

ORDER-INVARIANT TYPES AND THEIR APPLICATIONS

PABLO BARCELÓ AND LEONID LIBKIN

Center for Semantic Web Research & Department of Computer Science, University of Chile
e-mail address: pbarcelo@dcc.uchile.cl

Laboratory for Foundations of Computer Science, School of Informatics, University of Edinburgh
e-mail address: libkin@inf.ed.ac.uk

ABSTRACT. Our goal is to show that the standard model-theoretic concept of types can be applied in the study of order-invariant properties, i.e., properties definable in a logic in the presence of an auxiliary order relation, but not actually dependent on that order relation. This is somewhat surprising since order-invariant properties are more of a combinatorial rather than a logical object. We provide two applications of this notion. One is a proof, from the basic principles, of a theorem by Courcelle stating that over trees, order-invariant MSO properties are expressible in MSO with counting quantifiers. The other is an analog of the Feferman-Vaught theorem for order-invariant properties.

1. INTRODUCTION

Invariant queries are an intriguing object that appear in the study of the expressive power of logics over finite models. The interest in them stems from the fact that to describe complexity classes by logical means, one often needs an underlying linear order. For instance, the Immerman-Vardi theorem characterizes polynomial time properties of graphs as those expressible in least-fixpoint logic in the presence of an order relation, cf. [10, 14]. However, the ordering can be chosen arbitrarily: its only goal is to enable the logic to simulate a Turing machine, which of course has the input on its tape in *some* order. That is, one needs an ordering to express a property, but it does not matter which order to use; any order would do. Properties expressed in this fashion are called *order-invariant*.

Since many results on capturing complexity classes require an ordering that is used in an invariant fashion, the notion is of interest. Before studying it for expressive logics like least-fixpoint, one would want to understand its behavior in simpler settings, like first-order logic (FO). Several attempts to do so, however, show that the notion is much harder to deal with than it initially appears.

Key words and phrases: finite model theory; invariance; types.

Barceló is funded by the Millennium Nucleus Center for Semantic Web Research under Grant NC120004. Libkin is supported by EPSRC grants J015377 and M025268. We are grateful to the reviewers for suggesting helpful pointers to the literature and modifications to improve the readability of the paper.

To start with, it only makes sense for *finite* structures: over infinite structures, a simple application of the interpolation theorem shows that any order-invariant property can be expressed without the order itself. But for finite structures, order-invariance does add power. This was noticed (although not published) by Yuri Gurevich a while ago, but by now this is a textbook result [14]. The most commonly used separating property is checking whether the number of atoms of a Boolean algebra is even: it cannot be done in FO, but can be if an arbitrary order on elements of the Boolean algebra is added. More complex examples show that the separation continues to hold if order is replaced by weaker devices such as the ability to choose an element from a set [19] or the successor relation [20].

These observations led to the study of the power of order-invariant properties in logics such as FO and monadic second-order logic (MSO). Despite much effort, we still know relatively little about order-invariant properties, and results that we know typically require a very significant effort (see [21, 22] for overviews). We do know nonetheless that order-invariant properties are local, much like FO-definable properties themselves [11], that over some tame structures such as words and trees, order-invariance does not add power [3, 18], and there are results that extend invariance beyond order, for instance to arithmetic predicates [23, 1], or prove strong separation results for auxiliary relations that are slightly weaker than order [15]. For the more powerful logic MSO we know that order-invariance on tame structures such as trees boils down to adding counting to the logic [7].

One of the reasons that the progress in understanding order-invariance is rather slow is the *lack of logical tools* – indeed, the set of order-invariant properties is not really a logic, i.e., it is not recursively enumerable. However, perhaps somewhat surprisingly in view of this observation, some logical tools can be adapted to deal with order-invariance. Showing this is our goal. We take the standard model-theoretic concept of *types* (i.e., sets of formulae of a logic or a fragment of a logic that hold in a given structure), which play a prominent role in both classical and finite model theory [4, 14], and demonstrate their usefulness in the study of order-invariance.

Specifically, we do the following.

- (1) We define a notion of *order-invariant types* that extends the notion of types to the order-invariant setting and study its basic properties.
- (2) We show that, despite order-invariant properties not forming a logic, a logic-based notion of order-invariant types can actually be useful. We provide two applications of it.
 - (a) First, we provide a proof, from the basic principles, of a result by Courcelle [7] saying that over trees, order-invariant MSO properties are the ones expressible in MSO with counting quantifiers. This was reported in our conference paper [2] but the proof was never published.
 - (b) Second, we prove an analog of the Feferman-Vaught theorem [8, 16] for order-invariant properties, and use it to extend the list of known classes where order-invariance does not increase the expressive power.

While not claiming a breakthrough, the goal of this note is to show that standard model-theoretic techniques are applicable in this notoriously difficult area, and perhaps offer a new avenue of attack on a host of unsolved problems related to order-invariance.

Organization. Basic concepts are defined in Section 2. In Section 3 we define order-invariant types and study their basic properties. The proof of Courcelle’s theorem based on

order-invariant types is given in Section 4, and the order-invariant version of the Feferman-Vaught theorem is given in Section 5.

2. PRELIMINARIES

We now present basic background concepts, following [14]. We assume familiarity with first-order logic (FO) and with its extension with monadic second-order quantification known as monadic second-order logic (MSO). First-order variables are denoted x, y, z, \dots , while second-order variables are denoted X, Y, Z, \dots . We assume that vocabularies are relational, i.e., they contain only relation and constant symbols. All structures are assumed to be finite, and will be denoted by letters $\mathcal{A}, \mathcal{B}, \dots$; the domain of a structure \mathcal{A} will be denoted by $\text{dom}(\mathcal{A})$.

The *quantifier rank* $\text{qr}(\varphi)$ of a formula (FO or MSO) is the depth of quantifier nesting in φ . Up to logical equivalence, there are only finitely many different formulae of quantifier rank k (FO or MSO) for each given vocabulary.

With each structure \mathcal{A} of vocabulary σ we associate its *rank- k FO type*

$$\text{tp}_{\text{FO}}^k(\mathcal{A}) = \{\varphi \mid \mathcal{A} \models \varphi \text{ and } \text{qr}(\varphi) \leq k\},$$

where φ ranges over FO sentences over σ . Similarly we define $\text{tp}_{\text{MSO}}^k(\mathcal{A})$. From the previous paragraph, both $\text{tp}_{\text{FO}}^k(\mathcal{A})$ and $\text{tp}_{\text{MSO}}^k(\mathcal{A})$ can be assumed to be finite. There are finitely many rank- k types, and for each rank- k (FO or MSO) type τ there is a sentence φ_τ of quantifier rank k in the logic that defines it, i.e., $\mathcal{A} \models \varphi_\tau$ if and only if $\text{tp}_{\text{FO}}^k(\mathcal{A}) = \tau$ (or $\text{tp}_{\text{MSO}}^k(\mathcal{A}) = \tau$). In particular, $\varphi_\tau = \bigwedge_{\varphi \in \text{tp}_{\text{FO}}^k(\mathcal{A})} \varphi$. We thus associate types with formulas that define them, and simply write τ instead of φ_τ . Every FO sentence of quantifier rank k is equivalent to a disjunction of rank- k FO types (and likewise for MSO).

EF games. For FO and MSO logical equivalence up to quantifier rank k can be captured using Ehrenfeucht-Fraïssé (EF) games. The game is played in two structures \mathcal{A} and \mathcal{B} , over the same vocabulary σ , by two players, the *spoiler* and the *duplicator*, for $k \geq 0$ of *rounds*. In round i the spoiler selects a structure, say \mathcal{A} , and an element $a_i \in \text{dom}(\mathcal{A})$; the duplicator responds by selecting an element b_i in the other structure, in this case $b_i \in \text{dom}(\mathcal{B})$. The duplicator *wins* if the mapping $a_i \mapsto b_i$, for $i \leq k$ defines a partial isomorphism between \mathcal{A} and \mathcal{B} . That is, for purely relational vocabulary this map is an isomorphism of substructures induced by $\{a_1, \dots, a_k\}$ and $\{b_1, \dots, b_k\}$. If constant symbols are involved, these sets are expanded by the interpretation of constants in \mathcal{A} and \mathcal{B} , and the interpretation of each constant symbol c in \mathcal{A} must be mapped into the interpretation of that symbol in \mathcal{B} .

The duplicator has a winning strategy in the k -round game if he wins in k rounds no matter how the spoiler plays. It is well known that this happens if and only if \mathcal{A} and \mathcal{B} agree on all FO sentences of quantifier rank up to k , and we write $\mathcal{A} \equiv_k^{\text{FO}} \mathcal{B}$ to denote this. In particular, $\mathcal{A} \equiv_k^{\text{FO}} \mathcal{B}$ iff $\text{tp}_{\text{FO}}^k(\mathcal{A}) = \text{tp}_{\text{FO}}^k(\mathcal{B})$.

An extension of the Ehrenfeucht-Fraïssé game also permits us to determine whether two structures have the same MSO type. This extension permits, in addition to the usual moves, also set moves, i.e., the spoiler can play a subset of a structure, say $U_i \subseteq \text{dom}(\mathcal{A})$, and the duplicator must then respond with a set in the other structure, i.e., $V_i \subseteq \text{dom}(\mathcal{B})$. The winning condition is that the usual element moves form a partial isomorphism of substructures expanded with predicates for the set moves. In particular, if a_i, U_j are an element and a set moves in \mathcal{A} , and b_i, V_j are the responses in \mathcal{B} , then $a_i \in U_j$ iff $b_i \in V_j$.

As for FO, the spoiler has a winning strategy in this k -round MSO game iff \mathcal{A} and \mathcal{B} agree on all MSO sentences of quantifier rank up to k , i.e., iff $\text{tp}_{\text{MSO}}^k(\mathcal{A}) = \text{tp}_{\text{MSO}}^k(\mathcal{B})$. In this case we write $\mathcal{A} \equiv_k^{\text{MSO}} \mathcal{B}$.

\mathfrak{C} -invariant sentences. Let \mathcal{L} be either FO or MSO and assume that σ and σ' are disjoint vocabularies. Consider structures \mathcal{A} and \mathcal{B} over σ and σ' , respectively, such that \mathcal{A} and \mathcal{B} share the same domain (i.e., $\text{dom}(\mathcal{A}) = \text{dom}(\mathcal{B})$). We denote by $(\mathcal{A}, \mathcal{B})$ the structure over $\sigma \cup \sigma'$ whose domain coincides with that of \mathcal{A} and \mathcal{B} , and the interpretation of σ (resp. σ') is inherited from \mathcal{A} (resp. \mathcal{B}).

Assume that \mathfrak{C} and \mathfrak{D} are classes of structures over σ' and σ , respectively. An \mathcal{L} sentence φ over $\sigma \cup \sigma'$ is *\mathfrak{C} -invariant over \mathfrak{D}* , if for each structure $\mathcal{A} \in \mathfrak{D}$ and any two structures $\mathcal{B}_1, \mathcal{B}_2 \in \mathfrak{C}$ with the same domain than \mathcal{A} , the following holds:

$$(\mathcal{A}, \mathcal{B}_1) \models \varphi \iff (\mathcal{A}, \mathcal{B}_2) \models \varphi.$$

We denote by $(\mathcal{L} \cup \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}$ the set of \mathfrak{C} -invariant \mathcal{L} sentences over \mathfrak{D} . We omit the superscript when \mathfrak{D} is the class of all structures over σ .

A \mathfrak{C} -invariant \mathcal{L} sentence φ over \mathfrak{D} defines a *query* Q_φ which is a set of structures in \mathfrak{D} as follows:

$$\begin{aligned} \mathcal{A} \in Q_\varphi &\iff (\mathcal{A}, \mathcal{B}) \models \varphi \text{ for some } \mathcal{B} \in \mathfrak{C} \text{ with } \text{dom}(\mathcal{A}) = \text{dom}(\mathcal{B}) \\ &\iff (\mathcal{A}, \mathcal{B}) \models \varphi \text{ for all } \mathcal{B} \in \mathfrak{C} \text{ with } \text{dom}(\mathcal{A}) = \text{dom}(\mathcal{B}). \end{aligned}$$

The most important case for us is when \mathfrak{C} is the class of linear orders. We then write $<$ instead of \mathfrak{C} , and denote by $(\mathcal{L} + \{<\})_{\text{inv}}^{\mathfrak{D}}$ the set of *$<$ -invariant \mathcal{L} sentences over \mathfrak{D}* . In other words, $(\mathcal{L} + \{<\})_{\text{inv}}^{\mathfrak{D}}$ consists of all \mathcal{L} sentences φ over vocabulary $\sigma \cup \{<\}$ such that for every structure $\mathcal{A} \in \mathfrak{D}$, and any two linear orders $<_1, <_2$ interpreting $<$ over $\text{dom}(\mathcal{A})$, we have:

$$(\mathcal{A}, <_1) \models \varphi \iff (\mathcal{A}, <_2) \models \varphi.$$

We say that $<$ -invariant \mathcal{L} *collapses to \mathcal{L}* over \mathfrak{D} if for every sentence φ in $(\mathcal{L} + \{<\})_{\text{inv}}^{\mathfrak{D}}$, the query Q_φ is definable in \mathcal{L} over \mathfrak{D} . It is known, for instance, that over words and trees, $<$ -invariant FO collapses to FO, see [3, 18]. But sometimes invariance adds power, as the example below demonstrates.

Example 1 A linear order can be used to define an MSO sentence φ_{even} that checks if the domain has even cardinality. Indeed, all one needs to do is to check the existence of a subset S of the ordering that corresponds to even positions (i.e., S consists of every other element starting from the second element of the ordering $<$) such that S also contains its last element. These are easily expressible in MSO, and hence this is an $(\text{MSO} + \{<\})_{\text{inv}}$ sentence, since it does not matter which linear order to use.

Using exactly the same idea, we can define a sentence $\mathbf{Q}_2x \psi(x)$ checking if the number of elements a satisfying a given formula $\psi(x)$ is divisible by 2, and in fact a sentence $\mathbf{Q}_px \psi(x)$ checking if the number of elements satisfying $\psi(x)$ is divisible by p . These are known as *counting quantifiers*, and they will be important for us in the next sections.

Finally, assume that the vocabulary σ is empty. Then the $(\text{MSO} + \{<\})_{\text{inv}}$ sentence φ_{even} is not definable in MSO alone, which shows that even over empty vocabularies, $<$ -invariant MSO does *not* collapse to MSO. The same property shows that over the class of Boolean algebras, $<$ -invariant FO does not collapse to FO (since in Boolean algebras one can mimic MSO quantification, c.f., [14]). \square

3. \mathfrak{C} -INVARIANT TYPES

Let \mathcal{L} be either FO or MSO. As before, σ and σ' are disjoint vocabularies, \mathfrak{C} is a class of structures over σ' and \mathfrak{D} a class of structures over σ . With each structure $\mathcal{A} \in \mathfrak{D}$ over σ we define its *rank- k \mathfrak{C} -invariant \mathcal{L} type over \mathfrak{D}* , which we denote by $\text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A})$, to be the set of all $(\mathcal{L} \cup \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}$ sentences φ of quantifier rank at most k such that $\mathcal{A} \in Q_\varphi$.

As in the case of rank- k \mathcal{L} types, each rank- k \mathfrak{C} -invariant \mathcal{L} type over \mathfrak{D} is definable by an $(\mathcal{L} \cup \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}$ sentence of quantifier rank at most k . This is because the conjunction of all (finitely many, up to equivalence) \mathfrak{C} -invariant \mathcal{L} sentences in $\text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A})$ is an $(\mathcal{L} \cup \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}$ sentence of quantifier rank k (because invariant sentences are closed under Boolean connectives). It follows that for any vocabularies σ and σ' , and every k , there are only finitely many rank- k \mathfrak{C} invariant \mathcal{L} types of \mathfrak{D} , and every rank- k \mathfrak{C} invariant \mathcal{L} sentence over \mathfrak{D} is equivalent to a disjunction of such rank- k types.

In this section, we provide a combinatorial characterization of rank- k invariant types which is crucial for our results. It can be described in terms of finite sequences of transformations that either replace a pair of structures $(\mathcal{A}, \mathcal{B})$ with another one of the same rank- k type (the usual, not invariant), or simply replacing the second component of the structure.

More precisely, let $(\mathcal{A}, \mathcal{B})$ and $(\mathcal{A}', \mathcal{B}')$ be two pairs of structures sharing the domain (i.e., $\text{dom}(\mathcal{A}) = \text{dom}(\mathcal{B})$ and $\text{dom}(\mathcal{A}') = \text{dom}(\mathcal{B}')$) and let $\mathcal{L}, \mathfrak{C}, \mathfrak{D}$ be as above. We write

$$(\mathcal{A}, \mathcal{B}) \sim_k^{\mathcal{L}, \mathfrak{D}, \mathfrak{C}} (\mathcal{A}', \mathcal{B}')$$

if either $\mathcal{A} = \mathcal{A}'$, or $\text{tp}_{\mathcal{L}}^k(\mathcal{A}, \mathcal{B}) = \text{tp}_{\mathcal{L}}^k(\mathcal{A}', \mathcal{B}')$. When $(\mathcal{A}', \mathcal{B}')$ is reachable from $(\mathcal{A}, \mathcal{B})$ by a finite sequence of $\sim_k^{\mathcal{L}, \mathfrak{D}, \mathfrak{C}}$ steps, we say that $(\mathcal{A}', \mathcal{B}')$ is a *k -flip* of $(\mathcal{A}, \mathcal{B})$ under $(\mathcal{L}, \mathfrak{D}, \mathfrak{C})$. The following simple observation establishes that k -flips preserve the k -invariant type of \mathcal{A} :

Lemma 3.1. If $(\mathcal{A}', \mathcal{B}')$ is a k -flip of $(\mathcal{A}, \mathcal{B})$ under $(\mathcal{L}, \mathfrak{D}, \mathfrak{C})$, then $\text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}) = \text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}')$.

Proof. We show that if $(\mathcal{A}, \mathcal{B}) \sim_k^{\mathcal{L}, \mathfrak{D}, \mathfrak{C}} (\mathcal{A}', \mathcal{B}')$ then $\text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}) = \text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}')$. The lemma follows then by a straightforward induction on the length of the finite sequence of $\sim_k^{\mathcal{L}, \mathfrak{D}, \mathfrak{C}}$ steps that constitutes a k -flip. The case when $\mathcal{A} = \mathcal{A}'$ is trivial. Assume then that $\mathcal{A} \neq \mathcal{A}'$ but $\text{tp}_{\mathcal{L}}^k(\mathcal{A}, \mathcal{B}) = \text{tp}_{\mathcal{L}}^k(\mathcal{A}', \mathcal{B}')$. Consider an arbitrary sentence $\varphi \in \text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A})$. Then $\mathcal{A} \in Q_\varphi$, and therefore $(\mathcal{A}, \mathcal{B}) \models \varphi$. But since $\text{tp}_{\mathcal{L}}^k(\mathcal{A}, \mathcal{B}) = \text{tp}_{\mathcal{L}}^k(\mathcal{A}', \mathcal{B}')$, we conclude that $(\mathcal{A}', \mathcal{B}') \models \varphi$, which implies that $\mathcal{A}' \in Q_\varphi$. Therefore, $\text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}) \subseteq \text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}')$. The proof that $\text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}') \subseteq \text{tp}_{(\mathcal{L}+\{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A})$ is symmetric. \square

The notion of k -flip describes the equivalence of invariant types. More formally, for $\mathcal{A} \in \mathfrak{D}$ and $\mathcal{B} \in \mathfrak{C}$ such that $\text{dom}(\mathcal{A}) = \text{dom}(\mathcal{B})$ let us define:

$$\psi_{(\mathcal{A}, \mathcal{B})}^k = \bigvee \{ \text{tp}_{\mathcal{L}}^k(\mathcal{A}', \mathcal{B}') \mid \mathcal{A}' \in \mathfrak{D}, \mathcal{B}' \in \mathfrak{C}, \text{dom}(\mathcal{A}) = \text{dom}(\mathcal{B}), \\ \text{and } (\mathcal{A}', \mathcal{B}') \text{ is a } k\text{-flip of } (\mathcal{A}, \mathcal{B}) \text{ under } (\mathcal{L}, \mathfrak{D}, \mathfrak{C}) \}.$$

We write $\psi_{(\mathcal{A}, \mathcal{B})}$ instead of $\psi_{(\mathcal{A}, \mathcal{B})}^k$ when k is clear from the context. Then we can establish the following:

Proposition 3.2. Let $k \geq 0$ and assume that $\mathcal{A} \in \mathfrak{D}$ and $\mathcal{B} \in \mathfrak{C}$ are structures such that $\text{dom}(\mathcal{A}) = \text{dom}(\mathcal{B})$. Then for every $\mathcal{A}' \in \mathfrak{D}$ and $\mathcal{B}' \in \mathfrak{C}$ such that $\text{dom}(\mathcal{A}') = \text{dom}(\mathcal{B}')$ the following are equivalent:

- (1) $(\mathcal{A}', \mathcal{B}') \models \psi_{(\mathcal{A}, \mathcal{B})}$.
- (2) $(\mathcal{A}', \mathcal{B}')$ is a k -flip of $(\mathcal{A}, \mathcal{B})$ under $(\mathcal{L}, \mathfrak{D}, \mathfrak{C})$.

Furthermore, $\psi_{(\mathcal{A}, \mathcal{B})}$ is an $(\mathcal{L} \cup \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}$ sentence of quantifier rank at most k and $Q_{\psi_{(\mathcal{A}, \mathcal{B})}}$ defines $\text{tp}_{(\mathcal{L} + \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A})$. That is, for another structure $\mathcal{A}' \in \mathfrak{D}$, we have $\text{tp}_{(\mathcal{L} + \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}) = \text{tp}_{(\mathcal{L} + \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}')$ iff $\mathcal{A}' \in Q_{\psi_{(\mathcal{A}, \mathcal{B})}}$.

Proof. Assume first that $(\mathcal{A}', \mathcal{B}') \models \psi_{(\mathcal{A}, \mathcal{B})}$. Then $\text{tp}_{\mathcal{L}}^k(\mathcal{A}', \mathcal{B}') = \text{tp}_{\mathcal{L}}^k(\mathcal{A}'', \mathcal{B}'')$, for some $\mathcal{A}'' \in \mathfrak{D}$ and $\mathcal{B}'' \in \mathfrak{C}$ such that $(\mathcal{A}'', \mathcal{B}'')$ is a k -flip of $(\mathcal{A}, \mathcal{B})$ under $(\mathcal{L}, \mathfrak{D}, \mathfrak{C})$. But then $(\mathcal{A}', \mathcal{B}')$ is also a k -flip of $(\mathcal{A}, \mathcal{B})$ under $(\mathcal{L}, \mathfrak{D}, \mathfrak{C})$. Assume, on the other hand, that $(\mathcal{A}', \mathcal{B}')$ is a k -flip of $(\mathcal{A}, \mathcal{B})$ under $(\mathcal{L}, \mathfrak{D}, \mathfrak{C})$. Then $(\mathcal{A}', \mathcal{B}') \models \psi_{(\mathcal{A}, \mathcal{B})}$ since $(\mathcal{A}', \mathcal{B}') \models \text{tp}_{\mathcal{L}}^k(\mathcal{A}', \mathcal{B}')$.

Clearly, $\psi_{(\mathcal{A}, \mathcal{B})}$ is of quantifier rank at most k . We prove next that it is an $(\mathcal{L} \cup \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}$ sentence. Let $\mathcal{A}' \in \mathfrak{D}$ and $\mathcal{B}_1, \mathcal{B}_2 \in \mathfrak{C}$ such that $\text{dom}(\mathcal{A}) = \text{dom}(\mathcal{B}_1) = \text{dom}(\mathcal{B}_2)$, and assume that $(\mathcal{A}', \mathcal{B}_1) \models \psi_{(\mathcal{A}, \mathcal{B})}$. Then, from the previous characterization we have that $(\mathcal{A}', \mathcal{B}_1)$ is a k -flip of $(\mathcal{A}, \mathcal{B})$ under $(\mathcal{L}, \mathfrak{D}, \mathfrak{C})$. Therefore, $(\mathcal{A}', \mathcal{B}_2)$ is also a k -flip of $(\mathcal{A}, \mathcal{B})$ under $(\mathcal{L}, \mathfrak{D}, \mathfrak{C})$, from which we conclude that $(\mathcal{A}', \mathcal{B}_2) \models \psi_{(\mathcal{A}, \mathcal{B})}$.

Finally, we prove that $Q_{\psi_{(\mathcal{A}, \mathcal{B})}}$ defines $\text{tp}_{(\mathcal{L} + \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A})$. Let \mathcal{A}' be a structure in \mathfrak{D} . Assume first that $\text{tp}_{(\mathcal{L} + \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}) = \text{tp}_{(\mathcal{L} + \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}')$. Since $\psi_{(\mathcal{A}, \mathcal{B})}$ is an $(\mathcal{L} \cup \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}$ sentence of quantifier rank k and $(\mathcal{A}, \mathcal{B}) \models \psi_{(\mathcal{A}, \mathcal{B})}$ for any \mathcal{B} such that $\text{dom}(\mathcal{A}) = \text{dom}(\mathcal{B})$, we conclude that $(\mathcal{A}', \mathcal{B}') \models \psi_{(\mathcal{A}, \mathcal{B})}$ for any \mathcal{B}' such that $\text{dom}(\mathcal{A}') = \text{dom}(\mathcal{B}')$. Therefore, $\mathcal{A}' \in Q_{\psi_{(\mathcal{A}, \mathcal{B})}}$. Assume, on the other hand, that $\mathcal{A}' \in Q_{\psi_{(\mathcal{A}, \mathcal{B})}}$. Therefore, for any $\mathcal{B}' \in \mathfrak{C}$ such that $\text{dom}(\mathcal{A}') = \text{dom}(\mathcal{B}')$ we have that $(\mathcal{A}', \mathcal{B}') \models \psi_{(\mathcal{A}, \mathcal{B})}$. But then $(\mathcal{A}', \mathcal{B}')$ is a k -flip of $(\mathcal{A}, \mathcal{B})$ under $(\mathcal{L}, \mathfrak{D}, \mathfrak{C})$. We conclude that $\text{tp}_{(\mathcal{L} + \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}) = \text{tp}_{(\mathcal{L} + \{\mathfrak{C}\})_{\text{inv}}^{\mathfrak{D}}}^k(\mathcal{A}')$ from Lemma 3.1. \square

4. COURCELLE'S THEOREM FROM INVARIANT TYPES

We have seen in Example 1 that counting quantifiers \mathbf{Q}_p can be defined in order-invariant MSO. Such quantifiers extend MSO by the following formation rule: if $\psi(x, \bar{y})$ is a formula, then $\varphi(\bar{y}) = \mathbf{Q}_p x \psi(x, \bar{y})$ is a formula. We have $\mathcal{A} \models \varphi(\bar{a})$ if $|\{a' \in \text{dom}(\mathcal{A}) \mid \mathcal{A} \models \psi(a', \bar{a})\}|$ is divisible by p . MSO extended with such quantifiers for all p is referred to as *counting* MSO, or CMSO. What Example 1 tells us is that CMSO is definable in order-invariant MSO. Courcelle's result from [7] says that over trees, the two coincide.

Courcelle's proof was quite involved; it used graph grammars and an algebraic approach to recognizability. We now provide a much simpler proof that uses tree automata techniques based on the invariant types machinery. The simplest way to prove that over trees, MSO captures tree automata, is to define, for each k , a deterministic tree automaton that assigns each node the rank- k MSO type of the subtree rooted at it, cf. [14]. Our proof extends this to the order-invariant setting: we show how to define an automaton that computes order-invariant rank- k types, and then prove that such an automaton can be encoded in CMSO over trees.

Trees and tree automata. We consider rooted trees with oriented edges. To be able to take advantage of automata machinery, we define them as unranked trees (cf. [6]) without a sibling ordering. More precisely, an *unranked tree domain* D is a prefix-closed finite set of words of positive natural numbers such that $s \cdot i \in D$ implies $s \cdot j \in D$ for each $1 \leq j < i$. An *unranked tree* over a finite alphabet Σ is a structure $T = (D, \prec_{\text{desc}}, (P_a)_{a \in \Sigma})$, where $D = \text{dom}(T)$ is an unranked tree domain, \prec_{desc} is interpreted as the descendant relation (i.e., $s \prec_{\text{desc}} s \cdot s'$ for each $s \cdot s' \in D$ such that s' is nonempty), and P_a as the set of nodes labeled a , for each $a \in \Sigma$. As usual, we require the P_a 's to form a partition of D , i.e., each element of the domain is assigned a unique label in Σ . The empty word will be denoted by ε ; hence $\varepsilon \in D$ is the root of T . In order to avoid notation clutter, throughout this section we simply write $(\text{MSO} + \{<\})_{\text{inv}}$ for $(\text{MSO} + \{<\})_{\text{inv}}^{\text{Trees}}$, where **Trees** is the set of all unranked trees.

Courcelle's theorem [7] says the following.

Theorem 4.1. [7] A set of unranked trees is definable in $(\text{MSO} + \{<\})_{\text{inv}}$ iff it is definable in CMSO.

As the first observation towards the proof, we note that order-invariance can be replaced by sibling-order invariance. A sibling order on a tree is a binary relation \prec_{sb} such that $s' \prec_{\text{sb}} s''$ implies that $s' = s \cdot i$ and $s'' = s \cdot j$ for some node s and distinct numbers i and j , and that on the set of all the children of every node s (i.e., $\{s \cdot i \mid s \cdot i \in D\}$), the relation \prec_{sb} is a linear order. We denote the extension of the unranked tree T with sibling order \prec_{sb} by (T, \prec_{sb}) , and call (T, \prec_{sb}) a \prec_{sb} -ordered unranked tree. We slightly abuse notation and denote simply by \prec_{sb} the class \mathfrak{C} of structures that represent sibling-orders on unranked trees. We can restrict our attention to sibling-order invariance due to the following. Indeed, it is well known that a linear order can be defined from \prec_{sb} , and of course vice versa, a linear order defines a sibling order. Since the two are inter-definable, we have the following (as before, we write $(\text{MSO} + \{\prec_{\text{sb}}\})_{\text{inv}}$ for $(\text{MSO} + \{\prec_{\text{sb}}\})_{\text{inv}}^{\text{Trees}}$):

Lemma 4.2. A set of unranked trees is definable in $(\text{MSO} + \{<\})_{\text{inv}}$ iff it is definable in $(\text{MSO} + \{\prec_{\text{sb}}\})_{\text{inv}}$.

Sibling order allows us to bring in tree automata over \prec_{sb} -ordered unranked trees. A *tree automaton* (TA) [6] is a tuple $\mathcal{N} = (\Sigma, Q, F, \delta)$, where Σ is a finite alphabet, Q is a set of states, $F \subseteq Q$ is the set of final states, and $\delta : Q \times \Sigma \rightarrow 2^{Q^*}$ is the transition function such that $\delta(q, a)$ is a regular language over Q for every $q \in Q$ and $a \in \Sigma$. A *run* of a TA \mathcal{N} on a \prec_{sb} -ordered unranked tree (T, \prec_{sb}) with domain D is a function $\rho : D \rightarrow Q$ such that, for every element $s \in D$ labeled a with children $s_1 \prec_{\text{sb}} \dots \prec_{\text{sb}} s_n$, the word $\rho(s_1) \dots \rho(s_n)$ is in $\delta(\rho(s), a)$ (if s is a leaf labeled a , then the condition enforces that the empty word belongs to $\delta(\rho(s), a)$). The \prec_{sb} -ordered tree is *accepted* by \mathcal{N} if there is a run ρ of \mathcal{N} on (T, \prec_{sb}) such that $\rho(\varepsilon) \in F$ (i.e., the root is in a state in F). A TA \mathcal{N} is *deterministic* if for each $q, q' \in Q$ such that $q \neq q'$ and $a \in \Sigma$, there is no word that belongs to both $\delta(q, a)$ and $\delta(q', a)$. A set of \prec_{sb} -ordered unranked trees is *regular* if and only if it is precisely the set of trees accepted by a TA. It is well-known that TAs can be determinized, that is, a set of \prec_{sb} -ordered unranked trees is regular iff it is accepted by a deterministic TA.

We prove here that the sets of unranked trees that are $(\text{MSO} + \{\prec_{\text{sb}}\})_{\text{inv}}$ -definable can be recognized by a particular class of TAs, which we call *invariant*. This is done by extending traditional techniques used to establish connections between MSO definability and automata recognizability over words and trees.

Definition 4.1. A TA $\mathcal{N} = (\Sigma, Q, F, \delta)$ is \prec_{sb} -invariant if for each state $q \in Q$ and symbol $a \in \Sigma$, every permutation of a word in $\delta(q, a)$ is also in $\delta(q, a)$. (Thus, a run of \mathcal{N} on a \prec_{sb} -ordered unranked tree does not depend on the actual interpretation of \prec_{sb}).

The next lemma establishes the desired connection between $(\text{MSO} + \{\prec_{\text{sb}}\})_{\text{inv}}$ -definability and \prec_{sb} -invariant TA recognizability:

Lemma 4.3. Let S be an $(\text{MSO} + \{\prec_{\text{sb}}\})_{\text{inv}}$ -definable set of unranked trees. There is a deterministic \prec_{sb} -invariant TA \mathcal{N} for which S is precisely the set of unranked trees T such that some \prec_{sb} -ordered extension (T, \prec_{sb}) of T is accepted by \mathcal{N} (or, equivalently, each \prec_{sb} -ordered extension (T, \prec_{sb}) of T is accepted by \mathcal{N}).

The automaton, as we already explained, will be computing invariant types of subtrees in its run. Before proving the lemma, we show how Courcelle's theorem easily follows from it. One direction is immediate from the observation made in Example 1: we saw that counting quantifiers can be expressed in MSO using an order relation. For the other direction, we make use of a description of regular languages closed under permutations. Let $S_{k,p}$, for $k, p \geq 0$, be the semilinear set $\{k + np \mid n \in \mathbb{N}\}$. For an alphabet $\Sigma = \{a_1, \dots, a_r\}$, consider the Parikh map $\Pi : \Sigma^* \rightarrow \mathbb{N}^r$ where the i -th component of $\Pi(w)$ is the number of occurrences of a_i in w . Then:

Lemma 4.4. A regular language $L \subseteq \Sigma^*$ is closed under permutation iff there exists a finite family \mathfrak{S} of r -tuples of sets of the form $S_{k,p}$, such that for each word w over Σ it is the case that $w \in L$ iff for some $(S_1, \dots, S_r) \in \mathfrak{S}$, we have $\Pi(w) \in S_1 \times \dots \times S_r$.

This can be obtained from results in [9] and is also an immediate consequence of Pillig's normal form [12] which describes permutations of words in regular languages. We also provide a simple direct model-theoretic proof of this result in the appendix.

Assume now that we have a set S of trees that is $(\text{MSO} + \{<\})_{\text{inv}}$ -definable. By Lemma 4.2, it is $(\text{MSO} + \{\prec_{\text{sb}}\})_{\text{inv}}$ -definable, and by Lemma 4.3, there is a \prec_{sb} -invariant TA \mathcal{N} for which S is precisely the set of unranked trees T for which there is a sibling-order \prec_{sb} over $\text{dom}(T)$ such that (T, \prec_{sb}) is accepted by \mathcal{N} (or, equivalently, for each sibling-order \prec_{sb} over $\text{dom}(T)$ it is the case that (T, \prec_{sb}) is accepted by \mathcal{N}). We construct a CMSO sentence φ_S that precisely defines those unranked trees T . This can be done by using standard techniques for translating from tree automata into MSO (see, e.g., [17]). In particular, φ_S expresses the existence of an accepting run of \mathcal{N} over some \prec_{sb} -ordered extension of T . That is, φ_S expresses that there is an assignment of states of \mathcal{N} to the nodes of T that respects the transition function and assigns a final state to the root of T . The only problem here is that the sentence φ_S is defined over the unranked tree T , and hence there is no \prec_{sb} -order available to check whether the transitions performed by the run of \mathcal{N} are valid. However, we know that \mathcal{N} is \prec_{sb} -invariant, and, therefore, that each transition of the form $\delta(q, a)$ in \mathcal{N} is described by a regular language L that is closed under permutation. From Lemma 4.4, in order to check whether a word w belongs to L we can simply check whether $\Pi(w)$ belongs to some r -tuple of sets of the form $S_{k,p}$ in \mathfrak{S} . This can clearly be defined with a CMSO formula since sets of the form $S_{k,p}$ are semilinear. Hence, Courcelle's theorem follows.

Thus, it remains to prove Lemma 4.3:

Proof. Assume that S is definable by sentence φ in $(\text{MSO} + \{\prec_{\text{sb}}\})_{\text{inv}}$ over finite alphabet Σ . Let $k \geq 1$ be the quantifier rank of φ . We construct a deterministic \prec_{sb} -invariant TA

\mathcal{N} over alphabet Σ such that the unique run of \mathcal{N} on an arbitrary \prec_{sb} -ordered extension (T, \prec_{sb}) of an unranked tree T labels the root of T with $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(T)$.

Let \mathcal{T} be the set of all $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(T)$, for T an unranked tree over Σ . Assume that T_1, \dots, T_p are unranked trees over Σ . We denote by $\mathcal{A}(T_1, \dots, T_p)$ the structure over vocabulary $(P_\tau)_{\tau \in \mathcal{T}}$ whose domain is $\{1, \dots, p\}$ and element i ($1 \leq i \leq p$) belongs to P_τ ($\tau \in \mathcal{T}$) iff $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(T_i) = \tau$. Notice then that the interpretation of the P_τ 's defines a partition over the domain of $\mathcal{A}(T_1, \dots, T_p)$. Let us also denote by $a(T_1, \dots, T_p)$ the unranked tree over Σ that has a root labeled a and trees T_1, \dots, T_p hanging from this root. The following claim is crucial for our construction of TA \mathcal{N} :

Claim 4.5. For every $a \in \Sigma$ and trees T_1, \dots, T_p over Σ , the type $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(T_1, \dots, T_p))$ is uniquely determined by $\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(T_1, \dots, T_p))$.

Proof. We slightly abuse notation and say that a sibling-ordered tree (resp., a word) is a k -flip of another sibling-ordered tree (resp., word), but formally mean that it is a k -flip under $(\text{MSO}, \text{Trees}, \prec_{\text{sb}})$ (resp., under $(\text{MSO}, \text{All}, \prec)$, where All is the set of all structures over the given vocabulary).

Let $T_1, \dots, T_p, T_{p+1}, \dots, T_r$ be unranked trees over alphabet Σ . We show that for each $a \in \Sigma$ it is the case that:

$$\begin{aligned} \text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(T_1, \dots, T_p)) = \text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(T_{p+1}, \dots, T_r)) &\implies \\ \text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(T_1, \dots, T_p)) = \text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(T_{p+1}, \dots, T_r)). & \end{aligned}$$

For each $\tau \in \mathcal{T}$, let T_τ be an arbitrary unranked tree such that $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(T_\tau) = \tau$. Assume that for each $1 \leq j \leq r$, $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(T_j) = \tau_j$. Then, from Proposition 3.2, for each $1 \leq j \leq r$ there exist sibling-orders $\prec_{\text{sb}}^{j,1}$ and $\prec_{\text{sb}}^{j,2}$ on $\text{dom}(T_{\tau_j})$ and $\text{dom}(T_j)$, respectively, such that:

$$(T_{\tau_j}, \prec_{\text{sb}}^{j,1}) \text{ is a } k\text{-flip of } (T_j, \prec_{\text{sb}}^{j,2}). \quad (4.1)$$

It is then possible to conclude from (4.1) that there exist sibling-orders \prec_{sb}^1 and \prec_{sb}^2 such that:

$$(a(T_{\tau_1} \cdots T_{\tau_p}), \prec_{\text{sb}}^1) \text{ is a } k\text{-flip of } (a(T_1 \cdots T_p), \prec_{\text{sb}}^2).$$

In fact, \prec_{sb}^1 can be defined as $\prec_{\text{sb}}^{j,1}$ on the elements of $a(T_{\tau_1} \cdots T_{\tau_p})$ that appear inside T_{τ_j} ($1 \leq j \leq p$) and as $(s_1 \prec_{\text{sb}}^1 \cdots \prec_{\text{sb}}^1 s_p)$ on the roots s_1, \dots, s_p of T_1, \dots, T_p , respectively. Analogously, we can define \prec_{sb}^2 over $a(T_1 \cdots T_p)$, this time using sibling-orders $\prec_{\text{sb}}^{j,2}$, for $1 \leq j \leq p$. Now the result follows by using the fact that $(T_{\tau_j}, \prec_{\text{sb}}^{j,1})$ is a k -flip of $(T_j, \prec_{\text{sb}}^{j,2})$, for each $1 \leq j \leq p$, and the fact that replacing in a sibling-ordered unranked tree (T, \prec_{sb}) a subtree rooted at a children of the root of T with a sibling-ordered unranked tree of its same rank- k MSO type preserves the rank- k MSO type of (T, \prec_{sb}) . This fact can be proved using standard composition arguments for the EF MSO game over sibling-ordered unranked trees (for a proof, see, e.g., [17]).

Analogously, we can prove that there exist sibling-orders \prec_{sb}^3 and \prec_{sb}^4 such that:

$$(a(T_{\tau_{p+1}} \cdots T_{\tau_r}), \prec_{\text{sb}}^3) \text{ is a } k\text{-flip of } (a(T_{p+1} \cdots T_r), \prec_{\text{sb}}^4).$$

Therefore, from Proposition 3.2 we have:

- (1) $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(T_{\tau_1}, \dots, T_{\tau_p})) = \text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(T_1, \dots, T_p))$.
- (2) $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(T_{\tau_{p+1}}, \dots, T_{\tau_r})) = \text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(T_{p+1}, \dots, T_r))$.

Hence, in order to finish the proof of the claim it is sufficient to prove that:

$$\begin{aligned} \text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(T_1, \dots, T_p)) &= \text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(T_{p+1}, \dots, T_r)) \implies \\ &\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(T_{\tau_1}, \dots, T_{\tau_p})) = \text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(T_{\tau_{p+1}}, \dots, T_{\tau_r})). \end{aligned}$$

This is what we do next. In order to simplify notation, we sometimes write $a(\kappa_1, \dots, \kappa_m)$ and $\mathcal{A}(\kappa_1, \dots, \kappa_m)$, for $\kappa_1, \dots, \kappa_m$ (not necessarily distinct) elements in \mathcal{T} , instead of $a(T_{\kappa_1}, \dots, T_{\kappa_m})$ and $\mathcal{A}(T_{\kappa_1}, \dots, T_{\kappa_m})$, respectively.

With each unranked tree of the form T_τ , for $\tau \in \mathcal{T}$, we associate an arbitrary sibling-order \prec_{sb}^τ over $\text{dom}(T)$. Given a linear order \prec over $\{1, \dots, m\}$, we associate with the unranked tree $a(\kappa_1, \dots, \kappa_m)$, where each κ_i is an element in \mathcal{T} ($1 \leq i \leq m$), a sibling-order $\prec_{\text{sb}}^{\prec}$ over $\text{dom}(a(\kappa_1, \dots, \kappa_m))$ defined in the following way:

- (1) The interpretation of $\prec_{\text{sb}}^{\prec}$ over T_{κ_j} , for $1 \leq j \leq m$, corresponds to $\prec_{\text{sb}}^{\kappa_j}$.
- (2) Over the roots s_1, \dots, s_m of $T_{\kappa_1}, \dots, T_{\kappa_m}$, respectively, it is the case that $s_i \prec_{\text{sb}}^{\prec} s_j$ iff $i < j$, for each $1 \leq i, j \leq m$.

Since $\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(T_1, \dots, T_p)) = \text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(T_{p+1}, \dots, T_r))$, Proposition 3.2 implies that there exist linear orders \prec_1 and \prec_2 over $\text{dom}(\mathcal{A}(T_1, \dots, T_p))$ and $\text{dom}(\mathcal{A}(T_{p+1}, \dots, T_r))$, respectively, such that:

$$(\mathcal{A}(T_1, \dots, T_p), \prec_1) \text{ is a } k\text{-flip of } (\mathcal{A}(T_{p+1}, \dots, T_r), \prec_2).$$

In other words, $(\mathcal{A}(T_1, \dots, T_p), \prec_1)$ belongs to the same connected component than $(\mathcal{A}(T_{p+1}, \dots, T_r), \prec_2)$ in the undirected graph defined by relation $\sim_k^{\text{MSO, All}, \prec}$. We prove next that this implies that $(a(\tau_1, \dots, \tau_p), \prec_{\text{sb}}^{\prec_1})$ belongs to the same connected component than $(a(\tau_{p+1}, \dots, \tau_r), \prec_{\text{sb}}^{\prec_2})$ in the undirected graph defined by relation $\sim_k^{\text{MSO, Trees}, \prec_{\text{sb}}}$, i.e., that $(a(\tau_1, \dots, \tau_p), \prec_{\text{sb}}^{\prec_1})$ is a k -flip of $(a(\tau_{p+1}, \dots, \tau_r), \prec_{\text{sb}}^{\prec_2})$. From Proposition 3.2, this implies that $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(\tau_1, \dots, \tau_p)) = \text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(\tau_{p+1}, \dots, \tau_r))$. We abuse notation and write \sim_k instead of $\sim_k^{\text{MSO, All}, \prec}$ and $\sim_k^{\text{MSO, Trees}, \prec_{\text{sb}}}$, as the superscript will always be clear from the context.

Notice that it is sufficient to prove the following. If $\kappa_1, \dots, \kappa_m, \kappa_{m+1}, \dots, \kappa_n$ are (not necessarily distinct) elements in \mathcal{T} , and \prec_1 and \prec_2 are linear orders over $\text{dom}(\mathcal{A}(\kappa_1, \dots, \kappa_m)) = \{1, \dots, m\}$ and $\text{dom}(\mathcal{A}(\kappa_{m+1}, \dots, \kappa_n)) = \{1, \dots, n - m\}$, respectively, then

$$\begin{aligned} (\mathcal{A}(\kappa_1, \dots, \kappa_m), \prec_1) \sim_k (\mathcal{A}(\kappa_{m+1}, \dots, \kappa_n), \prec_2) \implies \\ (a(\kappa_1, \dots, \kappa_m), \prec_{\text{sb}}^{\prec_1}) \sim_k (a(\kappa_{m+1}, \dots, \kappa_n), \prec_{\text{sb}}^{\prec_2}). \end{aligned}$$

Assume then that $(\mathcal{A}(\kappa_1, \dots, \kappa_m), \prec_1) \sim_k (\mathcal{A}(\kappa_{m+1}, \dots, \kappa_n), \prec_2)$. From the definition of relation \sim_k , there are two possibilities:

- (1) $\mathcal{A}(\kappa_1, \dots, \kappa_m) = \mathcal{A}(\kappa_{m+1}, \dots, \kappa_n)$. and \prec_2 is a permutation of \prec_1 over $\{1, \dots, m\}$. Clearly then $a(\kappa_1, \dots, \kappa_m) = a(\kappa_{m+1}, \dots, \kappa_n)$, and hence:

$$(a(\kappa_1, \dots, \kappa_m), \prec_{\text{sb}}^{\prec_1}) \sim_k (a(\kappa_{m+1}, \dots, \kappa_n), \prec_{\text{sb}}^{\prec_2}).$$

- (2) $\mathcal{A}(\kappa_1, \dots, \kappa_m) \neq \mathcal{A}(\kappa_{m+1}, \dots, \kappa_n)$ but $(\mathcal{A}(\kappa_1, \dots, \kappa_m), \prec_1) \stackrel{\text{MSO}}{\equiv_k} (\mathcal{A}(\kappa_{m+1}, \dots, \kappa_n), \prec_2)$. A standard composition argument for the EF MSO game over unranked trees shows in this case that $(a(\kappa_1, \dots, \kappa_m), \prec_{\text{sb}}^{\prec_1})$

$$\begin{aligned} &\equiv_k^{\text{MSO}} (a(\kappa_{m+1}, \dots, \kappa_n), \prec_{\text{sb}}^{<_2}) \text{ (for a proof, see, e.g., [17]). Therefore:} \\ &\quad (a(T_{\kappa_1}, \dots, T_{\kappa_m}), \prec_{\text{sb}}^{<_1}) \sim_k (a(T_{\kappa_{m+1}}, \dots, T_{\kappa_n}), \prec_{\text{sb}}^{<_2}). \end{aligned}$$

This finishes the proof of the claim. \square

Assume that $\mathcal{T}' \subseteq \mathcal{T}$ is the set of \prec_{sb} -invariant rank- k MSO types of unranked trees T such that $T \in Q_\varphi$. Let us define a TA $\mathcal{N} := (\Sigma, \mathcal{T}, \mathcal{T}', \delta)$, such that the word $\tau_1 \dots \tau_m \in \mathcal{T}^*$ belongs to $\delta(\tau, a)$ (for $\tau \in \mathcal{T}$ and $a \in \Sigma$) iff $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(T_{\tau_1}, \dots, T_{\tau_p})) = \tau$. Next claim shows that \mathcal{N} is indeed a TA:

Claim 4.6. For each $a \in \Sigma$ and $\tau \in \mathcal{T}$, the set $\delta(\tau, a)$ is a regular language over \mathcal{T} .

Proof. From Claim 4.5, we know that the type $\text{tp}_{(\text{MSO}+\{\prec_{\text{sb}}\})_{\text{inv}}}^k(a(\tau_1, \dots, \tau_p))$ is determined by $\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\tau_1, \dots, \tau_p))$. Therefore, it is sufficient to construct a deterministic NFA \mathcal{N}' over alphabet \mathcal{T} that, given a word $\tau_1 \dots \tau_p$ in \mathcal{T}^* , the unique run of \mathcal{N}' over $\tau_1 \dots \tau_p$ labels position p with $\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\tau_1, \dots, \tau_p))$. The states of \mathcal{N}' are all the types of the form $\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\kappa_1, \dots, \kappa_q))$, for $\kappa_1, \dots, \kappa_q$ (not necessarily distinct) elements in \mathcal{T} ; the initial state of \mathcal{N}' is $\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}())$, the rank- k (MSO + $\{\prec\}$)_{inv} type of the empty structure over vocabulary $(P_\tau)_{\tau \in \mathcal{T}}$; and the transition relation δ' of \mathcal{N}' satisfies that $\delta(\chi, \tau)$, for $\chi = \text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\kappa_1, \dots, \kappa_q))$ and $\tau \in \mathcal{T}$, is $\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\kappa_1, \dots, \kappa_q, \tau))$. We prove next that the transition function δ' is well-defined.

We prove using invariant types that for each $\tau \in \mathcal{T}$, $\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\kappa_1, \dots, \kappa_q, \tau))$ is uniquely determined by $\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\kappa_1, \dots, \kappa_q))$. This shows that δ' is well-defined and also that \mathcal{N}' is deterministic. Consider then (not necessarily distinct) elements $\kappa_1, \dots, \kappa_q, \kappa_{q+1}, \dots, \kappa_t$ in \mathcal{T} , and assume that:

$$\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\kappa_1, \dots, \kappa_q)) = \text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\kappa_{q+1}, \dots, \kappa_t)).$$

It then follows from Proposition 3.2 that there exist linear orders $<_1$ and $<_2$ over $\text{dom}(\mathcal{A}(\kappa_1, \dots, \kappa_q)) = \{1, \dots, q\}$ and $\text{dom}(\mathcal{A}(\kappa_{q+1}, \dots, \kappa_t)) = \{1, \dots, t - q\}$, respectively, such that:

$$(\mathcal{A}(\kappa_1, \dots, \kappa_q), <_1) \text{ is a } k\text{-flip of } (\mathcal{A}(\kappa_{q+1}, \dots, \kappa_t), <_2).$$

But then clearly

$$(\mathcal{A}(\kappa_1, \dots, \kappa_q, \tau), <'_1) \text{ is a } k\text{-flip of } (\mathcal{A}(\kappa_{q+1}, \dots, \kappa_t, \tau), <'_2),$$

where $<'_1$ is the linear order over $\{1, \dots, q+1\}$ such that the restriction of $<'_1$ over $\{1, \dots, q\}$ is $<_1$ and $i <'_1 q+1$, for each $1 \leq i \leq q$, and equivalently for $<'_2$, this time over $\{1, \dots, t+1-q\}$. We conclude from Proposition 3.2 that

$$\text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\kappa_1, \dots, \kappa_q, \tau)) = \text{tp}_{(\text{MSO}+\{\prec\})_{\text{inv}}}^k(\mathcal{A}(\kappa_{q+1}, \dots, \kappa_t, \tau)).$$

This concludes the proof of the claim. \square

By definition, \mathcal{N} is deterministic. Furthermore, Claim 4.5 implies that \mathcal{N} is invariant. In fact, if $\tau_1 \cdots \tau_m$ is a word in \mathcal{T}^* and π is a permutation over $\{1, \dots, m\}$, then it is the case that $\text{tp}_{(\text{MSO}+\{<\})_{\text{inv}}}^k(\mathcal{A}(T_{\tau_1}, \dots, T_{\tau_m})) = \text{tp}_{(\text{MSO}+\{<\})_{\text{inv}}}^k(\mathcal{A}(T_{\tau_{\pi_1}}, \dots, T_{\tau_{\pi_m}}))$ (since $\mathcal{A}(T_{\tau_1}, \dots, T_{\tau_m}) \sim_k \mathcal{A}(T_{\tau_{\pi_1}}, \dots, T_{\tau_{\pi_m}})$, and, therefore, Claim 4.5 tells us that $\text{tp}_{(\text{MSO}+\{<_{\text{sb}}\})_{\text{inv}}}^k(a(T_{\tau_1}, \dots, T_{\tau_m})) = \text{tp}_{(\text{MSO}+\{<\})_{\text{inv}}}^k(a(T_{\tau_{\pi_1}}, \dots, T_{\tau_{\pi_m}}))$, for each $a \in \Sigma$).

Finally, it is possible to prove by induction that the unique run of \mathcal{N} on a $<_{\text{sb}}$ -ordered unranked tree $(T, <_{\text{sb}})$ labels each node s of T with $\text{tp}_{(\text{MSO}+\{<_{\text{sb}}\})_{\text{inv}}}^k(T_s)$, where T_s is the subtree of T rooted on s . In fact, consider an arbitrary a -labeled node s of T with children $s_1 <_{\text{sb}} \dots <_{\text{sb}} s_m$, and assume that the unique run of \mathcal{N} on $(T, <_{\text{sb}})$ assigns states τ_1, \dots, τ_m to s_1, \dots, s_m , respectively. Then by induction hypothesis $\text{tp}_{(\text{MSO}+\{<_{\text{sb}}\})_{\text{inv}}}^k(T_{s_i}) = \tau_i$, for each $1 \leq i \leq m$. By definition, the run of \mathcal{N} on $(T, <_{\text{sb}})$ labels s with $\text{tp}_{(\text{MSO}+\{<_{\text{sb}}\})_{\text{inv}}}^k(a(T_{\tau_1}, \dots, T_{\tau_m}))$. But notice that $\mathcal{A}(T_{\tau_1}, \dots, T_{\tau_m}) = \mathcal{A}(T_{s_1}, \dots, T_{s_m})$, which implies that $\text{tp}_{(\text{MSO}+\{<_{\text{sb}}\})_{\text{inv}}}^k(a(T_{\tau_1}, \dots, T_{\tau_m})) = \text{tp}_{(\text{MSO}+\{<\})_{\text{inv}}}^k(a(T_{s_1}, \dots, T_{s_m}))$ from Claim 4.5. The result now follows since $T_s = a(T_{s_1}, \dots, T_{s_m})$.

The latter means that for an arbitrary sibling-order $<_{\text{sb}}$ over $\text{dom}(T)$ we have that \mathcal{N} accepts $(T, <_{\text{sb}})$ if and only if $\text{tp}_{(\text{MSO}+\{<_{\text{sb}}\})_{\text{inv}}}^k(T) \in c'$ if and only if $T \in Q_\varphi$. This concludes our proof. \square

5. AN ORDER-INVARIANT FEFERMAN-VAUGHT THEOREM

The classical Feferman-Vaught theorem shows how theories of complex structures can be recovered using theories of simpler structures they are built from [8, 16]. In a simple version, it says that FO theories of product $\mathcal{A} \times \mathcal{B}$ and disjoint union $\mathcal{A} \sqcup \mathcal{B}$ are determined by theories of \mathcal{A} and \mathcal{B} . In the case of disjoint unions, we assume that the vocabulary is augmented with unary predicates for the universes of structures \mathcal{A} and \mathcal{B} . We now use order-invariant types to show that the same is true for order-invariant FO theories of structures. As a consequence, we obtain new classes of structures where $<$ -invariant FO collapses to FO.

Let $\text{Th}_{(\text{FO}+\{<\})_{\text{inv}}}(\mathcal{A})$ be the $<$ -invariant FO theory of \mathcal{A} , i.e., the set of all $<$ -invariant FO sentences true in \mathcal{A} , and $\text{Th}_{(\text{FO}+\{<\})_{\text{inv}}}^k(\mathcal{A})$ be its restriction to sentences of quantifier rank up to k (note that without such a restriction, the theory will have a sentence describing \mathcal{A} up to isomorphism). We then prove the following.

Theorem 5.1. Let \mathcal{A}, \mathcal{B} be structures over the same vocabulary. Then both $\text{Th}_{(\text{FO}+\{<\})_{\text{inv}}}^k(\mathcal{A} \sqcup \mathcal{B})$ and $\text{Th}_{(\text{FO}+\{<\})_{\text{inv}}}^k(\mathcal{A} \times \mathcal{B})$ are uniquely determined by $\text{Th}_{(\text{FO}+\{<\})_{\text{inv}}}^k(\mathcal{A})$ and $\text{Th}_{(\text{FO}+\{<\})_{\text{inv}}}^k(\mathcal{B})$.

Proof: Given structures \mathcal{A} and \mathcal{B} , and linear orders $<^{\mathcal{A}}, <^{\mathcal{B}}$ on \mathcal{A} and \mathcal{B} , respectively, we define a linear order $(<^{\mathcal{A}} \times <^{\mathcal{B}})$ on $\mathcal{A} \times \mathcal{B}$ such that (a, b) precedes (a', b') in this order whenever $b <_{\mathcal{B}} b'$, or $b = b'$ and $a <_{\mathcal{A}} a'$ (i.e., lexicographically, starting with the second component). It follows immediately by a straightforward EF game argument that if we have structures $\mathcal{A}_i, \mathcal{B}_i$, for $i = 1, 2$, over the same vocabulary, and linear orders $<^{\mathcal{A}_i}$ and $<^{\mathcal{B}_i}$ on them such that $(\mathcal{A}_1, <^{\mathcal{A}_1}) \equiv_k (\mathcal{A}_2, <^{\mathcal{A}_2})$ and $(\mathcal{B}_1, <^{\mathcal{B}_1}) \equiv_k (\mathcal{B}_2, <^{\mathcal{B}_2})$, then $(\mathcal{A}_1 \times \mathcal{B}_1, (<^{\mathcal{A}_1} \times <^{\mathcal{B}_1})) \equiv_k (\mathcal{A}_2 \times \mathcal{B}_2, (<^{\mathcal{A}_2} \times <^{\mathcal{B}_2}))$.

We use this observation, and the notion of a k -flip, to show the following.

Lemma 5.2. Let $\mathcal{A}_i, \mathcal{B}_i$ be structures over the same vocabulary, $i = 1, 2$. Then:

- (1) $\text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{A}_1 \sqcup \mathcal{B}_1) = \text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{A}_2 \sqcup \mathcal{B}_2)$ whenever $\text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{A}_1) = \text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{A}_2)$ and $\text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{B}_1) = \text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{B}_2)$.
- (2) $\text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{A}_1 \times \mathcal{B}_1) = \text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{A}_2 \times \mathcal{B}_2)$ whenever $\text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{A}_1) = \text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{A}_2)$ and $\text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{B}_1) = \text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{B}_2)$.

We only prove the case of the product in Lemma 5.2 since the case of the disjoint sum is completely analogous. From Proposition 3.2 we have to show that there exist linear orders $\langle_{\mathcal{A}_1 \times \mathcal{B}_1}, \langle_{\mathcal{A}_2 \times \mathcal{B}_2}$ on $(\mathcal{A}_1 \times \mathcal{B}_1)$ and $(\mathcal{A}_2 \times \mathcal{B}_2)$, respectively, such that $(\mathcal{A}_2 \times \mathcal{B}_2, \langle_{\mathcal{A}_2 \times \mathcal{B}_2})$ is a k -flip of $(\mathcal{A}_1 \times \mathcal{B}_1, \langle_{\mathcal{A}_1 \times \mathcal{B}_1})$.

Since $\text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{A}_1) = \text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{A}_2)$ and $\text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{B}_1) = \text{tp}_{(\mathcal{L}+\{\langle\cdot\rangle\})_{\text{inv}}}^k(\mathcal{B}_2)$, we know from Proposition 3.2 that there exist linear orders $\langle_{\mathcal{A}_i}, \langle_{\mathcal{B}_i}$, $i \in [1, 2]$, on \mathcal{A}_i and \mathcal{B}_i , respectively, such that $(\mathcal{A}_2, \langle_{\mathcal{A}_2})$ is a k -flip of $(\mathcal{A}_1, \langle_{\mathcal{A}_1})$, and $(\mathcal{B}_2, \langle_{\mathcal{B}_2})$ is a k -flip of $(\mathcal{B}_1, \langle_{\mathcal{B}_1})$. We now prove that $(\mathcal{A}_2 \times \mathcal{B}_2, (\langle_{\mathcal{A}_2} \alpha \langle_{\mathcal{B}_2}))$ is a k -flip of $(\mathcal{A}_1 \times \mathcal{B}_1, (\langle_{\mathcal{A}_1} \alpha \langle_{\mathcal{B}_1}))$, which, by Proposition 3.2, implies the result.

Assume that

$$(\mathcal{A}^1, \langle_{\mathcal{A}}^1) \sim_k (\mathcal{A}^2, \langle_{\mathcal{A}}^2) \sim_k \cdots \sim_k (\mathcal{A}^m, \langle_{\mathcal{A}}^m)$$

is a sequence witnessing the fact that $(\mathcal{A}_2, \langle_{\mathcal{A}_2})$ is a k -flip of $(\mathcal{A}_1, \langle_{\mathcal{A}_1})$. Then $(\mathcal{A}^1, \langle_{\mathcal{A}}^1) = (\mathcal{A}_1, \langle_{\mathcal{A}_1})$, and $(\mathcal{A}^m, \langle_{\mathcal{A}}^m) = (\mathcal{A}_2, \langle_{\mathcal{A}_2})$. Also, assume that

$$(\mathcal{B}^1, \langle_{\mathcal{B}}^1) \sim_k (\mathcal{B}^2, \langle_{\mathcal{B}}^2) \sim_k \cdots \sim_k (\mathcal{B}^n, \langle_{\mathcal{B}}^n)$$

is a sequence witnessing the fact that $(\mathcal{B}_2, \langle_{\mathcal{B}_2})$ is a k -flip of $(\mathcal{B}_1, \langle_{\mathcal{B}_1})$. Then $(\mathcal{B}^1, \langle_{\mathcal{B}}^1) = (\mathcal{B}_1, \langle_{\mathcal{B}_1})$, and $(\mathcal{B}^n, \langle_{\mathcal{B}}^n) = (\mathcal{B}_2, \langle_{\mathcal{B}_2})$. Then in order to show $(\mathcal{A}_2 \times \mathcal{B}_2, (\langle_{\mathcal{A}_2} \alpha \langle_{\mathcal{B}_2}))$ is a k -flip of $(\mathcal{A}_1 \times \mathcal{B}_1, (\langle_{\mathcal{A}_1} \alpha \langle_{\mathcal{B}_1}))$, it is enough to show that

$$\begin{aligned} (\mathcal{A}^1 \times \mathcal{B}^1, (\langle_{\mathcal{A}}^1 \alpha \langle_{\mathcal{B}}^1)) &\sim_k (\mathcal{A}^2 \times \mathcal{B}^1, (\langle_{\mathcal{A}}^2 \alpha \langle_{\mathcal{B}}^1)) \sim_k \cdots \sim_k (\mathcal{A}^m \times \mathcal{B}^1, (\langle_{\mathcal{A}}^m \alpha \langle_{\mathcal{B}}^1)) \sim_k \\ &(\mathcal{A}^m \times \mathcal{B}^2, (\langle_{\mathcal{A}}^m \alpha \langle_{\mathcal{B}}^2)) \sim_k (\mathcal{A}^m \times \mathcal{B}^3, (\langle_{\mathcal{A}}^m \alpha \langle_{\mathcal{B}}^3)) \sim_k \cdots \sim_k (\mathcal{A}^m \times \mathcal{B}^n, (\langle_{\mathcal{A}}^m \alpha \langle_{\mathcal{B}}^n)). \end{aligned}$$

Assume first that transition is from $(\mathcal{A}^i \times \mathcal{B}^1, (\langle_{\mathcal{A}}^i \alpha \langle_{\mathcal{B}}^1))$ to $(\mathcal{A}^{i+1} \times \mathcal{B}^1, (\langle_{\mathcal{A}}^{i+1} \alpha \langle_{\mathcal{B}}^1))$ for some $i < m$. We analyze two cases:

- $(\mathcal{A}^i, \langle_{\mathcal{A}}^i) \equiv_k (\mathcal{A}^{i+1}, \langle_{\mathcal{A}}^{i+1})$: We conclude that $(\mathcal{A}^i \times \mathcal{B}^1, (\langle_{\mathcal{A}}^i \alpha \langle_{\mathcal{B}}^1)) \equiv_k (\mathcal{A}^{i+1} \times \mathcal{B}^1, (\langle_{\mathcal{A}}^{i+1} \alpha \langle_{\mathcal{B}}^1))$ from the observation at the beginning of the proof and the fact that $(\mathcal{B}_1, \langle_{\mathcal{B}}^1) \equiv_k (\mathcal{B}_1, \langle_{\mathcal{B}}^1)$.
- $(\mathcal{A}^{i+1}, \langle_{\mathcal{A}}^{i+1})$ is a permutation of $(\mathcal{A}^i, \langle_{\mathcal{A}}^i)$ (i.e., one reinterprets the order on the same structure): Clearly, $(\mathcal{A}^{i+1} \times \mathcal{B}^1, (\langle_{\mathcal{A}}^{i+1} \alpha \langle_{\mathcal{B}}^1))$ is a permutation of $(\mathcal{A}^i \times \mathcal{B}^1, (\langle_{\mathcal{A}}^i \alpha \langle_{\mathcal{B}}^1))$.

The case when the transition is from $(\mathcal{A}^m \times \mathcal{B}^j, (\langle_{\mathcal{A}}^m \alpha \langle_{\mathcal{B}}^j))$ to $(\mathcal{A}^m \times \mathcal{B}^{j+1}, (\langle_{\mathcal{A}}^m \alpha \langle_{\mathcal{B}}^{j+1}))$ for some $j < n$ is completely analogous. This proves the lemma.

We now conclude the proof of the theorem. Again, we only prove it for the product. Let φ be a $(\text{FO} + \langle)_{\text{inv}}$ sentence, and assume that

$$(\text{tp}_{(\text{FO}+\langle)_{\text{inv}}}^k(\mathcal{A}_1), \text{tp}_{(\text{FO}+\langle)_{\text{inv}}}^k(\mathcal{B}_1)), \dots, (\text{tp}_{(\text{FO}+\langle)_{\text{inv}}}^k(\mathcal{A}_m), \text{tp}_{(\text{FO}+\langle)_{\text{inv}}}^k(\mathcal{B}_m))$$

is an enumeration of all different pairs of rank- k $(\text{FO} + \langle)_{\text{inv}}$ types of structures $\mathcal{A}_i, \mathcal{B}_i$ such that $(\mathcal{A}_i \times \mathcal{B}_i) \models \varphi$, $i \leq m$. Associate with each sentence $\text{tp}_{(\text{FO}+\langle)_{\text{inv}}}^k(\mathcal{A}_i)$ a propositional variable α_i , and with each sentence $\text{tp}_{(\text{FO}+\langle)_{\text{inv}}}^k(\mathcal{B}_i)$ a propositional variable β_i . Then it is

possible to show that the boolean function Φ defined over propositional variables α_i, β_i in the following way

$$\Phi(\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_m) = 1 \iff \bigvee_i (\alpha_i \wedge \beta_i) = 1,$$

where $\alpha_i = 1$ if and only if $\mathcal{A} \models \text{tp}_{(\text{FO}+\langle \cdot \rangle)_{\text{inv}}}^k(\mathcal{A}_i)$ and $\beta_i = 1$ if and only if $\mathcal{B} \models \text{tp}_{(\text{FO}+\langle \cdot \rangle)_{\text{inv}}}^k(\mathcal{B}_i)$, $i \leq m$, satisfies that

$$\Phi(\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_m) = 1 \iff (\mathcal{A} \times \mathcal{B}) \models \varphi.$$

In fact, assume first that $(\mathcal{A} \times \mathcal{B}) \models \varphi$. Then $\mathcal{A} = \mathcal{A}_i$ and $\mathcal{B} = \mathcal{B}_i$ for some $i \leq m$, and, therefore, $\alpha_i \wedge \beta_i = 1$, and $\Phi = 1$. Assume on the other hand that $\Phi = 1$. Then for some $i \leq m$, $\alpha_i \wedge \beta_i = 1$, implying that $\mathcal{A} \models \text{tp}_{(\text{FO}+\langle \cdot \rangle)_{\text{inv}}}^k(\mathcal{A}_i)$ and $\mathcal{B} \models \text{tp}_{(\text{FO}+\langle \cdot \rangle)_{\text{inv}}}^k(\mathcal{B}_i)$. Hence, $\mathcal{A} \equiv_k^{\text{inv}} \mathcal{A}_i$ and $\mathcal{B} \equiv_k^{\text{inv}} \mathcal{B}_i$, and from Lemma 5.2, $(\mathcal{A} \times \mathcal{B}) \equiv_k^{\text{inv}} (\mathcal{A}_i \times \mathcal{B}_i)$. But $(\mathcal{A}_i \times \mathcal{B}_i) \models \varphi$, and thus $(\mathcal{A} \times \mathcal{B}) \models \varphi$. This completes the proof. \square

We now use Theorem 5.1 to describe classes of structures on which $\langle \cdot \rangle$ -invariant FO collapses to FO. Let $\mathcal{C}, \mathcal{C}'$ be classes of structures. By $\prod(\mathcal{C}, \mathcal{C}')$ (respectively, $\coprod(\mathcal{C}, \mathcal{C}')$) we denote classes of structures of the form $\mathcal{A} \times \mathcal{B}$ (respectively, $\mathcal{A} \sqcup \mathcal{B}$) where $\mathcal{A} \in \mathcal{C}$ and $\mathcal{B} \in \mathcal{C}'$.

Corollary 5.3. Let $\mathcal{C}, \mathcal{C}'$ be classes of structures on which $\langle \cdot \rangle$ -invariant FO collapses to FO. Then $\langle \cdot \rangle$ -invariant FO collapses to FO over both $\prod(\mathcal{C}, \mathcal{C}')$ and $\coprod(\mathcal{C}, \mathcal{C}')$.

Indeed, every $\langle \cdot \rangle$ -invariant sentence over $\mathcal{A} \times \mathcal{B}$ (or $\mathcal{A} \sqcup \mathcal{B}$) is given by a finite set of order-invariant types, which in turn, by Theorem 5.1, are given by sets of pairs of $\langle \cdot \rangle$ -invariant types over \mathcal{A} and \mathcal{B} . Since these are expressible in FO by the assumption, we get that every $\langle \cdot \rangle$ -invariant sentence over products or disjoint unions is expressible in FO too.

For example, combining this with the results of [3, 18] we get that $\langle \cdot \rangle$ -invariant FO collapses to FO over grids (products of successor relations) or even products of words, i.e., grids colored in a way that is uniquely determined by coloring of its components.

One may wonder whether we can get the collapse result for arbitrarily colored grids. While we do not know the answer, we provide an example that indicates not only that it is hard to obtain such a result from the Feferman-Vaught theorem, but also that the collapse of $\langle \cdot \rangle$ -invariant FO is a very fragile notion.

Let \mathcal{C} be a class of structures. By $\text{Un}(\mathcal{C})$ we denote the class of structures of the form (\mathcal{A}, C) , where $\mathcal{A} \in \mathcal{C}$ and C is a subset of \mathcal{A} (i.e., structures of \mathcal{C} extended with a single unary predicate). Already such a tiny extension can destroy the collapse.

Proposition 5.4. There is a class of structures \mathcal{C} such that $\langle \cdot \rangle$ -invariant FO collapses to FO over \mathcal{C} , but it does not collapse to FO over $\text{Un}(\mathcal{C})$.

Proof: The class \mathcal{C} of structures we consider will have two unary relations V and E partitioning the domain, and two binary relations L and R . The interpretation is that elements of E provide names for edges in the complete directed graph $V \times V$, and if e is the name of an edge (x, y) , then $L(e, x)$ and $R(e, y)$ hold (i.e., these stand for left and right vertices of a directed edge). In any FO sentence over such structures (even with extra predicates), we can first make quantification relativized to V and E (i.e., $\exists x \in V$, $\exists e \in E$) and then replace each $\exists e \in E$ with $\exists e_l, e_r \in V$, and then change each $L(e, x)$ to $x = e_l$ and each $R(e, y)$ to $y = e_r$. If we have an order-invariant sentence, we can assume that

the order is given on V and extended to E lexicographically, while each element of V is below each element of E (since we have a complete freedom in choosing the order). That is, each $v < e$ is replaced by *true* if $v \in V$ and $e \in E$, and each $e < e'$ is replaced by $e_l < e'_l \vee (e_l = e'_l \wedge e_r < e'_r)$. Thus, every order-invariant sentence over \mathcal{C} is equivalent to an order-invariant sentence in the language of only one unary predicate V and the order $<$, and a simple counting argument shows that such sentences collapse to FO over V (see, e.g., [14]). Hence, we have the collapse over \mathcal{C} .

To show the lack of collapse over $\text{Un}(\mathcal{C})$, we construct an order-invariant sentence that

- (1) checks that the new relation C is a subset of E , and that V is of the form 2^X for some set X , and C contains precisely the edges of the subset relation;
- (2) and on structures of such form, uses the order to check whether $|X|$ is even.

The second item is done exactly as in the proof of the fact that on Boolean algebras, with an order one can check whether the number of atoms is even (see Example 1), and the proof that such a sentence is not expressible in FO alone is done in exactly the same way as the original proof separating $<$ -invariant FO from FO on Boolean algebras, see [14].

To ensure that C is of the right form, we must check the following, in addition to $C \subseteq E$.

- The relation C is reflexive, transitive, and anti-symmetric.
- There is a single element $v_0 \in V$ such that v_0 is connected by an edge in C to every other element of V (it plays the role of the empty set).
- We then define a set A of elements $v \neq v_0$ such that there is no C -edge between v_0 and some v' , and also a C -edge between v' and v . These play the role of atoms of the Boolean algebra. We also define $A(v)$ as the set of elements of A so that there is a C -edge from them to v (i.e., atoms under v).
- For all $v \neq v'$ which are different from v_0 , we must have $A(v) \neq A(v')$.
- For all $v \neq v'$ which are different from v_0 , we must have an element v'' so that $A(v'') = A(v) \cup A(v')$.

It is routine to verify that all of these are expressible in FO over V, E, L, R , and ensure that $C \subseteq E$ gives V and E the structure of a Boolean algebra. This concludes the proof. \square

Final remarks. The order-invariant Feferman-Vaught theorem presented in this section also holds for MSO in the case of disjoint unions, but fails for products (see, e.g., [16]). There are several other relevant operations (e.g., transductions or interpretations) that preserve the types of structures [16]. It is an interesting open problem to establish if such operations continue to preserve $<$ -invariant types as well. It also appears possible to use the technique of invariant types to look at containment of $<$ -invariant FO in MSO (as was done in [3]), for instance, for unbounded disjoint unions without the extra predicates for structures. Another possibility is to give a direct proof of a generalization of Courcelle's theorem on linear-time data complexity of MSO over structures of bounded treewidth to $<$ -invariant MSO [5], using the observation that the MSO version can be proved using Feferman-Vaught techniques [16].

REFERENCES

- [1] M. Anderson, D. van Melkebeek, N. Schweikardt, and L. Segoufin. Locality of queries definable in invariant first-order logic with arbitrary built-in predicates. In *International Colloquium on Automata, Languages and Programming (ICALP)*, pages 368–379, 2011.
- [2] P. Barceló and L. Libkin. Temporal logics over unranked trees. In *Symposium on Logic in Computer Science (LICS)*, pages 31–40, 2005.
- [3] M. Benedikt and L. Segoufin. Towards a characterization of order-invariant queries over tame graphs. *Journal of Symbolic Logic*, 74(1):168–186, 2009.
- [4] C. Chang and H. Keisler. *Model Theory*. North Holland, 1990.
- [5] Y. Chen and J. Flum. On the ordered conjecture. In *Proceedings of the 27th Annual IEEE Symposium on Logic in Computer Science*, pages 225–234, 2012.
- [6] H. Comon, M. Dauchet, R. Gilleron, C. Löding, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. *Tree Automata Techniques and Applications*. Available on: <http://www.grappa.univ-lille3.fr/tata>, 2007. Release October, 12th 2007.
- [7] B. Courcelle. The monadic second-order logic of graphs V: on closing the gap between definability and recognizability. *Theoretical Computer Science*, 80(2):153–202, 1991.
- [8] S. Feferman and R. Vaught. The first order properties of algebraic systems. *Fund. Math.*, 47:57–103, 1959.
- [9] S. Ginsburg and E. H. Spanier. Bounded regular sets. *Proceedings of the AMS*, 17(5):1043–1049, 1966.
- [10] E. Grädel, P. Kolaitis, L. Libkin, M. Marx, J. Spencer, M. Vardi, Y. Venema, and S. Weinstein. *Finite Model Theory and Its Applications*. Springer, 2007.
- [11] M. Grohe and T. Schwentick. Locality of order-invariant first-order formulas. *ACM Transactions on Computational Logic*, 1(1):112–130, 2000.
- [12] D. Kozen. On two letters versus three. In *Fixed Points in Computer Science (FICS)*, pages 44–50, 2002.
- [13] R. E. Ladner. Application of model theoretic games to discrete linear orders and finite automata. *Information and Control*, 33(4):281–303, 1977.
- [14] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [15] L. Libkin and L. Wong. Lower bounds for invariant queries in logics with counting. *Theoretical Computer Science*, 288(1):153–180, 2002.
- [16] J. A. Makowsky. Algorithmic uses of the Feferman-Vaught theorem. *Annals of Pure and Applied Logic*, 126(1-3):159–213, 2004.
- [17] F. Neven and T. Schwentick. Query automata over finite trees. *Theoretical Computer Science*, 275(1-2):633–674, 2002.
- [18] H. Niemistö. On locality and uniform reduction. In *Symposium on Logic in Computer Science (LICS)*, pages 41–50, 2005.
- [19] M. Otto. Epsilon-logic is more expressive than first-order logic over finite structures. *Journal of Symbolic Logic*, 65(4):1749–1757, 2000.
- [20] B. Rossman. Successor-invariant first-order logic on finite structures. *Journal of Symbolic Logic*, 72(2):601–618, 2007.
- [21] N. Schweikardt. On the expressive power of logics with invariant uses of arithmetic predicates. In *Workshop on Logic, Language, Information and Computation (WoLLIC)*, pages 85–87, 2012.
- [22] N. Schweikardt. A short tutorial on order-invariant first-order logic. In *International Computer Science Symposium in Russia (CSR)*, pages 112–126, 2013.
- [23] N. Schweikardt and L. Segoufin. Addition-invariant FO and regularity. In *Symposium on Logic in Computer Science (LICS)*, pages 273–282, 2010.

APPENDIX A. A DIRECT PROOF OF LEMMA 4.4.

We prove the only if direction, as the other one is immediate. A word w over alphabet $\Sigma = \{a_1, a_2, \dots, a_r\}$ is said to be *partitioned*, if it belongs to the regular language $(a_1)^*(a_2)^*\dots(a_r)^*$. Notice that for every word w there is a unique permutation w^p of w (up to isomorphism) that is partitioned.

We start by proving that rank- k MSO types of partitioned words can be defined by means of finite collections of r -tuples of sets of the form $S_{k,p}$. Formally, we prove:

(*) Let τ be the rank- k MSO type of a partitioned word over Σ . There exists a finite family \mathcal{S}_τ of r -tuples of sets of the form $S_{k,p}$, such that for every word w over Σ we have that $\text{tp}_{\text{MSO}}^k(w_p) = \tau$ iff for some $(S_1, \dots, S_r) \in \mathcal{S}_\tau$ it is the case that $\Pi(w) \in S_1 \times \dots \times S_r$.

We prove (*) next. For $a \in \Sigma$ and w a partitioned word over Σ , we denote by w_a the maximal subword of w that is of the form a^* . By using a standard composition argument for the MSO EF game, one can show that for any two partitioned words w and w' over Σ ,

$$\text{tp}_{\text{MSO}}^k(w) = \text{tp}_{\text{MSO}}^k(w') \iff \text{tp}_{\text{MSO}}^k(w_{a_i}) = \text{tp}_{\text{MSO}}^k(w'_{a_i}), \text{ for each } 1 \leq i \leq r.$$

Hence, for each rank- k MSO type τ of a partitioned word there is a finite family \mathcal{F}_τ of tuples of the form $(\tau_{a_1}, \dots, \tau_{a_r})$, where each τ_{a_i} is the rank- k MSO type of some word that only uses symbol a_i ($1 \leq i \leq r$), such that for each partitioned word w the rank- k MSO type of w is τ iff for some $(\tau_{a_1}, \dots, \tau_{a_r}) \in \mathcal{F}_\tau$ the rank- k MSO type of w_{a_i} is τ_{a_i} , for each $1 \leq i \leq r$. Therefore, in order to prove (*) it is enough to show that for each rank- k MSO type of the form τ_a ($a \in \Sigma$) there exists a set of the form $S_{k,p}$ ($k, p \geq 0$) such that for a word w of the form a^* ,

$$\text{the rank-}k \text{ MSO type of } w \text{ is } \tau_a \iff |w| \in S_{k,p}.$$

This is what we do next.

It is known (see [13] and [14] for the textbook treatment) that for a word w of the form a^* , the rank- k MSO type of w is τ_a if and only if it is accepted by the NFA $\mathcal{N} = (\{a\}, \Gamma, \tau_0, \{\tau_a\}, \delta)$, where the set of states Γ is the set of all rank- k MSO types of words in a^* ; the initial state of \mathcal{N} is the rank- k MSO type of the empty word, denoted by τ_0 ; the final state of \mathcal{N} is τ_a ; and for $\tau' \in \Gamma$ we have that $\delta(\tau', a)$ contains all the rank- k MSO types of words of the form $w' \cdot a$, for w' a word in a^* with rank- k MSO type τ' . Clearly, \mathcal{N} is a deterministic NFA since a simple composition argument for the MSO EF game shows that the rank- k MSO type of $w' \cdot a$ is completely determined by the rank- k MSO type of w' . Furthermore, since the alphabet of \mathcal{N} is unary, a simple inspection of the transition graph of \mathcal{N} reveals that there must exist integers $k, p \geq 0$ such that for a word w of the form a^* it is the case that \mathcal{N} accepts w iff $|w| = k + np$ for some $n \geq 0$. This finishes the proof of (*). We now continue with the proof of the lemma.

Let L be a regular language that is closed under permutation. From Büchi's theorem, there is an MSO sentence φ_L over vocabulary $(\langle, (P_a)_{a \in \Sigma})$ that *defines* L , i.e., for each word w over Σ it is the case that $w \in L$ iff $w \models \varphi_L$. Assume that the quantifier rank of φ is k , and let \mathcal{T} be the set of all rank- k MSO types of the partitioned words w that satisfy φ (i.e., the ones that belong to L). If \mathcal{T} is empty (which implies that L is also empty), we define \mathcal{S} to be the empty set, which clearly satisfies the statement of the lemma. Let us assume then that \mathcal{T} is not empty. We claim that a word w belongs to L if and only there is a rank- k MSO type $\tau \in \mathcal{T}$ and an r -tuple $(S_1, \dots, S_r) \in \mathcal{S}_\tau$ such that $\Pi(w) \in S_1 \times \dots \times S_r$. Indeed:

$$\begin{aligned} w \in L &\iff w^p \in L && (L \text{ is closed under permutation}) \\ &\iff \text{the rank-}k \text{ MSO type } \tau \text{ of } w^p \text{ is in } \mathcal{T} && (\text{tp}_{\text{MSO}}^k(w^p) \in \mathcal{T} \Rightarrow w_p \models \varphi) \\ &\iff \Pi(w) \in S_1 \times \dots \times S_r, \text{ for some } (S_1, \dots, S_r) \in \mathcal{S}_\tau. && (\text{from } (*)) \end{aligned}$$

This finishes the proof of the lemma since \mathcal{T} is finite. \square