160(1&2):87–143, 1996.

[8] B. Courcelle and J. Engelfriet. *Graph structure and monadic second-order logic, a language theoretic approach*. Cambridge University Press, To appear, 2012.

[9] B. Courcelle, J. A. Makowsky, and U. Rotics. On the fixed parameter complexity of graph enumeration problems definable in monadic second-order logic. *Discrete Applied Mathematics*, 108(1-2):23–52, 2001.

[10] A. Dawar and L. Hella. The expressive power of finitely many generalized quantifiers. *Information and Computation*, 123(2):172–184, 1995.

[11] A. Dawar, S. Lindell, and S. Weinstein. First order logic, fixed point logic and linear order. In *Computer Science Logic (CSL'95)*, pages 161–177, 1995.

[12] M. Elberfeld, A. Jakoby, and T. Tantau. Logspace versions of the theorems of Bodlaender and Courcelle. In *Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science, (FOCS'10)*, pages 143–152, 2010.

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

[14] T. Ganzow and S. Rubin. Order-invariant MSO is stronger than counting MSO in the finite. In *Proceedings of the 25th Annual Symposium on Theoretical Aspects of Computer Science (STACS'08)*, pages 313–324, 2008.

[15] M. Grohe and G. Turán. Learnability and definability in trees and similar structures. *Theory of Computing Systems*, 37(1):193–220, 2004.

[16] N. Immerman. Relational queries computable in polynomial time. *Information and Control*, 68:86–104, 1986.

[17] N. Immerman. Languages that capture complexity classes. *SIAM Journal on Computing*, 16(4):770–778, 1987.

[18] V. Kabanets. Recognizability equals definability for partial k-paths. In *Proc. of the 24th International Colloquium on Automata, Languages and Programming (ICALP'97)*, pages 805–815, 1997.

[19] P. Kolaitis and M. Y. Vardi. Fixpoint logic vs. infinitary logic in finite-model theory. In *Proceedings of the Seventh Annual IEEE Symposium on Logic in Computer Science (LICS'92)*, pages 46–57, 1992.

[20] D. Lapoire. Recognizability equals monadic second-order definability for sets of graphs of bounded tree-width. In *Proc. of the 15th Annual Symposium on Theoretical Aspects of Computer Science (STACS'98)*, pages 618–628, 1998.

[21] J. A. Makowsky. Coloured tutte polynomials and kauffman brackets for graphs of bounded tree width. *Discrete Applied Mathematics*, 145(2):276–290, 2005.

[22] G. L. McColm. When is arithmetic possible? *Annals of Pure and Applied Logic*, 50(1):29–51, 1990.

[23] S. Oum and P. D. Seymour. Approximating clique-width and branch-width. *Journal of Combinatorial Theory, Series B*, 96(4):514–528, 2006.

[24] K. Podelski and M. Ziegler. Stable graphs. *Fundamenta Mathematicae*, 100:101–107, 1978.

[25] M. Y. Vardi. The complexity of relational query languages. In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing (STOC'82)*, pages 137–146, 1982.

Appendix

A. Proof of Proposition II.4.

Let C be a class of τ-structures and $\varphi_1, \varphi_2, \dots$ an enumeration of the L-sentences of vocabulary τ. We define the sequence $(\mathrm{C}_i)_{i\geq 0}$ of infinite classes C_i with

$$\mathrm{C} = \mathrm{C}_0 \supseteq \mathrm{C}_1 \supseteq \cdots \supseteq \mathrm{C}_i \supseteq \cdots$$

inductively as follows. Assume we already have defined the infinite class C_i. Then, at least one of φ_i and $\neg\varphi_i$ holds in an infinite subclass of C_i; denote such a formula by ψ_i. We let C_{i+1} be the class

$$\big\{\mathcal{A} \in \mathrm{C}_i \bigm| \mathcal{A} \models \psi_i\big\} \cup \big\{\mathcal{B} \bigm| \mathcal{B} \text{ has a } (\leq i)\text{th smallest size in } \mathrm{C}_i\big\}.^4 \tag{1}$$

Then C_{i+1} is infinite and $\mathrm{C}_i \supseteq \mathrm{C}_{i+1}$. Finally we set $\mathrm{C}' := \bigcap_{i\geq 0} \mathrm{C}_i$.

Claim. C' is infinite and for every $i \geq 1$ only finitely many structures in C' are not a model of ψ_i.

Proof of the Claim. Let $i \geq 1$. Then, by definition of C_{i+1}, we have

$$\big\{\mathcal{B} \bigm| \mathcal{B} \text{ has a } (\leq i)\text{th smallest size in } \mathrm{C}_i\big\} \subseteq \big\{\mathcal{B} \bigm| \mathcal{B} \text{ has a } (\leq (i+1))\text{th smallest size in } \mathrm{C}_{i+1}\big\}.$$

Thus, $\big\{\mathcal{B} \bigm| \mathcal{B} \text{ has a } (\leq i)\text{th smallest size in } \mathrm{C}_i\big\} \subseteq \mathrm{C}'$ and therefore, C' contains at least i many structures. Since i is arbitrary, we conclude that C' is infinite. Let $\mathcal{A} \in \mathrm{C}'$ and assume that $\mathcal{A}$ is not a model of ψ_i, then, by (1), $\mathcal{A} \in \big\{\mathcal{B} \bigm| \mathcal{B} \text{ has a } (\leq i)\text{th smallest size in } \mathrm{C}_i\big\}$. Thus, there at most finitely many such $\mathcal{A}$. ⊣

By the claim the formula ψ_i is equivalent to any valid first-order sentence (say, to $\exists x\, x = x$) in all but finitely many structure of C'. Taking care of these finitely many structures separately, we get an FO-sentence equivalent to ψ_i in C'. □

B. Proof of Proposition II.4.

First we recall the definition of cliquewidth from [15]. A *k-colored τ-structure* is a pair $(\mathcal{A}, \gamma)$ consisting of a τ-structure and a mapping $\gamma : A \to [k]$. A *basic k-colored τ-structure* is a k-colored τ-structure $(\mathcal{A}, \gamma)$ where $|A| = 1$ and $R^{\mathcal{A}} = \emptyset$ for all $R \in \tau$. We let $\Gamma_k[\tau]$ be the smallest class of k-colored τ-structures that contains all basic k-colored τ-structures and is closed under the following operations:

– *Union*: Take the union of two k-colored τ-structures with disjoint universes.
– *$(i \to j)$-recoloring,* for $i, j \in [k]$: Take a k-colored τ-structure and recolor all vertices colored i to j.
– *$(R, i_1, \dots, i_r)$-connecting, for every $r \geq 1$, every r-ary $R \in \tau$ and $i_1, \dots, i_r \in [k]$*: Take a k-colored τ-structure $(\mathcal{A}, \gamma)$ and add all tuples $(a_1, \dots, a_r) \in A^r$ with $\gamma(a_j) = i_j$ for $j \in [r]$ to $R^{\mathcal{A}}$.

The *cliquewidth* of a τ-structure $\mathcal{A}$, denoted by $\mathrm{cw}(\mathcal{A})$, is the minimum k such that there exists a k-coloring $\gamma : A \to [k]$ such that $(\mathcal{A}, \gamma) \in \Gamma_k[\tau]$.

Claim 1: Let τ be a vocabulary of arity ≤ 2. There is a polynomial time algorithm that assigns to every τ-structure $\mathcal{A}$ a graph $\mathcal{G}_{\mathcal{A}}$ such that

$$\mathrm{cw}(\mathcal{G}_{\mathcal{A}}) \leq 3 \cdot |\tau| \cdot (\mathrm{cw}(\mathcal{A}) + 1).$$

Proof of the Claim: Let $\tau = \{R_1, \dots, R_\ell\}$ and, for notational simplicity, assume that all R_m are binary. The vertices and edges of $\mathcal{G}_{\mathcal{A}}$ are obtained by the following steps:

(1) Every $a \in A$ is a vertex in $\mathcal{G}_{\mathcal{A}}$;
(2) To every $a \in A$ we add a new vertex a' and the edge $\{a, a'\}$ to $\mathcal{G}_{\mathcal{A}}$;
(3) For every $a \in A$ and $m \in [\ell]$ we add two cliques $\mathcal{C}_{2m}(a)$ and $\mathcal{C}_{2m+1}(a)$ of new vertices to $\mathcal{G}_{\mathcal{A}}$; all vertices of $\mathcal{C}_{2m}(a)$ or $\mathcal{C}_{2m+1}(a)$ are linked to a by an edge;
(4) For every $m \in [\ell]$ and every $(a, b) \in R_m^{\mathcal{A}}$ we link every vertex of $\mathcal{C}_{2m}(a)$ to every vertex of $\mathcal{C}_{2m+1}(b)$.

Assume that $\mathrm{cw}(\mathcal{A}) = k$. For $\mathcal{G}_{\mathcal{A}}$, besides the colors in $[k]$, we use colors denoted by $(m, i, 1)$ and $(m, i, 2)$ for $m \in [\ell]$ and $i \in [k]$, and "auxiliary" colors b, c, d. Then $k' := k + 2 \cdot |\tau| \cdot |k| + 3$ is the number of all these colors.

Consider a sequence of operations leading to a k-colored $(\mathcal{A}, \gamma)$ for some γ. Along this sequence it is not hard to obtain a k'-colored $(\mathcal{G}_{\mathcal{A}}, \gamma')$ for some γ'. For example, whenever in the sequence for $(\mathcal{A}, \gamma)$ a new element a of A of color i is introduced, we ensure that we get the vertices a' and, for $m \in [\ell]$, the vertices of the cliques $\mathcal{C}_{2m}(a)$ and $\mathcal{C}_{2m+1}(a)$ that (after some recoloring) they are colored by b, $(m, i, 1)$ and $(m, i, 2)$, respectively, and that they are all linked with a. ⊣

The preceding proof cannot be extended in a straightforward way to vocabularies of higher arity.

[4] The structure $\mathcal{B}$ has a $(\leq i)$th smallest size in a class D of structures if $\mathcal{B} \in \mathrm{D}$ and there is a $j \leq i$ and a sequence $\mathcal{B}_1, \dots, \mathcal{B}_{j-1}$ of structures in D such that $|B_1| < |B_2| < \dots |B_{j-1}| < |B|$, and there is no such sequence for $j = i+1$.

In this context we draw the reader's attention to [2] where a class of structures with a ternary relation is exhibited that has bounded cliquewith but unbounded treewidth.

Claim 2: Let τ be a vocabulary of arity ≤ 2. Then every τ-structure $\mathcal{A}$ is FO-definable in the graph $\mathcal{G}_\mathcal{A}$; that is, (we assume again that $\tau = \{R_1, \dots, R_\ell\}$ and that all R_m are binary), there are first-order formulas $\varphi(x)$ and $\varphi_1(x,y), \dots, \varphi_\ell(x,y)$ such that for all τ-structures $\mathcal{A}$ we have

$$\mathcal{A} = (\varphi(\mathcal{G}_\mathcal{A}), \varphi_1(\mathcal{G}_\mathcal{A}), \dots, \varphi_\ell(\mathcal{G}_\mathcal{A})),$$

where

$$\varphi(\mathcal{G}_\mathcal{A}) := \big\{a \in G_\mathcal{A} \bigm| \mathcal{G}_\mathcal{A} \models \varphi(a)\big\}$$

and for $m \in [\ell]$,

$$\varphi_m(\mathcal{G}_\mathcal{A}) := \big\{(a,b) \in G_\mathcal{A}^2 \bigm| \mathcal{G}_\mathcal{A} \models \varphi_m(a,b)\big\}.$$

Proof of the Claim 2: We observe that

- the elements of A are precisely those vertices of the graph $\mathcal{G}_\mathcal{A}$ that have a neighbor of degree 1 (namely a');
- a tuple (a,b) of elements of A is in R_m if and only if a forms together with $2m$ new elements a maximal clique and b forms with $2m+1$ new elements a maximal clique, too, and the new elements of these two cliques form a clique themselves.

As these properties can be expressed in FO, the claim follows. $\dashv$

Now let O be a class of structures in a vocabulary τ of arity ≤ 2 of bounded cliquewidth. Let $\mathrm{C_O} := \{\mathcal{G}_\mathcal{A} \mid \mathcal{A} \in \mathrm{O}\}$. By Claim 1 and Claim 2, p-MC(O, MSO) is fpt-reducible to p-MC($\mathrm{C_O}$, MSO) and the class $\mathrm{C_O}$ has bounded cliquewidth, too. By results of [9] and [23], we know that p-MC($\mathrm{C_O}$, MSO) $\in$ FPT. Then, the claim follows by Corollary III.2. $\Box$

C. Proof of Theorem IV.3.

We first introduce some concepts and derive some simple results which will be used in the proof of this theorem. If $\mathcal{A}$ is an ordered $\tau_<$-structure and $a, a' \in A$ we define the "open" intervals $]a, a'[^\mathcal{A}$, $]-\infty, a[^\mathcal{A}$, and $]a, \infty[^\mathcal{A}$ in the obvious way, e.g.,

$$]-\infty, a[^\mathcal{A} := \big\{b \in A \bigm| b <^\mathcal{A} a\big\}.$$

Definition C.1. Let $\mathcal{A}$ and $\mathcal{B}$ be ordered $\tau_<$-structures. Assume that $A \cap B = \{e_1, \dots, e_m\}$ with $m \geq 1$ and $e_1 <^\mathcal{A} \cdots <^\mathcal{A} e_m$. We say that $\mathcal{A}$ and $\mathcal{B}$ are *order-consistent* if

(i) $e_1 <^\mathcal{B} \cdots <^\mathcal{B} e_m$;
(ii) at least one of the intervals $]-\infty, e_1[^\mathcal{A}$ and $]-\infty, e_1[^\mathcal{B}$ is empty;
(iii) for every $i \in [m-1]$, at least one of the intervals $]e_i, e_{i+1}[^\mathcal{A}$ and $]e_i, e_{i+1}[^\mathcal{B}$ is empty;
(iv) at least one of the intervals $]e_m, \infty[^\mathcal{A}$ and $]e_m, \infty[^\mathcal{B}$ is empty.

For order-consistent structures $\mathcal{A}$ and $\mathcal{B}$ we introduce the ordered union $\mathcal{A} \cup_< \mathcal{B}$ (along $A \cap B$) and show:

Lemma C.2. *If $\mathcal{A}$ and $\mathcal{B}$ are order-consistent $\tau_<$-structures with $A \cap B = \{e_1, \dots, e_m\}$ and $m \geq 1$, then the* ordered union *$\mathcal{A} \cup_< \mathcal{B}$ is an ordered structure, where $\mathcal{A} \cup_< \mathcal{B}$ is the $\tau_<$-structure obtained from the union $\mathcal{A} \cup \mathcal{B}$ by adding, for $a \in A \setminus B$ and $b \in B \setminus A$, to the interpretation in the union of $<$ the tuple (a,b) if $\big(a <^\mathcal{A} e_i$ and $e_i <^\mathcal{B} b\big)$ for some $i \in [m]$ and otherwise, the tuple (b,a).*

Proof: The conditions (i)–(iv) of Definition C.1 guarantee that for distinct $d, d' \in A \cup B$ exactly one of $d <^{\mathcal{A} \cup_< \mathcal{B}} d'$ and $d' <^{\mathcal{A} \cup_< \mathcal{B}} d$ holds. The rest is simple. $\Box$

The following result extends Proposition IV.2 and can be obtained by the Ehrenfeucht-Fraïssé method, too.

Proposition C.3. *Let $\mathcal{A}$ and $\mathcal{B}$ be order-consistent $\tau_<$-structures with $A \cap B = \{e_1, \dots, e_m\}$. Similarly, let $\mathcal{A}'$ and $\mathcal{B}'$ be order-consistent $\tau_<$-structures with $A' \cap B' = \{e'_1, \dots, e'_m\}$. Set $\bar{e} := e_1, \dots, e_m$ and $\bar{e}' := e'_1, \dots, e'_m$. Assume that $\bar{a} = a_1, \dots, a_\ell \in A \setminus B$ and $\bar{b} = b_1, \dots, b_s \in B \setminus A$ and similarly that $\bar{a}' = a'_1, \dots, a'_\ell \in A' \setminus B'$ and $\bar{b}' = b'_1, \dots, b'_s \in B' \setminus A'$. Let $i \in \mathbb{N}$. If*

$$(\mathcal{A}, \bar{a}, \bar{e}) \equiv^i (\mathcal{A}', \bar{a}', \bar{e}') \text{ and } (\mathcal{B}, \bar{b}, \bar{e}) \equiv^i (\mathcal{B}', \bar{b}', \bar{e}'),$$

then

$$(\mathcal{A} \cup_< \mathcal{B}, \bar{a}, \bar{b}, \bar{e}) \equiv^i (\mathcal{A}' \cup_< \mathcal{B}', \bar{a}', \bar{b}', \bar{e}').$$

Trees are directed graphs $\mathcal{T} = (T, E^\mathcal{T})$ that have a distinguished vertex $r \in T$, the *root of* $\mathcal{T}$, such that for every vertex $t \in T$ there is exactly one (directed) path from r to t. We usually call the vertices of a tree *nodes*. The partial ordering $\preceq^\mathcal{T}$ on T is defined by $\big(t' \preceq^\mathcal{T} t \iff t'$ is on the path from r to $t\big)$. For $t \in T$ we denote by $\mathcal{T}_t$ the subtree rooted at t; its set of nodes is $T_t = \big\{t' \bigm| t \preceq^\mathcal{T} t'\big\}$. A tree is *binary* if every node has at most two children. Recall that we write tree decompositions in the form $(\mathcal{T}, (B_t)_{t\in T})$, where $\mathcal{T}$ is the corresponding tree and B_t the *bag* at t.

For every bag B of a tree decomposition of an ordered structure $\mathcal{A}$ we may write $\mathcal{A}$ as the ordered union along B; more precisely:

Lemma C.4. *Let $(\mathcal{A}, <^{\mathcal{A}})$ be an ordered $\tau_<$-structure and $(\mathcal{T}, (B_t)_{t\in T})$ a tree decomposition of the structure $(\mathcal{A}, S(<^{\mathcal{A}}))$. For every $t \in T$ with nonempty B_t we set*

- *$(\mathcal{A}_t^-, <^{\mathcal{A}_t^-})$ the substructure of $(\mathcal{A}, <^{\mathcal{A}})$ with universe $A_t^- := \bigcup_{t'\in T_t} B_{t'}$;*
- *$(\mathcal{A}_t^+, <^{\mathcal{A}_t^+})$ the substructure of $(\mathcal{A}, <^{\mathcal{A}})$ with universe $A_t^+ := B_t \cup \bigcup_{t'\notin T_t} B_{t'}$.*

Then, $(\mathcal{A}_t^-, <^{\mathcal{A}_t^-})$ and $(\mathcal{A}_t^+, <^{\mathcal{A}_t^+})$ are order-consistent and $(\mathcal{A}, <^{\mathcal{A}}) = (\mathcal{A}_t^-, <^{\mathcal{A}_t^-}) \cup_< (\mathcal{A}_t^+, <^{\mathcal{A}_t^+})$.

Proof: Let $B_t = \{e_1, \ldots, e_m\}$ with $e_1 <^{\mathcal{A}} \ldots <^{\mathcal{A}} e_m$. For example, we verify (iii) in Definition C.1. So, for contradiction, assume that there is an $i \in [m-1]$ such that both, $]e_i, e_{i+1}[^{(\mathcal{A}_t^-, <^{\mathcal{A}_t^-})}$ and $]e_i, e_{i+1}[^{(\mathcal{A}_t^+, <^{\mathcal{A}_t^+})}$, are nonempty. Then there are $a_1, a_2 \in A$ such that

$$e_1 <^{\mathcal{A}} a_1 <^{\mathcal{A}} a_2 <^{\mathcal{A}} e_2,$$

with $(a_1, a_2) \in S(<^{\mathcal{A}})$ and

$$\begin{aligned}&\big(a_1 \in A_t^- \setminus A_t^+ \text{and } a_2 \in A_t^+ \setminus A_t^-\big)\\ &\quad\text{or } \big(a_1 \in A_t^+ \setminus A_t^- \text{ and } a_2 \in A_t^- \setminus A_t^+\big).\end{aligned} \tag{2}$$

By $(a_1, a_2) \in S(<^{\mathcal{A}})$ there is a bag $B_{t'}$ with $\{a_1, a_2\} \subseteq B_{t'}$. It follows that $\{a_1, a_2\} \subseteq A_t^-$ or$\{a_1, a_2\} \subseteq A_t^+$ contradicting (2). Now, the equality $(\mathcal{A}, <^{\mathcal{A}}) = (\mathcal{A}_t^-, <^{\mathcal{A}_t^-}) \cup_< (\mathcal{A}_t^+, <^{\mathcal{A}_t^+})$ is easily verified. □

After these preliminaries we start with the proof of Theorem IV.3. Let O be an infinite class of ordered $\tau_<$-structures and assume that, for some $k \in \mathbb{N}$, tw(S(O)) $\leq k$. Wlog. we may assume that the universe of every structure in O has cardinality greater than $2^{2^{k+1}}$.

We only consider tree decompositions where the corresponding tree has a "canonical" ordering of the children of every node, so that we can speak, say, of the left-most path of the tree.

Using the algorithm of [12] we see that there is an algorithm $\mathbb{D}$ that in logarithmic space computes for every ordered $\tau_<$-structure $(\mathcal{A}, <^{\mathcal{A}})$ with

$$2^{2^{k+1}} < |A| \quad \text{and} \quad \text{tw}\big((\mathcal{A}, S(<^{\mathcal{A}}))\big) \leq k \tag{3}$$

a tree decomposition $\big(\mathcal{T}, (B_t)_{t\in T}\big)$ of $(\mathcal{A}, S(<^{\mathcal{A}}))$ with binary tree $\mathcal{T}$ and with $|B_t| = k+1$ for every $t \in T$ and determines the unique node $t_0 \in T$ satisfying (i) and (ii):

(i) $2^{2^{c(t_0)}} \leq |A| < k + 1 + 2 \cdot 2^{2^{c(t_0)}}$ where $c(t_0) := |A_{t_0}^-|$ (as $\mathcal{T}$ has binary branching there is a node satisfying these inequalities);

(ii) t_0 is on the left-most path of $\mathcal{T}$ containing a node satisfying (i) and on that path t_0 is the $\preceq^{\mathcal{T}}$-smallest node with this property.

Let $B_{t_0} = \{a_1, \ldots, a_{k+1}\}$ with $a_1 <^{\mathcal{A}} \ldots <^{\mathcal{A}} a_{k+1}$. We set

$$\begin{aligned}(\mathcal{A}^-, <^{\mathcal{A}^-}) &:= (\mathcal{A}_{t_0}^-, <^{\mathcal{A}_{t_0}^-}),\\ (\mathcal{A}^+, <^{\mathcal{A}^+}) &:= (\mathcal{A}_{t_0}^+, <^{\mathcal{A}_{t_0}^+}),\\ \text{and} \qquad \bar{a} &:= a_1, \ldots, a_{k+1}.\end{aligned}$$

The algorithm $\mathbb{D}$ assigns to $(\mathcal{A}, <^{\mathcal{A}}) \in$ O the ordered structures $(\mathcal{A}^-, <^{\mathcal{A}^-})$, $(\mathcal{A}^+, <^{\mathcal{A}^+})$, and the sequence $\bar{a}$. By Lemma C.4, we know that

$$(\mathcal{A}, <^{\mathcal{A}}) = (\mathcal{A}^-, <^{\mathcal{A}^-}) \cup_< (\mathcal{A}^+, <^{\mathcal{A}^+}). \tag{4}$$

To ordered $\tau_<$-structures $(\mathcal{A}, <^{\mathcal{A}})$ not satisfying (3) the algorithm $\mathbb{D}$ assigns as $(\mathcal{A}^-, <^{\mathcal{A}^-})$ and as $(\mathcal{A}^+, <^{\mathcal{A}^+})$ the substructure of $(\mathcal{A}, <^{\mathcal{A}})$ whose universe only contains the $<^{\mathcal{A}}$-first element a. It sets $\bar{a} := a, \ldots, a$.

As the algorithm $\mathbb{D}$ is only applied to *ordered* $\tau_<$-structures, we can assume that

$$\begin{aligned}(\mathcal{A}^-, <^{\mathcal{A}^-}) &= (\mathcal{B}^-, <^{\mathcal{B}^-}),\\ (\mathcal{A}^+, <^{\mathcal{A}^+}) &= (\mathcal{B}^+, <^{\mathcal{B}^+}),\\ \text{and} \qquad \bar{a} &= \bar{b}\end{aligned}$$

if $(\mathcal{A}, <^{\mathcal{A}}) \cong (\mathcal{B}, <^{\mathcal{B}})$.

By contradiction, let us assume that DTC > FO on O. We let $\tau_<(k) := \tau_< \cup \{c_1, \ldots, c_{k+1}\}$ with constant symbols $c_1, \ldots, c_{k+1}$. We define a sequence $(\mathrm{C}_i)_{i\geq 0}$ of infinite classes C_i of ordered $\tau_<(k)$-structures with

$$\mathrm{C}_0 \supseteq \mathrm{C}_1 \supseteq \cdots \supseteq \mathrm{C}_i \supseteq \cdots \tag{5}$$

inductively. We set

$$\mathrm{C}_0 := \big\{(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \bigm| (\mathcal{A}, <^{\mathcal{A}}) \in \mathrm{O}\big\}.^5$$

By (i), the class C_0 is infinite. Assume we already have defined the infinite subclass C_i of C_0. Then there is at least one first-order i-type θ_i such that for infinitely many structures $(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \in \mathrm{C}_i$ we have $\text{tp}^i((\mathcal{A}^+, <^{\mathcal{A}^+}, \bar{a})) = \theta_i$. We let C_{i+1} be

$$\begin{aligned}&\{(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \in \mathrm{C}_i \mid \text{tp}^i((\mathcal{A}^+, <^{\mathcal{A}^+}, \bar{a})) = \theta_i\}\\ &\quad \cup \big\{\mathcal{B} \mid \mathcal{B} \text{ has a } (\leq i)\text{th smallest size in } \mathrm{C}_i\big\}.\end{aligned}$$

[5] We may assume that for every $\mathcal{B} \in \mathrm{C}_0$ there is at most one $(\mathcal{A}, <^{\mathcal{A}}) \in$ O with $\mathcal{B} = (\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a})$; if this is not the case we start with a corresponding infinite subclass of O.

Finally we set $\mathrm{C} = \bigcap_{i \geq 0} \mathrm{C}_i$.

The proof of the following claim is straightforward.

Claim 1. C is infinite and for every $i \geq 0$, $\mathrm{tp}^i((\mathcal{A}^+, <^{\mathcal{A}^+}, \bar{a})) = \theta_i$ for all but finitely many structures $(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \in \mathrm{C}$.

Claim 2. For every PFP$[\tau_<(k)]$-formula $\varphi(\bar{x})$ there is an $i \in \mathbb{N}$ and an FO$[\tau_<(k)]$-formula $\rho(\bar{x})$ such that for $(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \in \mathrm{C}$ with $\mathrm{tp}^i((\mathcal{A}^+, <^{\mathcal{A}^+}, \bar{a})) = \theta_i$ and $\bar{b} \in (A^-)^{|\bar{x}|}$ we have

$$(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \models \varphi(\bar{b}) \iff (\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \models \rho(\bar{b}).$$

Proof of Claim 2. Let $\varphi(\bar{x})$ be a PFP-formula and for notational simplicity, assume that $\bar{x} = x_1, x_2$. We consider the following property of arbitrary ordered $\tau_<$-structures $(\mathcal{A}, <^{\mathcal{A}})$ and $b_1, b_2 \in A$:

$$b_1, b_2 \in A^- \text{ and } (\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \models \varphi(b_1, b_2). \quad (6)$$

By (i), we know that $2^{2^{|A^-|}} \leq |A|$ if A^- contains more than one element. Hence, for given $(\mathcal{A}, <^{\mathcal{A}})$ and $b_1, b_2 \in A$, the property (6) can be checked in logarithmic space and, by the properties of $\mathbb{D}$, the class of $(\mathcal{A}, <^{\mathcal{A}}, b_1, b_2)$ with this property is closed under isomorphism. Therefore, there is an DTC-formula $\psi(x_1, x_2)$ such that for all ordered $\tau_<$-structures $(\mathcal{A}, <^{\mathcal{A}})$ and $b_1, b_2 \in A$,

$$(\mathcal{A}, <^{\mathcal{A}}) \models \psi(b_1, b_2) \iff \mathcal{A} \text{ and } b_1, b_2 \text{ satisfy (6)}. \quad (7)$$

By our assumption DTC $\equiv$ FO on O, there is an FO-formula $\xi(x_1, x_2)$ such that for every structure $(\mathcal{A}, <^{\mathcal{A}}) \in \mathrm{O}$ and $b_1, b_2 \in A$ we have

$$(\mathcal{A}, <^{\mathcal{A}}) \models \psi(b_1, b_2) \iff (\mathcal{A}, <^{\mathcal{A}}) \models \xi(b_1, b_2). \quad (8)$$

Let i be the quantifier rank of ξ. We let $\rho(x_1, x_2)$ be the FO-formula

$$\bigvee \big\{ \mathrm{tp}^i_{d_1,d_2}((\mathcal{E}^-, <^{\mathcal{E}^-}, \bar{e}))(x_1, x_2) \bigm| (\mathcal{E}, <^{\mathcal{E}}) \in \mathrm{O},\ (\mathcal{E}^-, <^{\mathcal{E}^-}, \bar{e}) \in \mathrm{C},\ d_1, d_2 \in E^-,\ (\mathcal{E}, <^{\mathcal{E}}) \models \xi(d_1, d_2), \text{ and } \mathrm{tp}^i((\mathcal{E}^+, <^{\mathcal{E}^+}, \bar{e})) = \theta_i \big\}.$$

Then, for every $(\mathcal{A}, <^{\mathcal{A}}) \in \mathrm{O}$ with $(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \in \mathrm{C}$ and $\mathrm{tp}^i((\mathcal{A}^+, <^{\mathcal{A}^+}, \bar{a})) = \theta_i$ and all $b_1, b_2 \in A^-$, we have

$$(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \models \rho(b_1, b_2) \iff (\mathcal{A}, <^{\mathcal{A}}) \models \xi(b_1, b_2). \quad (9)$$

In fact, the implication from right to left is easy. Now assume that $(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \models \rho(b_1, b_2)$. Then there is $(\mathcal{E}, <^{\mathcal{E}}) \in \mathrm{O}$ with $(\mathcal{E}^-, <^{\mathcal{E}^-}, \bar{e}) \in \mathrm{C}$ and there are $d_1, d_2 \in E^-$ with

$$(\mathcal{E}, <^{\mathcal{E}}) \models \xi(d_1, d_2) \text{ and } \mathrm{tp}^i((\mathcal{E}^+, <^{\mathcal{E}^+}, \bar{e})) = \theta_i$$

such that

$$(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \models \mathrm{tp}^i_{d_1,d_2}((\mathcal{E}^-, <^{\mathcal{E}^-}, \bar{e}))(b_1, b_2).$$

Thus,

$$(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}, b_1, b_2) \equiv^i (\mathcal{E}^-, <^{\mathcal{E}^-}, \bar{e}, d_1, d_2)$$

and

$$(\mathcal{A}^+, <^{\mathcal{A}^+}, \bar{a}) \equiv^i (\mathcal{E}^+, <^{\mathcal{E}^+}, \bar{e})$$

as both structures have i-type θ_i. By Proposition C.3 and Lemma C.4,

$$(\mathcal{A}, <^{\mathcal{A}}, \bar{a}, b_1, b_2) \equiv^i (\mathcal{E}, <^{\mathcal{E}}, \bar{e}, d_1, d_2)$$

Therefore, $(\mathcal{A}, <^{\mathcal{A}}) \models \xi(b_1, b_2)$.

Putting the equivalences (7)–(9) together we get for $(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \in \mathrm{C}$ with $\mathrm{tp}^i((\mathcal{A}^+, <^{\mathcal{A}^+}, \bar{a})) = \theta_i$ and $b_1, b_2 \in A^-$,

$$\begin{aligned} &(\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \models \varphi(b_1, b_2) \\ &\iff (\mathcal{A}, <^{\mathcal{A}}) \models \psi(b_1, b_2) && \text{(by (7)} \\ &\iff (\mathcal{A}, <^{\mathcal{A}}) \models \xi(b_1, b_2) && \text{(by (8)} \\ &\iff (\mathcal{A}^-, <^{\mathcal{A}^-}, \bar{a}) \models \rho(b_1, b_2) && \text{(by (9).} \end{aligned}$$

This finishes the proof of Claim 2. ⊣

Now it is easy to show that PFP $\equiv$ FO on C. Given a PFP-formula $\varphi(\bar{x})$ we determine $i \geq 1$ and an FO-formula $\rho(\bar{x})$ satisfying Claim 2. Taking care of the finitely many structures not covered by Claim 2 separately, we get an FO-formula equivalent to $\varphi(\bar{x})$ in C. □

D. Proof of Theorem VI.4.

We use the notation and conventions introduced for trees after Proposition C.3. Let $\mathcal{A}$ be a structure of treewidth k. Let $\big(\mathcal{T}, (B_t)_{t \in T}\big)$ be a smooth tree decomposition of width k of $\mathcal{A}$; recall that *smooth* means that $|B_t| = k+1$ for every $t \in T$ and that

$$\big|B_t \cap B_{t_0}\big| = k$$

for every $(t_0, t) \in E^{\mathcal{T}}$.

If $(t_0, t) \in E^{\mathcal{T}}$, we let a_t be the element of A with

$$\{a_t\} = B_t \setminus B_{t_0}.$$

and

$$W := W\big(\mathcal{A}, \mathcal{T}, (B_t)_{t \in T}\big) := \big\{a_t \bigm| t \in T \setminus \{r\}\big\}.$$

Note that $A \setminus W \subseteq B_r$.

For every subtree $\mathcal{T}'$ of $\mathcal{T}$ let $\mathcal{A}_{\mathcal{T}'}$ be the substructure of $\mathcal{A}$ induced on the set

$$\bigcup_{t' \in T'} B_{t'}.$$

One easily verifies that

$$\big(\mathcal{T}, (B_{t'})_{t' \in T'}\big)$$

is a smooth tree decomposition of $\mathcal{A}_{\mathcal{T}'}$.

Again we assume that the children of every node $t \in T$ are ordered. Therefore we can speak of the first child $f(t)$ of t, of the last child $\ell(t)$ of t, and of the the *rightmost leaf* $\textit{rleaf}(\mathcal{T})$ of T.

We first define an ordering $<^W$ on W. We present it by the enumeration $e(W)$ of the nodes of W according to this ordering. We do this by induction (on the number of nodes of $\mathcal{T}$). At the same time we present a tree decomposition of $\big(\mathcal{A}, S(<^W)\big)$. Thereby we will ensure that:

(T1) If $W \neq \emptyset$ and $e(W) = a_1, \ldots, a_m$, then $a_1 = a_{f(r)}$ and $a_m = a_{\textit{rleaf}(\mathcal{T})}$.

(T2) The width of the tree decomposition is at most $k + 5$.

(T3) The bag at the root is

$$\begin{cases} B_r, & \text{if } \mathcal{T} \text{ contains only one node} \\ B_r \cup \big\{a_{f(r)}, a_{\textit{rleaf}(\mathcal{T})}\big\}, & \text{otherwise} \end{cases}$$

Clearly, if then $e(A \setminus W)$ is *any* enumeration of $A \setminus W$, the enumeration $e(A \setminus W), e(W)$ corresponds to an ordering of A satisfying the claim of Theorem VI.4.

We give the inductive definition:

(a) If $T = \{r\}$, then $W = \emptyset$ and $e(W)$ is the empty sequence.

(b) If the root r has a unique child t, then $e(W) := a_t, e(W')$ where

$$\begin{aligned} W' &:= W\big(\mathcal{A}_{\mathcal{T}_t}, \mathcal{T}_t.(B_{t'})_{t' \in T_t}\big) \\ &= \big\{a_{t'} \mid t' \in T_t \setminus \{t\}\big\}. \end{aligned}$$

(c) If the root r has at least two children, let $\ell(r)$ be the last child of r. Consider the tree $\mathcal{T}_1 := \mathcal{T} \setminus \mathcal{T}_{\ell(r)}$ (the subtree of $\mathcal{T}$ induced on the set $T \setminus T_{\ell(r)}$) and the subtree $\mathcal{T}_2 := \mathcal{T}_{\ell(r)}$ rooted at $\ell(r)$.

The substructures

$$\mathcal{A}_1 := \mathcal{A}_{\mathcal{T}_1} \quad \text{and} \quad \mathcal{A}_2 := \mathcal{A}_{\mathcal{T}_2},$$

of $\mathcal{A}$ have the smooth tree decompositions

$$\big(\mathcal{T}_1, (B_t)_{t \in T_1}\big) \quad \text{and} \quad \big(\mathcal{T}_2, (B_t)_{t \in T_2}\big).$$

Let

$$W_1 := W\big(\mathcal{A}_1, \mathcal{T}_1, (B_t)_{t \in T_1}\big)$$

and

$$W_2 := W\big(\mathcal{A}_2, \mathcal{T}_2, (B_t)_{t \in T_2}\big).$$

It is easy to see that

$$W = W_1 \,\dot\cup\, W_2 \,\dot\cup\, \{a_{\ell(r)}\}.$$

Applying the induction hypothesis to the tree decompositions $(\mathcal{T}_1, (B_t)_{t \in T_1})$ and $(\mathcal{T}_2, (B_t)_{t \in T_2})$ of $\mathcal{A}_1$ and $\mathcal{A}_2$, we get enumerations $e(W_1)$ and $e(W_2)$. We let

$$e(W) := e(W_1), a_{\ell(r)}, e(W_2).$$

We only verify (T2) and (T3) in case (c). By induction hypothesis, there are tree decomposition of $(\mathcal{A}_1, <^{W_1})$ and of $(\mathcal{A}_2, <^{W_2})$, both of width $\leq k + 5$, and whose bags at the root are

$$B_r \cup \{a_{f(r)}, a_{\textit{rleaf}(\mathcal{T}_1)}\}$$

and

$$B_{\ell(r)} \cup \{a_{f(\ell(r))}, a_{\textit{rleaf}(\mathcal{T}_2)}\}$$

respectively. (Here, we assume that $\mathcal{T}_2$ contains at least two nodes, otherwise the proof is easier.) Note that $a_{\ell(r)} \in B_{\ell(r)}$. We first join these tree decompositions by adding a new node t on the top of the two roots with bag

$$B_r \cup B_{\ell(r)} \cup \big\{a_{f(r)}, a_{\textit{rleaf}(\mathcal{T}_1)}, a_{f(\ell(r))}, a_{\textit{rleaf}(\mathcal{T}_2)}\big\}$$

of size $k+6$ as $|B_r \cup B_{\ell(r)}| = k+2$ by the smoothness of the tree decomposition $\big(\mathcal{T}, (B_t)_{t \in T}\big)$. Then, we already have a tree decomposition of $\big(\mathcal{A}, S(<^W)\big)$ of width $\leq k+5$ as the "new" tuples $(a_{\textit{rleaf}(\mathcal{T}_1)}, a_{\ell(r)})$ and $(a_{\ell(r)}, a_{f(\ell(r))})$ of $S(<^W)$ are realized in this bag. To satisfy (T2), we only need to put a root on the top of t with bag

$$B_r \cup \big\{a_{f(r)}, a_{\textit{rleaf}(\mathcal{T})}\big\}. \qquad \square$$