

ON THE STRICTNESS OF THE QUANTIFIER STRUCTURE HIERARCHY IN FIRST-ORDER LOGIC*

YUGUO HE

School of Computer Science, Beijing Institute of Technology, Beijing 100081, China.
e-mail address: hugo274@gmail.com

ABSTRACT. We study a natural hierarchy in first-order logic, namely the quantifier structure hierarchy, which gives a systematic classification of first-order formulas based on structural quantifier resource. We define a variant of Ehrenfeucht-Fraïssé games that characterizes quantifier classes and use it to prove that this hierarchy is strict over finite structures, using strategy compositions. Moreover, we prove that this hierarchy is strict even over ordered finite structures, which is interesting in the context of descriptive complexity.

1. INTRODUCTION

One of the major interests of finite model theory is to separate the expressive power of different logics or fragments of logics. Quantifiers are an important logical resource for measuring the logical complexity of problems. The study of fragments of first-order logic (FO) based on quantifier structures, especially quantifier prefixes, has a long history in model theory [2]. However, so far there are few results about the expressive power of such fragments. Walkoe [10] proved that there exists a sentence with prefix p which is different from any sentence with prefix q if p and q are different but with the same length. In the proof, the structures are assumed to be infinite. Afterwards Keisler and Walkoe [3] improved this result by showing its validity over finite structures. Chandra and Harel [1] proved that $\Sigma_k \subsetneq \Sigma_{k+1}$ over finite digraphs. Sipser [9] proved a similar result in the context of unbounded fan-in bounded depth circuits.

In 1996, Grädel and McColm [2] established a strict hierarchy based on quantifier classes in the infinitary logic over finite structures and resolved a conjecture of Immerman, i.e. $\Sigma_i^{TC} \subsetneq \Sigma_{i+1}^{TC}$ for each i . At the same time, they proposed a conjecture on the expressive power of the fragments of FO based on prefixes, which generalized the previous results [10], [3] and [1]. In 1998, Rosen [6] confirmed this conjecture and called the strict hierarchy based on these fragments of FO the first-order prefix hierarchy. Actually, Rosen proved a stronger result, which states that, over a single binary relation, for any prefix p there is a first-order

2012 ACM CCS: [Theory of computation]: Logic—Finite Model Theory; [Mathematics of computing]: Discrete mathematics—Graph theory—Trees.

Key words and phrases: Finite Model Theory, Quantifier Structure, Quantifier Class, Ehrenfeucht-Fraïssé Games, Strategy Composition, Point-expansion, Ordered Structures.

* This paper is from a part of the author's PhD thesis. It is an extension and revision of a conference paper published in Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science.

sentence φ_p in prenex normal form with prefix p , such that for any sentences ψ in infinitary logic, φ_p is not equivalent to ψ if p is not embeddable in the “quantifier structure” of ψ .¹ However, a stronger version of the conjecture remains open, i.e. whether it holds over finite structures or not [5, 6]. One way to prove the conjecture is to prove a finite version of Rosen’s main theorem.

In this paper, we continue this line of study. We define a variant of Ehrenfeucht-Fraïssé games that characterizes quantifier classes and prove the following main result:

Let S_1 and S_2 be two finite Γ -labeled forests. Over the class of all finite digraphs,

$$\text{if } S_1 \not\leq_e S_2, \text{ then } \text{FO}\{S_1\} \not\subseteq \text{FO}\{S_2\}.$$

The structures we use in the proof are finite trees, which makes it easy to prove a stronger result: The above main result holds even when the structures have a linear order. Here we introduce the ideas that are used to deal with linear order in a simpler context:

Over the class of all ordered finite digraphs,

$$\text{if } \mathcal{W}(S_1) \not\subseteq \mathcal{W}(S_2), \text{ then } \text{FO}\{S_1\} \not\subseteq \text{FO}\{S_2\}.$$

2. PRELIMINARIES

2.1. General background. Let \mathbb{N} and \mathbb{N}^+ denote the set of natural numbers (non-negative integers) and positive natural numbers respectively.

We assume that the readers have basic knowledge about finite model theory. In the following we briefly introduce some necessary background. The readers can cf. the textbook [4] for more of it.

A *relational signature* consists of a sequence of relation and constant symbols. In this paper, a signature is relational and finite, whenever mentioned.

Let $\sigma = \langle R_1, \dots, R_m, c_1, \dots, c_n \rangle$ be a signature, a σ -*structure* \mathfrak{A} consists of a universe $|\mathfrak{A}|$ together with an interpretation of

- each k -ary relation symbol $R_i \in \sigma$ as a k -ary relation on $|\mathfrak{A}|$, denoted by $R_i^{\mathfrak{A}}$;
- each constant symbol $c_i \in \sigma$ as an element in $|\mathfrak{A}|$.

A structure is called *finite* if its universe is a finite set.

A σ -structure \mathfrak{A}' is a *substructure* of \mathfrak{A} if the following hold:

- (1) $|\mathfrak{A}'| \subseteq |\mathfrak{A}|$;
- (2) For any k -ary relation $R \in \sigma \cup \{=\}$, $R^{\mathfrak{A}'} = R^{\mathfrak{A}} \cap |\mathfrak{A}'|^k$;
- (3) For any constant $c \in \sigma$, $c^{\mathfrak{A}'} = c^{\mathfrak{A}}$.

Let $\sigma' \subseteq \sigma$. The σ' -*reduct* of \mathfrak{A} , denoted $\mathfrak{A}|\sigma'$, is obtained from \mathfrak{A} by leaving all the symbols in $\sigma \setminus \sigma'$ uninterpreted.

Let \mathfrak{A} and \mathfrak{B} be two structures of the same signature. An *isomorphism* between \mathfrak{A} and \mathfrak{B} is a bijection $h : |\mathfrak{A}| \rightarrow |\mathfrak{B}|$ such that the following hold:

- (1) For any k -ary relation $R \in \sigma \cup \{=\}$ and $(a_1, \dots, a_k) \in |\mathfrak{A}|^k$,

$$(a_1, \dots, a_k) \in R^{\mathfrak{A}} \text{ iff } (h(a_1), \dots, h(a_k)) \in R^{\mathfrak{B}};$$

¹Here, the notion “quantifier structure” is from Grädel and McColm [2], which is different from ours (cf. Definition 2.6).

(2) For any constant $c \in \sigma$, $h(c^{\mathfrak{A}}) = c^{\mathfrak{B}}$.

Say that two structures \mathfrak{A} and \mathfrak{B} over the same signature are isomorphic if there is an isomorphism between them, denoted $\mathfrak{A} \cong \mathfrak{B}$.

Let $\bar{a} = (a_1, \dots, a_k) \in |\mathfrak{A}|^k$, $\bar{b} = (b_1, \dots, b_k) \in |\mathfrak{B}|^k$. Say that (\bar{a}, \bar{b}) defines a *partial isomorphism* between \mathfrak{A} and \mathfrak{B} if \bar{a} contains the elements that interpret all the constants of \mathfrak{A} , \bar{b} contains the elements that interpret all the constants of \mathfrak{B} , and the substructure of \mathfrak{A} that is generated by \bar{a} is isomorphic to the substructure of \mathfrak{B} that is generated by \bar{b} . More precisely, the following hold:

(1) for any m -ary relation symbol $R \in \sigma \cup \{=\}$ and any sequence (i_1, \dots, i_m) of numbers from $[k]$,

$$(a_{i_1}, \dots, a_{i_m}) \in R^{\mathfrak{A}} \text{ iff } (b_{i_1}, \dots, b_{i_m}) \in R^{\mathfrak{B}}.$$

(2) for any constant $c \in \sigma$ and any $i \in [k]$,

$$a_i = c^{\mathfrak{A}} \text{ iff } b_i = c^{\mathfrak{B}}.$$

We assume that the readers have basic knowledge about first-order logic, especially what is the meaning of “a formula is true in a structure”. Without loss of generality, we assume that all the formulas and sentences are in *negation normal form*, i.e. all negations can only occur immediately before atoms.

Let \mathfrak{A} be a σ -structure and ψ be a first-order sentence. We use $\mathfrak{A} \models \psi$ to denote that ψ is true in \mathfrak{A} , and we call \mathfrak{A} a *model* for ψ . Let $\text{Mod}(\psi)$ be the set of models of ψ . A *property* Q over σ is a set of σ -structures closed under isomorphism. Say that Q is *expressible*, or *definable*, in FO if there is a sentence φ in FO such that for every \mathfrak{A} , $\mathfrak{A} \in \text{Mod}(\varphi)$ iff $\mathfrak{A} \in Q$.

A *linear order* is a binary relation that is transitive, antisymmetric and total. Let τ be a signature. And let $\tau^{\text{ORD}} := \tau \cup \{\leq\}$ where \leq is interpreted in a τ^{ORD} -structure as a linear order of its universe.

2.2. Γ -labeled forests. Let $n \in \mathbb{N}^+$. Given a graph $\mathcal{G} = (V, E)$, a *directed path* P in \mathcal{G} is a sequence of vertices (v_0, \dots, v_n) such that there is an arc from v_i to v_{i+1} for any $i < n$. The length of P is n . A directed path is *nontrivial* if the length of the path is nonzero.

Trees are defined in the usual way in computer science. If there is an arrow from a node a to a node b , then we call a a *father* of b and b a *child* of a . In a tree, each node has zero or more children and each node has at most one father. A node which has no father is called a *root* and a node which has no child is called a *leaf*. An *inner node* is any node that has child nodes. A *tree* is a connected acyclic digraph that has a root and some leaves. A *degenerate tree* is a directed path.

The *height of a tree* is the length of a longest directed path in the tree. A *forest* is composed of disjoint trees. Let S be a forest. Define its height, denoted $h(S)$, as the maximum height of its trees. And define its rank, denoted $rk(S)$, as $h(S) + 1$ when S is not empty and 0 otherwise. Let $\Gamma = \{\exists, \forall\}$. A forest is a Γ -*labeled forest* if all its nodes are labeled with “ \exists ” or “ \forall ”. We call those nodes labeled with “ \exists ” \mathcal{E} nodes and the other nodes \mathcal{A} nodes.

A Γ -*labeled perfect binary tree* is a Γ -labeled tree where each node, except the leaves, has exactly one \mathcal{E} child and one \mathcal{A} child, and all the leaves are at the same depth.

A \exists_n -*perfect binary tree*, denoted $*\mathcal{T}_n^{\exists}$, is a Γ -labeled perfect binary tree, whose root is labeled with \exists and height is $(n - 1)$. Likewise, a \forall_n -*perfect binary tree*, denoted $*\mathcal{T}_n^{\forall}$, is a Γ -labeled perfect binary tree, whose root is labeled with \forall and height is $(n - 1)$.

2.3. Prefixes. The following terminology and conventions come from Grädel-McColm [2] and Rosen [6]. A *prefix* p is a finite string in Γ^* . The dual of p , denoted by \bar{p} , is the prefix obtained from p by swapping \exists with \forall . Let $\mathcal{P} \subseteq \Gamma^*$. Then $\bar{\mathcal{P}} := \{\bar{p} \in \Gamma^* \mid p \in \mathcal{P}\}$. A prefix p is a *subsequence* of a prefix q if p can be obtained from q by possibly deleting some elements of q , without changing the order of the remaining elements of q . A partial order on Γ^* , called prefix embedding, can be defined as follows: $p \preceq q$ iff p is a subsequence of q . Here we use the curly symbol to distinguish it from the usual symbol of linear orders. Nevertheless, whether a symbol stands for a linear order or (prefix) embedding should be easily decided from the context. We use the same notation “ \preceq ” to denote the embedding relation between two sets of prefixes. For $\mathcal{P}_1, \mathcal{P}_2 \subseteq \Gamma^*$, $\mathcal{P}_1 \preceq \mathcal{P}_2 \Leftrightarrow \forall p \in \mathcal{P}_1, \exists q \in \mathcal{P}_2$ s.t. $p \preceq q$. $\mathcal{P}_1 \prec \mathcal{P}_2$ if $\mathcal{P}_1 \preceq \mathcal{P}_2$ but $\mathcal{P}_2 \not\preceq \mathcal{P}_1$. We use “ $*$ ” to denote the concatenation of words. For any $\alpha \in \Gamma$ and $\mathcal{P} \subseteq \Gamma^*$, $\alpha * \mathcal{P} := \{\alpha * p \mid p \in \mathcal{P}\}$. We define $\mathcal{P}^- := \{p \mid \exists q \in \mathcal{P} \text{ s.t. } p \preceq q\}$ as the downward closure of \mathcal{P} . Let $\Gamma_c = \{\exists, \forall, \exists^*, \forall^*\}$ where \exists^* and \forall^* are characters. We interpret a word in Γ_c^* as a regular expression. $\gamma : \Gamma_c^* \rightarrow \wp(\Gamma^*)$ maps such a regular expression to the regular language it denotes, where $\wp(\Gamma^*)$ is the power set of Γ^* . We define $\gamma^- : \Gamma_c^* \rightarrow \wp(\Gamma^*)$ so that for any $v \in \Gamma_c^*$, $\gamma^-(v) = \{q \in \Gamma^* \mid \text{there is } q' \in \gamma(v) \text{ and } q \preceq q'\}$, the downward closure of $\gamma(v)$.

For a prefix p , $|p|$ is the length of p . $p[i]$ is the i -th letter of p . Let $l(p)$ be the last letter of p . For $0 \leq i < |p|$, let p^{-i} be the prefix obtained from p by removing the first i letters in p , i.e. $p = p[1] * \dots * p[i] * p^{-i}$.

Finally, let ϵ be the empty string.

Lemma 2.1. *Let p, q be prefixes. The following hold:*

- (1) $\overline{\bar{p}} = p$.
- (2) $\overline{\bar{p} * \bar{q}} = \overline{\bar{p} * \bar{q}}$.

Proof.

- (1) By definition.
- (2) Let $|p| = n, |q| = m$.

By definition, $\overline{\bar{p} * \bar{q}} = \overline{p[1] * \dots * p[n] * q[1] * \dots * q[m]} = \overline{p[1] * \dots * p[n]} * \overline{q[1] * \dots * q[m]}$.

Because the concatenation operation on words satisfies the associative law, it follows that $\overline{p[1] * \dots * p[n]} * \overline{q[1] * \dots * q[m]} = \overline{p[1] * \dots * p[n] * q[1] * \dots * q[m]}$. Hence, by definition, $\overline{\bar{p} * \bar{q}} = \overline{\bar{p} * \bar{q}}$. \square

Definition 2.2. (Rosen, [6]). Let $s \in \Gamma$ and s^i denote the string consisting of i repetitions of s . Define $f : \Gamma^* \rightarrow \Gamma_c^*$ as follows:

- (1) If $p = \exists^n$, then $f(p) := a_1 * \dots * a_{2n-1}$, where $a_i = \forall^*$ for i odd, and $a_i = \exists$ for i even;
- (2) If $p = \forall^n$, then $f(p) := a_1 * \dots * a_{2n-1}$, where $a_i = \exists^*$ for i odd, and $a_i = \forall$ for i even;
- (3) If $p = s_1^{i_1} * \dots * s_n^{i_n}$ ($s_i \in \{\exists, \forall\}, s_i \neq s_{i+1}, i_j \in \mathbb{N}^+$), then $f(p) := f(s_1^{i_1}) * \dots * f(s_n^{i_n})$.

Lemma 2.3. *Let p be a prefix. Then $f(p) = \overline{f(\bar{p})}$.*

Proof. Assume that $p = s_1^{i_1} * \dots * s_n^{i_n}$ ($s_i \in \{\exists, \forall\}, s_i \neq s_{i+1}, i_j \in \mathbb{N}^+$). By the definition of the dual of a prefix, $\bar{p} = \bar{s}_1^{i_1} * \dots * \bar{s}_n^{i_n}$. Note that $\bar{s}_i \neq \bar{s}_{i+1}$ since $s_i \neq s_{i+1}$. By definition, $f(\bar{p}) = f(\bar{s}_1^{i_1}) * \dots * f(\bar{s}_n^{i_n})$. Note that by definition $f(\bar{s}_j^{i_j}) = \overline{f(s_j^{i_j})}$. Hence, $f(\bar{p}) = \overline{f(s_1^{i_1}) * \dots * f(s_n^{i_n})}$. By Lemma 2.1 (ii), it means that $f(\bar{p}) = \overline{f(s_1^{i_1}) * \dots * f(s_n^{i_n})} = \overline{f(p)}$. Therefore, $f(p) = f(\bar{p})$. \square

Lemma 2.4. (Rosen, [6]). For every prefix $p \in \Gamma^*$, $f(p)$ is the unique word in Γ_c^* such that $\gamma^-(f(p)) = \{q \in \Gamma^* \mid p \not\preceq q\}$.

Lemma 2.5. For any $\mathcal{P}, \mathcal{P}_1, \mathcal{P}_2 \subseteq \Gamma^*$, $\mathcal{P}^- = \mathcal{P}_1^- \cup \mathcal{P}_2^-$ if $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2$.

Proof. If $p \in \mathcal{P}^-$, then there exists $q \in \mathcal{P}_1 \cup \mathcal{P}_2$ such that $p \preceq q$. In other words, either $q \in \mathcal{P}_1$ or $q \in \mathcal{P}_2$ such that $p \preceq q$. That is, $p \in \mathcal{P}_1^-$ or $p \in \mathcal{P}_2^-$. Hence, $p \in \mathcal{P}_1^- \cup \mathcal{P}_2^-$.

If $p \in \mathcal{P}_1^- \cup \mathcal{P}_2^-$, then either $p \in \mathcal{P}_1^-$ or $p \in \mathcal{P}_2^-$. That is, there exists $q \in \mathcal{P}_1$, or $q \in \mathcal{P}_2$, such that $p \preceq q$. In other words, there exists $q \in \mathcal{P}$ such that $p \preceq q$. Therefore, $p \in \mathcal{P}^-$. \square

2.4. Quantifier classes. Let φ be a FO formula. Recall that we assume that any formula is in negation normal form. Let \uplus be the disjoint union and $\bigsqcup_i S_i$ be the disjoint union of S_i .

If Φ is a set of formulas, then $\bigwedge \Phi$ ($\bigvee \Phi$ resp.) is the conjunction (disjunction resp.) of all the formulas in Φ . Similarly, we use $\bigwedge_i \theta_i$ ($\bigvee_i \theta_i$ resp.) to represent the conjunction (disjunction resp.) of all θ_i .

Definition 2.6. The *quantifier structure* of φ , denoted $qs(\varphi)$, is a Γ -labeled forest, which is defined inductively as follows:

- If φ is a literal, then $qs(\varphi)$ is empty;
 - If $\varphi = \bigwedge_i \theta_i$ or $\bigvee_i \theta_i$, then $qs(\varphi)$ is $\bigsqcup_i qs(\theta_i)$;
 - If $\varphi = \exists x\theta$, then $qs(\varphi)$ is composed of an \mathcal{E} node and $qs(\theta)$ where there is an arc from this \mathcal{E} node to each root of $qs(\theta)$ (note that $qs(\theta)$ is a forest);
Similarly, If $\varphi = \forall x\theta$, then $qs(\varphi)$ is composed of an \mathcal{A} node and $qs(\theta)$ where there is an arc from this \mathcal{A} node to each root of $qs(\theta)$;
- In these two cases, if $qs(\theta)$ is empty, then $qs(\varphi)$ contains a single node.

Note that this definition is different from Grädel and McColm's [2], in which $qs(\varphi)$ is defined as a set of strings:

- If ψ is a literal, then $qs(\psi) = \{\epsilon\}$ where ϵ is the empty word;
- If $\psi := \bigwedge \Phi$ or $\psi := \bigvee \Phi$ where Φ is a set of formulas, then

$$qs(\psi) := \bigcup_{\varphi \in \Phi} qs(\varphi);$$

- $\psi := \exists x_i \varphi$, then $qs(\varphi) := \exists * qs(\varphi)$; likewise, if $\psi := \forall x_i \varphi$, then $qs(\psi) := \forall * qs(\varphi)$.

Definition 2.7. Let S_1, S_2 be two Γ -labeled forests. Define $S_1 \preceq_e S_2$ if there is a mapping ι , *not necessarily injective*, from the nodes of S_1 to the nodes of S_2 such that v and $\iota(v)$ have the same label for any v , and there is a nontrivial directed path from $\iota(x)$ to $\iota(y)$ in S_2 if there is an arc from node x to node y in S_1 .

Remark 2.8. Note that the relation \preceq_e is not necessary antisymmetric. That is, there are non-isomorphic Γ -labeled forests S_1, S_2 such that $S_1 \preceq_e S_2$ and $S_2 \preceq_e S_1$.

Definition 2.9. Suppose that we are given a Γ -labeled forest S . For any path $P := (v_0, \dots, v_n)$ in the forest, there is a word (s_0, \dots, s_n) in Γ^* associated with it such that the node v_i is labeled with s_i . We say that this word, as well as all its subsequences, can be *read off* this Γ -labeled forest. Let $\mathscr{W}(S)$ be the set of words that can be read off the forest S .

Definition 2.10. Suppose that we are given a set $\mathcal{P} \subseteq \Gamma^*$. Let $\mathcal{P} = \mathcal{P}_{\exists} \cup \mathcal{P}_{\forall}$, where $\mathcal{P}_{\exists} = \exists * \mathcal{P}_1$ and $\mathcal{P}_{\forall} = \forall * \mathcal{P}_2$. These sets can be empty. We can inductively define a Γ -labeled forest $\mathcal{F}(\mathcal{P})$ as follows:

- (1) If \mathcal{P} is empty, then $\mathcal{F}(\mathcal{P})$ is empty, i.e. this forest contains no node.
- (2) Let S_1 be a Γ -labeled forest such that its root is an \mathcal{E} node and there is an arc from this root to each root of $\mathcal{F}(\mathcal{P}_1)$. Likewise, let S_2 be a Γ -labeled forest such that its root is an \mathcal{A} node and there is an arc from this root to each root of $\mathcal{F}(\mathcal{P}_2)$.
- (3) $\mathcal{F}(\mathcal{P})$ is the disjoint union of S_1 and S_2 .

Note that $\mathcal{F}(\mathcal{P})$ is composed of at most two trees, the roots of which have different labels.

Lemma 2.11. $\mathcal{W}(\mathcal{F}(\mathcal{P})) = \mathcal{P}^-$, for all $\mathcal{P} \subseteq \Gamma^*$.

Proof. The base case when $\mathcal{F}(\mathcal{P})$ is empty, i.e. when $rk(\mathcal{F}(\mathcal{P})) = 0$, is trivial.

Assume that it holds when $rk(\mathcal{F}(\mathcal{P})) \leq k$ for some $k \geq 0$.

Assume that

$$rk(\mathcal{F}(\mathcal{P})) = k + 1 \text{ and } \mathcal{P} = \mathcal{P}_{\exists} \cup \mathcal{P}_{\forall},$$

where

$$\mathcal{P}_{\exists} = \exists * \mathcal{P}_1 \text{ and } \mathcal{P}_{\forall} = \forall * \mathcal{P}_2.$$

Clearly,

$$rk(\mathcal{F}(\mathcal{P}_1)) \leq k \text{ and } rk(\mathcal{F}(\mathcal{P}_2)) \leq k.$$

According to Definition 2.10, $\mathcal{F}(\mathcal{P})$ is the disjoint union of $\mathcal{F}(\mathcal{P}_{\exists})$ and $\mathcal{F}(\mathcal{P}_{\forall})$. In other words, $\mathcal{W}(\mathcal{F}(\mathcal{P}))$ equals $\mathcal{W}(\mathcal{F}(\mathcal{P}_{\exists})) \cup \mathcal{W}(\mathcal{F}(\mathcal{P}_{\forall}))$, hence equals

$$(\exists * \mathcal{W}(\mathcal{F}(\mathcal{P}_1))) \cup \mathcal{W}(\mathcal{F}(\mathcal{P}_1)) \cup (\forall * \mathcal{W}(\mathcal{F}(\mathcal{P}_2))) \cup \mathcal{W}(\mathcal{F}(\mathcal{P}_2)),$$

and by assumption equals

$$(\exists * \mathcal{P}_1^-) \cup \mathcal{P}_1^- \cup (\forall * \mathcal{P}_2^-) \cup \mathcal{P}_2^- = \mathcal{P}_{\exists}^- \cup \mathcal{P}_{\forall}^-.$$

By Lemma 2.5, $\mathcal{P}^- = \mathcal{P}_{\exists}^- \cup \mathcal{P}_{\forall}^-$. Therefore, $\mathcal{W}(\mathcal{F}(\mathcal{P})) = \mathcal{P}^-$. \square

Remark 2.12. This lemma implies that, for any $p \in \mathcal{P}^- \subseteq \Gamma^*$, p can be read off from some path of $\mathcal{F}(\mathcal{P})$.

Lemma 2.13. For any Γ -labeled forest S and $\mathcal{P} \subseteq \Gamma^*$, $S \preceq_e \mathcal{F}(\mathcal{P})$ if $\mathcal{W}(S) \subseteq \mathcal{P}^-$.

Proof. The base case when S is empty is trivial.

Assume that it holds when $rk(S) \leq k$ where $k \geq 0$.

Let S be a Γ -labeled forest such that $rk(S) = k + 1$ and $\mathcal{W}(S) \subseteq \mathcal{P}^-$. S is a disjoint union of at most two forests S_{\exists} and S_{\forall} : the roots of S_{\exists} are all \mathcal{E} nodes and the roots of S_{\forall} are all \mathcal{A} nodes. Then $\mathcal{W}(S) = \mathcal{W}(S_{\exists}) \cup \mathcal{W}(S_{\forall})$. Note that a substructure of a forest is also a forest. Because $\mathcal{W}(S) \subseteq \mathcal{P}^-$ and Lemma 2.11, $\mathcal{W}(S) \subseteq \mathcal{W}(\mathcal{F}(\mathcal{P}))$. It means that there is a forest (substructure) \mathcal{F}_{\exists} of $\mathcal{F}(\mathcal{P})$ such that all its roots are \mathcal{E} nodes and that $\mathcal{W}(S_{\exists}) \subseteq \mathcal{W}(\mathcal{F}_{\exists})$. Likewise, there is a forest (substructure) \mathcal{F}_{\forall} of $\mathcal{F}(\mathcal{P})$ such that all its roots are \mathcal{A} nodes and that $\mathcal{W}(S_{\forall}) \subseteq \mathcal{W}(\mathcal{F}_{\forall})$. Note that \mathcal{F}_{\exists} and \mathcal{F}_{\forall} are not necessary disjoint.

Now, if we remove all the \mathcal{E} roots from S_{\exists} , we get a forest called S_1 . Similarly, if we remove all the \mathcal{A} roots from S_{\forall} , we get another forest called S_2 .

Likewise, if we remove all the \mathcal{E} roots from \mathcal{F}_\exists , we get a forest called \mathcal{F}_1 . Similarly, if we remove all the \mathcal{A} roots from \mathcal{F}_\forall , we get another forest called \mathcal{F}_2 .

Observe that $\mathcal{W}(S_1) \subseteq \mathcal{W}(\mathcal{F}_1) \subseteq \mathcal{W}(\mathcal{F}_1)^-$ and $rk(S_1) \leq k$. By assumption, $S_1 \preceq_e \mathcal{F}_1$. Let us denote the map that embeds S_1 to \mathcal{F}_1 as ι_1 . Likewise, $S_2 \preceq_e \mathcal{F}_2$ and the embedding map is denoted ι_2 . Note that the domains of ι_1 and ι_2 are different. Therefore, we can merge these two maps easily, i.e. let $\iota_0 = \iota_1 \cup \iota_2$. Note that S_1 (S_2 resp.) is embeddable to \mathcal{F}_1 (\mathcal{F}_2 resp.) through ι_0 . Now, we can extend the embedding map ι_0 to ι such that: (i) the father of any root r_1 of S_1 is mapped to the father of $\iota_0(r_1)$; (ii) the father of any root r_2 of S_2 is mapped to the father of $\iota_0(r_2)$. Therefore, S is embeddable to $\mathcal{F}(P)$ through ι , i.e. $S \preceq_e \mathcal{F}(P)$. \square

Remark 2.14. Lemma 2.11 and Lemma 2.13 tell us that $\mathcal{F}(P)$ is the “maximal” Γ -labeled forest (in the sense of embeddings) among all these forests, from which the set of words that can be read off is a subset of \mathcal{P}^- .

Define the *quantifier rank* of FO formula φ , denoted $qr(\varphi)$, to be $rk(qs(\varphi))$. Note that this definition is equivalent to the usual definition of quantifier rank (see for instance Libkin [4]). Let $\text{FO}[k] := \{\varphi \in \text{FO} \mid qr(\varphi) \leq k\}$.

Let S be a Γ -labeled forest. Define the quantifier class $\text{FO}\{S\}$ to be the set of queries that are definable by the set of first-order sentences $\{\theta \in \text{FO} \mid qs(\theta) \preceq_e S\}$.

A first-order formula is in prenex normal form if it is a single string of quantifiers followed by a quantifier free formula. Its quantifier prefix, which is obtained from this string of quantifiers by removing the variables in the string, corresponds to a Γ -labeled degenerate tree.

Given a prefix p , we define the prefix class $\text{FO}(p)$ as the set of FO sentences in prenex normal form such that for any $\psi \in \text{FO}(p)$, its prefix is a subsequence of p (Grädel and McColm, [2]). Grädel-McColm’s conjecture says that the prefix classes form a strict hierarchy: For any prefix p, q , $\text{FO}(p) \not\subseteq \text{FO}(q)$ if $p \not\preceq q$ over arbitrary structures. Rosen [6] confirmed this conjecture over infinite structures and called it the first-order prefix hierarchy. Similarly, we can define a hierarchy formed by quantifier classes, which can be called the first-order quantifier structure hierarchy. These two hierarchies are independent.

3. QUANTIFIER STRUCTURE HIERARCHY: THE FIRST OBSERVATION

In this section, we define a variant of Ehrenfeucht-Fraïssé games that characterizes quantifier classes and prove that those quantifier classes form a natural and strict hierarchy:

Theorem 3.1. *Let S_1 and S_2 be two Γ -labelled forests. Over the class of all digraphs, if $\mathcal{W}(S_1) \not\subseteq \mathcal{W}(S_2)$, then $\text{FO}\{S_1\} \not\subseteq \text{FO}\{S_2\}$.*

3.1. Games that characterize quantifier classes. Let S be a Γ -labelled forest. We define an asymmetric variant of the Ehrenfeucht-Fraïssé games as follows. Let $k \in \mathbb{N}$, σ contains k constant symbols. Let \mathfrak{A} and \mathfrak{B} be two σ -structures. Let the k -tuple \bar{u} be the interpretation of the constants in \mathfrak{A} and the k -tuple \bar{v} be the interpretation of the constants in \mathfrak{B} . The game $G_S(\mathfrak{A}, \mathfrak{B})$ is played by two players, called the spoiler and duplicator, on a game board consisting of S , \mathfrak{A} and \mathfrak{B} . At the beginning of the game, the spoiler picks a tree \mathcal{T} in the forest S and puts a token on the root of \mathcal{T} . Assume that the depth of \mathcal{T} is $n - 1$. Afterwards, for every i where $1 \leq i \leq n$, in the i -th round the spoiler chooses an

element from the structure \mathfrak{A} if the current node, on which the token is put, is an \mathcal{E} node. Otherwise if it is an \mathcal{A} node he picks an element of \mathfrak{B} . Then the duplicator has to respond by picking an element from the other structure. Afterwards the spoiler chooses a child of the current node in S and moves the token to it. This completes one round.

Assume that after n' ($n' \leq n$) rounds a sequence $\bar{c} = (c_1, \dots, c_{n'})$ has been picked in \mathfrak{A} and a sequence $\bar{d} = (d_1, \dots, d_{n'})$ has been picked in \mathfrak{B} . The spoiler wins the game if $(\bar{u}\bar{c}, \bar{v}\bar{d})$ does not define a partial isomorphism between \mathfrak{A} and \mathfrak{B} .

The game ends whenever the spoiler wins or the token arrives at a leaf of \mathcal{T} . The duplicator wins if the spoiler fails to win in the end.

Informally, \bar{u} and \bar{v} can be regarded as a carry-over of past history of the game played before the beginning, which has to be taken care of.

A strategy of the duplicator is a scheme by which she knows how to choose an element in each round depending on the history of the play.

For any tuple $(\mathfrak{a}_1, \dots, \mathfrak{a}_n) \in (|\mathfrak{A}| \uplus |\mathfrak{B}|)^n$, we associate it with a prefix $(p[1], \dots, p[n])$ such that $\mathfrak{a}_i \in |\mathfrak{A}|$ iff $p[i] = \exists$ for any $1 \leq i \leq n$. A *strategy of the duplicator* in the game $G_S(\mathfrak{A}, \mathfrak{B})$ is a function

$$D_{\mathfrak{A}, \mathfrak{B}}^S : \uplus_{i=1}^{h(S)+1} (|\mathfrak{A}| \uplus |\mathfrak{B}|)^i \rightarrow (|\mathfrak{A}| \uplus |\mathfrak{B}|).$$

If the duplicator has a strategy guiding her choices in the game that ensures her winning in the end no matter how the spoiler plays, we call this strategy a winning strategy of the duplicator. If there exists such a winning strategy for the duplicator in the game $G_S(\mathfrak{A}, \mathfrak{B})$ then we write $\mathfrak{A} \rightsquigarrow_S \mathfrak{B}$. The winning strategy of the spoiler can be defined dually because in our games either the spoiler or the duplicator has a winning strategy. Let $\bar{a} \in |\mathfrak{A}|^t$ and $\bar{b} \in |\mathfrak{B}|^t$. We use $(\mathfrak{A}, \bar{a}) \rightsquigarrow_S (\mathfrak{B}, \bar{b})$ to denote that the duplicator has a winning strategy, in which \bar{a} is picked in \mathfrak{A} , and \bar{b} is picked in \mathfrak{B} , before the game starts. Equivalently, we say that the spoiler has a winning strategy in the game $G_S((\mathfrak{A}, \bar{a}), (\mathfrak{B}, \bar{b}))$.

Note that the standard Ehrenfeucht-Fraïssé game $G_n(\mathfrak{A}, \mathfrak{B})$ (see [4]) is exactly the game $G_S(\mathfrak{A}, \mathfrak{B})$ where $S = \{*\mathcal{T}_n^\exists, *\mathcal{T}_n^\forall\}$.

Definition 3.2. Let $n \in \mathbb{N}$ and \bar{c} be an n -tuple of elements from $|\mathfrak{A}|$. Then for a Γ -labelled forest S , the QS - S n -type of \bar{c} over σ -structure \mathfrak{A} is defined as:

$$tp_n^S(\mathfrak{A}, \bar{c}) = \{\varphi(\bar{c}) \in \text{FO}\{S\} \mid \mathfrak{A} \models \varphi(\bar{c})\}.$$

The following lemma is well-known, cf. [4] for a simple explanation.

Lemma 3.3. For fixed $k, n \in \mathbb{N}$, there are only finitely many formulas, in n free variables, in $\text{FO}[k]$ up to logical equivalence.

Corollary 3.4. Let $n \in \mathbb{N}$ and S be a Γ -labelled forest, and let \bar{c} be an n -tuple of elements from $|\mathfrak{A}|$, there are only finitely many formulas in $tp_n^S(\mathfrak{A}, \bar{c})$ up to logical equivalence.

Let S be a Γ -labelled forest.

Definition 3.5. Let $[\mathfrak{A}, \bar{a}]$ be an expansion of \mathfrak{A} to $\sigma \cup \{\bar{c}\}$ such that \bar{a} interprets the tuple of constants \bar{c} in $[\mathfrak{A}, \bar{a}]$.

Lemma 3.6. Let S and S' be two Γ -labelled forests such that $S' \preceq_e S$. Let \mathfrak{A} and \mathfrak{B} be two structures over the same signature. If the duplicator has a winning strategy in the game $G_S(\mathfrak{A}, \mathfrak{B})$, then she also has a winning strategy in the game $G_{S'}(\mathfrak{A}, \mathfrak{B})$.

Proof. Note that the duplicator can mimic her winning strategy in the game $G_{S'}(\mathfrak{A}, \mathfrak{B})$ to play the game $G_{S'}(\mathfrak{A}, \mathfrak{B})$. And a subset of a partial isomorphism is still a partial isomorphism. \square

Remark 3.7. Lemma 3.6 tells us that if the duplicator has a winning strategy in $G_S(\mathfrak{A}, \mathfrak{B})$, she also has a winning strategy when the players are allowed to skip playing arbitrary rounds of the game. Lemma 3.6 also tells us that if the spoiler has a winning strategy in $G_{S'}(\mathfrak{A}, \mathfrak{B})$, he also has a winning strategy in $G_S(\mathfrak{A}, \mathfrak{B})$. In other words, quantifiers are logical resources that can be exploited by the spoiler to detect the difference between two structures in the games.

It is obvious that $[\mathfrak{A}, \bar{a}] \rightsquigarrow_S [\mathfrak{B}, \bar{b}]$ iff $(\mathfrak{A}, \bar{a}) \rightsquigarrow_S (\mathfrak{B}, \bar{b})$, because in each round of both the game $G_S([\mathfrak{A}, \bar{a}], [\mathfrak{B}, \bar{b}])$ and $G_S((\mathfrak{A}, \bar{a}), (\mathfrak{B}, \bar{b}))$, if there is a partial isomorphism between two structures in the former, then this partial isomorphism is also a partial isomorphism between two structures in the latter.

In the following, we prove a connection between the games just defined and quantifier classes, which is a variant of the result of Grädel and McColm [2].

Recall that \mathfrak{A} has constants that are interpreted by \bar{u} . And \mathfrak{B} has constants that are interpreted by \bar{v} . We assume that \bar{u} (\bar{v} resp.) and \bar{a} (\bar{b} resp.) do not share any element.

Theorem 3.8. *For arbitrary finite σ -structures $\mathfrak{A}, \mathfrak{B}$, two tuples $\bar{a} \in |\mathfrak{A}|^t$, $\bar{b} \in |\mathfrak{B}|^t$, and a Γ -labelled forest S , the following are equivalent:*

- (i) $[\mathfrak{A}, \bar{a}] \rightsquigarrow_S [\mathfrak{B}, \bar{b}]$;
- (ii) $tp_t^S(\mathfrak{A}, \bar{a}) \subseteq tp_t^S(\mathfrak{B}, \bar{b})$.

Proof. (i) \rightarrow (ii):

When S is an empty forest, i.e. $rk(S) = 0$, $tp_t^S(\mathfrak{A}, \bar{a}) \not\subseteq tp_t^S(\mathfrak{B}, \bar{b})$ means there is a quantifier-free formula $\eta(\bar{x})$ such that $(\mathfrak{A}, \bar{a}) \models \eta(\bar{x})$ but $(\mathfrak{B}, \bar{b}) \not\models \eta(\bar{x})$. Hence, the mapping from $\bar{u}\bar{a}$ to $\bar{v}\bar{b}$ does not define a partial isomorphism. In other words, the spoiler wins the game and $[\mathfrak{A}, \bar{a}] \not\rightsquigarrow_S [\mathfrak{B}, \bar{b}]$.

Assume that (i) \rightarrow (ii) when $rk(S) \leq k$ for $k \geq 0$.

Assume that $rk(S) = k + 1$ and S consists of m trees S_1, \dots, S_m . Suppose that (ii) is false. Let $\varphi(\bar{x}) \in \text{FO}\{S\}$ such that $(\mathfrak{A}, \bar{a}) \models \varphi(\bar{x})$ but $(\mathfrak{B}, \bar{b}) \not\models \varphi(\bar{x})$. Then φ is a first-order formula that is a disjunction or conjunction of formulas $\text{FO}\{S_i\}$ ($1 \leq i \leq m$). There must exist one disjunct or conjunct ψ such that $(\mathfrak{A}, \bar{a}) \models \psi(\bar{x})$ while $(\mathfrak{B}, \bar{b}) \not\models \psi(\bar{x})$, where $qs(\psi) \preceq_e S_i$ for some $1 \leq i \leq m$. By Lemma 3.6, we may assume that the spoiler moves the token from the root of S_i . Assume that the root of S_i is an \mathcal{E} node, then ψ has the form $\exists y \psi_1(\bar{x}y)$. Hence, there is $c \in |\mathfrak{A}|$ s.t. $(\mathfrak{A}, \bar{a}c) \models \psi_1(\bar{x}y)$. Then the spoiler can pick c , and no matter which element, say d , the duplicator picks, $\psi_1(\bar{x}y)$ distinguishes the pair $[\mathfrak{A}, \bar{a}c]$ and $[\mathfrak{B}, \bar{b}d]$, where the variables $\bar{x}y$ are assigned the values $\bar{a}c$ and $\bar{b}d$ respectively, because $(\mathfrak{B}, \bar{b}) \not\models \exists y \psi_1(\bar{x}y)$. By induction assumption the spoiler has a winning strategy over the game $G_{qs(\psi_1)}([\mathfrak{A}, \bar{a}c], [\mathfrak{B}, \bar{b}d])$. Similarly, if S_i is a tree whose root is an \mathcal{A} node, the spoiler can pick $d \in |\mathfrak{B}|$ such that for any $c \in |\mathfrak{A}|$ picked by the duplicator $\psi_1(\bar{x}y)$ distinguishes the pair $[\mathfrak{A}, \bar{a}c]$ and $[\mathfrak{B}, \bar{b}d]$ where the variables $\bar{x}y$ are assigned the values $\bar{a}c$ and $\bar{b}d$ respectively. In other words, the spoiler can show that there is an element d such that it makes $(\mathfrak{B}, \bar{b}d) \not\models \psi_1(\bar{x}y)$ while $(\mathfrak{A}, \bar{a}c) \models \psi_1(\bar{x}y)$ is always true. By induction assumption the spoiler has a winning strategy over the game $G_{qs(\psi_1)}([\mathfrak{A}, \bar{a}c], [\mathfrak{B}, \bar{b}d])$. Therefore, $[\mathfrak{A}, \bar{a}] \not\rightsquigarrow_S [\mathfrak{B}, \bar{b}]$.

(ii) \rightarrow (i):

According to the definition of the game, when $rk(S) = 0$, $[\mathfrak{A}, \bar{a}] \not\prec_S [\mathfrak{B}, \bar{b}]$ means that the mapping from $\bar{u}\bar{a}$ to $\bar{v}\bar{b}$ is not a partial isomorphism, which implies that there exists a quantifier-free FO formula $\xi(\bar{x})$ s.t. either $(\mathfrak{A}, \bar{a}) \models \xi(\bar{x})$ but $(\mathfrak{B}, \bar{b}) \not\models \xi(\bar{x})$, or $(\mathfrak{A}, \bar{a}) \not\models \xi(\bar{x})$ but $(\mathfrak{B}, \bar{b}) \models \xi(\bar{x})$. Let $\neg\xi(\bar{x})$ be the negation of $\xi(\bar{x})$. In the former case, it implies that $tp_i^S(\mathfrak{A}, \bar{a}) \not\subseteq tp_i^S(\mathfrak{B}, \bar{b})$; in the latter case, $(\mathfrak{A}, \bar{a}) \models \neg\xi(\bar{x})$ but $(\mathfrak{B}, \bar{b}) \not\models \neg\xi(\bar{x})$, which also implies that $tp_i^S(\mathfrak{A}, \bar{a}) \not\subseteq tp_i^S(\mathfrak{B}, \bar{b})$. All in all, $[\mathfrak{A}, \bar{a}] \not\prec_S [\mathfrak{B}, \bar{b}]$ implies $tp_i^S(\mathfrak{A}, \bar{a}) \not\subseteq tp_i^S(\mathfrak{B}, \bar{b})$ when $rk(S) = 0$.

Assume that (ii) \rightarrow (i) when $rk(S) = k$ for $k \geq 0$.

Now assume that S is composed of trees S_1, \dots, S_m and $rk(S) = k + 1$. Assume that (i) is false. Then over one of the trees the spoiler has a winning strategy. Hence, the spoiler can first pick this tree to play. If this tree's root is an \mathcal{E} node r , we can regard this tree as a digraph composed of r and a forest S' such that there is an arc from r to each root of S' . In the first round the spoiler can pick an element $c \in |\mathfrak{A}|$ such that no matter which element $d \in |\mathfrak{B}|$ the duplicator picks, $[\mathfrak{A}, \bar{a}c]$ and $[\mathfrak{B}, \bar{b}d]$ form the new game board over which the spoiler will win the game $G_{S'}([\mathfrak{A}, \bar{a}c], [\mathfrak{B}, \bar{b}d])$. By Corollary 3.4, there are only finitely many formulas in $tp_{i+1}^{S'}(\mathfrak{A}, \bar{a}c)$ up to logical equivalence. Let T/E be a set of formulas where each equivalent class in $tp_{i+1}^{S'}(\mathfrak{A}, \bar{a}c)$ has exactly one formula in T/E . Let $\varphi(\bar{x}y)$ be the conjunction of all the formulas in T/E . By the induction hypothesis, for any d there is a formula $\eta(\bar{x}y) \in \text{FO}\{S'\}$ such that $(\mathfrak{A}, \bar{a}c) \models \eta(\bar{x}y)$ but $(\mathfrak{B}, \bar{b}d) \not\models \eta(\bar{x}y)$. Note that η is equivalent to one formula in T/E . Hence, $(\mathfrak{A}, \bar{a}c) \models \varphi(\bar{x}y)$ but $(\mathfrak{B}, \bar{b}d) \not\models \varphi(\bar{x}y)$, for any d . In other words, $(\mathfrak{A}, \bar{a}) \models \exists y\varphi(\bar{x}y)$ but $(\mathfrak{B}, \bar{b}) \not\models \exists y\varphi(\bar{x}y)$. Note that $\exists y\varphi(\bar{x}y) \in \text{FO}\{S\}$. Therefore, $tp_i^S(\mathfrak{A}, \bar{a}) \not\subseteq tp_i^S(\mathfrak{B}, \bar{b})$.

The case when the tree picked by the spoiler in the first step is a tree whose root is an \mathcal{A} node can be proved similarly. \square

Corollary 3.9. *Let \mathcal{K} be a class of finite structures and S be a Γ -labelled forest. If there is $\mathfrak{A} \in \mathcal{K}$ and $\mathfrak{B} \notin \mathcal{K}$ such that $\mathfrak{A} \rightsquigarrow_S \mathfrak{B}$, then there is no first-order sentence φ such that $qs(\varphi) \preceq_e S$ and $\mathcal{K} = \text{Mod}(\varphi)$.*

3.2. Point-expansions.

Definition 3.10. Let \mathfrak{A} be a structure over signature σ_A and \mathcal{K} be a set $\{\mathfrak{C}_i\}_{i \in I}$ of finite structures indexed by a set I . Let σ_i be the signature of \mathfrak{C}_i for each $i \in I$ such that σ_i contains a special constant c_i that is called hook. Assume that no two signatures share a constant. Let $\mathfrak{J}_{\mathfrak{A}} : |\mathfrak{A}| \rightarrow \mathcal{K}$ be a total function. Define the point-expansion of \mathfrak{A} by $\mathfrak{J}_{\mathfrak{A}}$ over \mathcal{K} , denoted $\mathcal{E}_{\mathcal{K}}^{\mathfrak{J}_{\mathfrak{A}}}(\mathfrak{A})$, as follows:

- (1) The signature of $\mathcal{E}_{\mathcal{K}}^{\mathfrak{J}_{\mathfrak{A}}}(\mathfrak{A})$, denoted $\sigma_{\mathcal{E}}$, is composed of the union of σ_i and σ_A , except for the hook in σ_i for any i .
- (2) Let $M := \bigsqcup_{a \in |\mathfrak{A}|} \mathfrak{J}_{\mathfrak{A}}(a)$ and $|\mathcal{E}_{\mathcal{K}}^{\mathfrak{J}_{\mathfrak{A}}}(\mathfrak{A})| := |M|$.
- (3) Let M' be the set of elements that interpret the hooks in respective disjoint substructures.
- (4) There is a bijection $g : M' \rightarrow |\mathfrak{A}|$ such that for any k -tuple $\bar{v} = (v_1, \dots, v_k) \in |M'|^k$ and any k -ary relation $R \in \sigma_A$,

$$\bar{v} \in R^{\mathcal{E}_{\mathcal{K}}^{\mathfrak{J}_{\mathfrak{A}}}(\mathfrak{A})} \text{ iff } (g(v_1), \dots, g(v_k)) \in R^{\mathfrak{A}}.$$

In other words, M' induces in $\mathcal{E}_{\mathcal{K}}^{\downarrow \mathfrak{A}}(\mathfrak{A})$ an isomorphic copy of \mathfrak{A} . Let

$$R_g := \{\bar{v} \in |M'|^k \mid (g(v_1), \dots, g(v_k)) \in R^{\mathfrak{A}}\}.$$

(5) Let $\{\mathcal{X}_i^{\hat{\mathcal{R}}}\}$ be the set of structures in M whose signatures contain $\hat{\mathcal{R}}$.

$$\hat{\mathcal{R}}^{\mathcal{E}_{\mathcal{K}}^{\downarrow \mathfrak{A}}(\mathfrak{A})} := \begin{cases} \bigcup_i \hat{\mathcal{R}}^{\mathcal{X}_i^{\hat{\mathcal{R}}}}, & \hat{\mathcal{R}} \notin \sigma_A; \\ \bigcup_i \hat{\mathcal{R}}^{\mathcal{X}_i^{\hat{\mathcal{R}}}} \cup \hat{\mathcal{R}}_g, & \hat{\mathcal{R}} \in \sigma_A. \end{cases}$$

(6) For any constant $c \in \sigma_{\mathcal{E}}$,

$$c^{\mathcal{E}_{\mathcal{K}}^{\downarrow \mathfrak{A}}(\mathfrak{A})} := \begin{cases} c^{\mathcal{E}_i}, & c \in \sigma_i \setminus \sigma_A; \\ g^{-1}(c^{\mathfrak{A}}), & c \in \sigma_A. \end{cases}$$

Informally speaking, $\mathcal{E}_{\mathcal{K}}^{\downarrow \mathfrak{A}}(\mathfrak{A})$ is a structure that is obtained from \mathfrak{A} by substituting each element $a \in |\mathfrak{A}|$ with $\mathfrak{J}_{\mathfrak{A}}(a)$, identifying a with the hook in $\mathfrak{J}_{\mathfrak{A}}(a)$. A point-expansion of the structure \mathfrak{A} can also be regarded as the result of a process that “glues” a small substructure $\mathfrak{J}_{\mathfrak{A}}(a)$ at each element a (also at the hook of $\mathfrak{J}_{\mathfrak{A}}(a)$) of the “prototype” structure \mathfrak{A} . And each small substructure shares only one element with the prototype structure, i.e. the “point” where they are “glued” together.

Lemma 3.11. *Suppose that we are given a forest \mathcal{F} , two structures \mathfrak{A} , \mathfrak{B} , a finite set \mathcal{K} of structures and two mappings $\mathfrak{J}_{\mathfrak{A}}$ and $\mathfrak{J}_{\mathfrak{B}}$ that expand \mathfrak{A} and \mathfrak{B} over \mathcal{K} respectively. Then the duplicator has a winning strategy in the game $G_{\mathcal{F}}(\mathcal{E}_{\mathcal{K}}^{\downarrow \mathfrak{A}}(\mathfrak{A}), \mathcal{E}_{\mathcal{K}}^{\downarrow \mathfrak{B}}(\mathfrak{B}))$ if the following is true: the duplicator has a winning strategy $D_{\mathfrak{A}, \mathfrak{B}}^{\mathcal{F}}$ in the game $G_{\mathcal{F}}(\mathfrak{A}, \mathfrak{B})$ such that*

(1) *for any $a \in |\mathfrak{A}|$ and sequence of elements $\bar{s} \in (|\mathfrak{A}| \uplus |\mathfrak{B}|)^{|\bar{s}|}$ ($|\bar{s}| \leq h(\mathcal{F})$), the duplicator has a winning strategy*

$$D_{\mathfrak{J}_{\mathfrak{A}}(a), \mathfrak{J}_{\mathfrak{B}}(D_{\mathfrak{A}, \mathfrak{B}}^{\mathcal{F}}(\bar{s}a))}^{\mathcal{F}}$$

in the game $G_{\mathcal{F}}(\mathfrak{J}_{\mathfrak{A}}(a), \mathfrak{J}_{\mathfrak{B}}(D_{\mathfrak{A}, \mathfrak{B}}^{\mathcal{F}}(\bar{s}a)))$.

(2) *for any $b \in |\mathfrak{B}|$ and sequence of elements $\bar{s} \in (|\mathfrak{A}| \uplus |\mathfrak{B}|)^{|\bar{s}|}$ ($|\bar{s}| \leq h(\mathcal{F})$), the duplicator has a winning strategy*

$$D_{\mathfrak{J}_{\mathfrak{A}}(D_{\mathfrak{A}, \mathfrak{B}}^{\mathcal{F}}(\bar{s}b)), \mathfrak{J}_{\mathfrak{B}}(b)}^{\mathcal{F}}$$

in the game $G_{\mathcal{F}}(\mathfrak{J}_{\mathfrak{A}}(D_{\mathfrak{A}, \mathfrak{B}}^{\mathcal{F}}(\bar{s}b)), \mathfrak{J}_{\mathfrak{B}}(b))$.

Proof. One winning strategy of the duplicator in the game $G_{\mathcal{F}}(\mathcal{E}_{\mathcal{K}}^{\downarrow \mathfrak{A}}(\mathfrak{A}), \mathcal{E}_{\mathcal{K}}^{\downarrow \mathfrak{B}}(\mathfrak{B}))$ is the composition of her winning strategies in $G_{\mathcal{F}}(\mathfrak{A}, \mathfrak{B})$ and $G_{\mathcal{F}}(\mathfrak{C}_i^A, \mathfrak{C}_j^B)$ where $\mathfrak{C}_i^A, \mathfrak{C}_j^B \in \mathcal{K}$.

Let $\mathfrak{A}_{\mathfrak{J}_{\mathfrak{A}}}$ be the isomorphic copy of \mathfrak{A} in $\mathcal{E}_{\mathcal{K}}^{\downarrow \mathfrak{A}}(\mathfrak{A})$, whose elements interpret the hooks of disjoint substructures in \mathcal{K} . And let $\mathfrak{B}_{\mathfrak{J}_{\mathfrak{B}}}$ be the isomorphic copy of \mathfrak{B} in $\mathcal{E}_{\mathcal{K}}^{\downarrow \mathfrak{B}}(\mathfrak{B})$, whose elements interpret the hooks of disjoint substructures in \mathcal{K} . Assume that the spoiler has already picked a sequence \bar{s} of elements in the game $G_{\mathcal{F}}(\mathfrak{A}_{\mathfrak{J}_{\mathfrak{A}}}, \mathfrak{B}_{\mathfrak{J}_{\mathfrak{B}}})$:

(ii) When the spoiler picks an element in $\mathfrak{A}_{\mathfrak{J}_{\mathfrak{A}}}$ or $\mathfrak{B}_{\mathfrak{J}_{\mathfrak{B}}}$, the duplicator uses the strategy $D_{\mathfrak{A}, \mathfrak{B}}^{\mathcal{F}}$; otherwise:

(iii) For any $a \in |\mathfrak{A}_{\mathfrak{J}_{\mathfrak{A}}}|$, if the spoiler picks an element in $\mathfrak{J}_{\mathfrak{A}}(a)$, the duplicator uses the strategy $D_{\mathfrak{J}_{\mathfrak{A}}(a), \mathfrak{J}_{\mathfrak{B}}(D_{\mathfrak{A}, \mathfrak{B}}^{\mathcal{F}}(\bar{s}a))}^{\mathcal{F}}$;

- (iii) For any $b \in |\mathfrak{B}_{\mathfrak{I}_{\mathfrak{B}}}|$, if the spoiler picks an element in $\mathfrak{I}_{\mathfrak{B}}(b)$, the duplicator uses the strategy $D_{\mathfrak{I}_{\mathfrak{A}}(D_{\mathfrak{A},\mathfrak{B}}^{\mathcal{F}}(\bar{s}b)), \mathfrak{I}_{\mathfrak{B}}(b)}^{\mathcal{F}}$.

Clearly, partial isomorphisms are preserved under such compositions, which provides the duplicator a winning strategy in the game $G_{\mathcal{F}}(\mathcal{E}_{\mathcal{K}}^{\mathfrak{I}_{\mathfrak{A}}}(\mathfrak{A}), \mathcal{E}_{\mathcal{K}}^{\mathfrak{I}_{\mathfrak{B}}}(\mathfrak{B}))$, which can be regarded as a main game together with a series of subgames. \square

3.3. Strictness of the FO quantifier structure hierarchy.

Definition 3.12. Let \mathcal{I} be a structure over signature $\langle c_I \rangle$ where c_I is the hook constant and \mathcal{I} has only one element, which is used to interpret c_I .

Definition 3.13. Let $m \in \mathbb{N}^+$ and $p \in \Gamma^*$. Let $\tau^+ = \langle U, R, B, r \rangle$ where R, B are binary relation symbols, U is a unary relation symbol, and r is a constant symbol. To make it vivid, we say that an element x is *black* if $x \in U$. All the elements in the structures are white unless explicitly labeled black. Likewise, an arc (x, y) is red if $(x, y) \in R$; an arc is blue if $(x, y) \in B$. Let $\tau_0^+ = \tau^+ \setminus \{U\}$.

We define $\tilde{\mathfrak{A}}_m^p$ and $\tilde{\mathfrak{B}}_m^p$ to be τ^+ -structures as follows:

- $\tilde{\mathfrak{A}}_m^p$ and $\tilde{\mathfrak{B}}_m^p$ are trees, whose edges are either red or blue. The constant r is interpreted by the root of the respective trees.
- When $|p| = 1$, all edges in $\tilde{\mathfrak{A}}_m^p$ and $\tilde{\mathfrak{B}}_m^p$ are red.
- If $p = \exists$,
 - (1) $\tilde{\mathfrak{A}}_m^p$ is a depth 1 tree that has $2m + 1$ leaves. One of its leaves is black.
 - (2) $\tilde{\mathfrak{B}}_m^p$ is a depth 1 tree that has $2m$ leaves. *None* of them is black.
$$\mathcal{T}_{\exists, m}^{\mathfrak{A}} := \tilde{\mathfrak{A}}_m^{\exists} | \tau_0^+; \mathcal{T}_{\exists, m}^{\mathfrak{B}} := \tilde{\mathfrak{B}}_m^{\exists} | \tau_0^+.$$
- If $p = \forall$,
 - (1) $\tilde{\mathfrak{A}}_m^p$ is a depth 1 tree that has $2m$ leaves. All of them are black.
 - (2) $\tilde{\mathfrak{B}}_m^p$ is a depth 1 tree that has $2m + 1$ leaves. All are black except one.
$$\mathcal{T}_{\forall, m}^{\mathfrak{A}} := \tilde{\mathfrak{A}}_m^{\forall} | \tau_0^+; \mathcal{T}_{\forall, m}^{\mathfrak{B}} := \tilde{\mathfrak{B}}_m^{\forall} | \tau_0^+.$$

Note that $\mathcal{T}_{\exists, m}^{\mathfrak{A}} = \mathcal{T}_{\forall, m}^{\mathfrak{B}}$ and $\mathcal{T}_{\exists, m}^{\mathfrak{B}} = \mathcal{T}_{\forall, m}^{\mathfrak{A}}$.
- When $|p| > 1$:

Let $\tilde{\mathfrak{A}}_m^{-q}$ be the same as $\tilde{\mathfrak{A}}_m^q$ except that the colours of all the edges are exchanged, i.e. red is interchanged with blue. Let $\mathfrak{D}_{q, m}^{A, A}$, a $(\tau^+ \cup \{\mathfrak{e}\} \setminus \{r\})$ -structure, be built from a copy of $\tilde{\mathfrak{A}}_m^q$ and a copy of $\tilde{\mathfrak{A}}_m^{-q}$ where they are joined together at their roots. Call their shared root a *junction point*, which interprets the hook constant \mathfrak{e} . Similarly, define $\mathfrak{D}_{q, m}^{A, B}$ to be the join of $\tilde{\mathfrak{A}}_m^q$ and $\tilde{\mathfrak{B}}_m^{-q}$; $\mathfrak{D}_{q, m}^{B, A}$ be the join of $\tilde{\mathfrak{B}}_m^q$ and $\tilde{\mathfrak{A}}_m^{-q}$; $\mathfrak{D}_{q, m}^{B, B}$ be the join of $\tilde{\mathfrak{B}}_m^q$ and $\tilde{\mathfrak{B}}_m^{-q}$. By ‘‘copies’’ we mean disjoint copies. Let $\mathcal{K} := \{\mathcal{I}, \mathfrak{D}_{q, m}^{A, A}, \mathfrak{D}_{q, m}^{A, B}, \mathfrak{D}_{q, m}^{B, A}\}$.

– If $p = \exists q$,

- (1 $^{\exists}$) $\tilde{\mathfrak{A}}_m^p$ is a point-expansion of $\mathcal{T}_{\exists, m}^{\mathfrak{A}}$ over \mathcal{K} as follows: The root of $\mathcal{T}_{\exists, m}^{\mathfrak{A}}$ is expanded by \mathcal{I} . It is also called the ‘‘root’’ of $\mathcal{T}_{\exists, m}^{\mathfrak{A}}$ that interprets r . One of its leaves is expanded by a copy of $\mathfrak{D}_{q, m}^{A, A}$. m leaves are expanded by a copy of $\mathfrak{D}_{q, m}^{A, B}$. Another m leaves are expanded by a copy of $\mathfrak{D}_{q, m}^{B, A}$.

- (2 $^{\exists}$) $\tilde{\mathfrak{B}}_m^p$ is a point-expansion of $\mathcal{T}_{\exists, m}^{\mathfrak{B}}$ over \mathcal{K} as follows: The root of $\mathcal{T}_{\exists, m}^{\mathfrak{B}}$ is expanded by \mathcal{I} . It is also called the ‘‘root’’ of $\mathcal{T}_{\exists, m}^{\mathfrak{B}}$ that interprets r . m leaves are expanded by a copy of $\mathfrak{D}_{q, m}^{A, B}$. The other m leaves are expanded by a copy of $\mathfrak{D}_{q, m}^{B, A}$.

- If $p = \forall q$,
 - (1 \forall) The *lem* of $\tilde{\mathfrak{A}}_m^p$ is the same as the above *lem* (2 \exists) of $\tilde{\mathfrak{B}}_m^p$.
 - (2 \forall) The *lem* of $\tilde{\mathfrak{B}}_m^p$ is the same as the *lem* (1 \exists) of $\tilde{\mathfrak{A}}_m^p$ except that $\mathfrak{D}_{q,m}^{A,A}$ is replaced by a copy of $\mathfrak{D}_{q,m}^{B,B}$.

Remark 3.14. $\tilde{\mathfrak{A}}_m^p$ is the same as $\tilde{\mathfrak{B}}_m^p$ except that the “colours” of leaves are flipped: black is interchanged with *not* black.

The structures $\tilde{\mathfrak{A}}_m^p$ and $\tilde{\mathfrak{B}}_m^p$ are trees with coloured edges and nodes. Note that only a leaf could be black by this definition because the root of a tree is not black by default. See Figure 1 for example, where $p = \exists\exists$ and $m = 1$. Here, we use a solid line to represent a red edge and a dashed line to represent a blue edge.

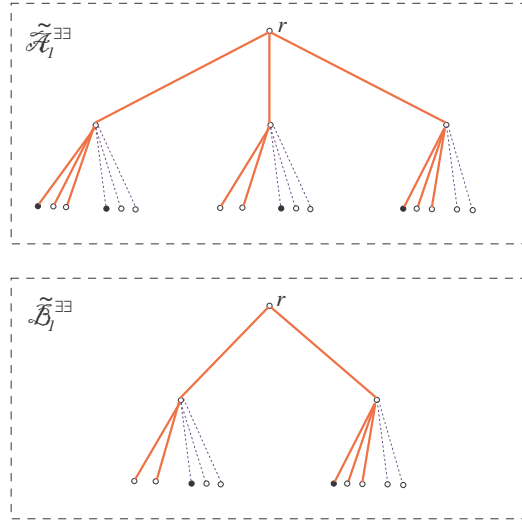


Figure 1: The structures $\tilde{\mathfrak{A}}_1^{\exists\exists}$ and $\tilde{\mathfrak{B}}_1^{\exists\exists}$.

We are going to define a formula $\tilde{\varphi}_p$ for each string $p \in \Gamma^*$.

Definition 3.15. Let $q \in \Gamma^*$ and assume that $|q| = d \geq 0$.

- (1) $\tilde{\psi}_\epsilon(x) = \tilde{\psi}_{-\epsilon}(x) := U(x)$;
- (2) $\tilde{\psi}_{\exists q}(y) := \exists x_{d+1}(Ryx_{d+1} \wedge \tilde{\psi}_q(x_{d+1}) \wedge \tilde{\psi}_{-q}(x_{d+1}))$;
 $\tilde{\psi}_{\forall q}(y) := \forall x_{d+1}(Ryx_{d+1} \rightarrow \tilde{\psi}_q(x_{d+1}) \vee \tilde{\psi}_{-q}(x_{d+1}))$;
- (3) $\tilde{\psi}_{-\exists q}(y) := \exists x_{d+1}(Byx_{d+1} \wedge \tilde{\psi}_q(x_{d+1}) \wedge \tilde{\psi}_{-q}(x_{d+1}))$;
 $\tilde{\psi}_{-\forall q}(y) := \forall x_{d+1}(Byx_{d+1} \rightarrow \tilde{\psi}_q(x_{d+1}) \vee \tilde{\psi}_{-q}(x_{d+1}))$;

Now we define a sentence over the signature τ^+ :

$$\tilde{\varphi}_p := \tilde{\psi}_p(r); \quad (3.1)$$

$$\tilde{\varphi}_{-p} := \tilde{\psi}_{-p}(r). \quad (3.2)$$

From now on, we assume that $m \in \mathbb{N}^+$ is an arbitrary natural number except where defined explicitly in context (see Theorem 3.23 for example).

Lemma 3.16. $\tilde{\mathfrak{A}}_m^p \models \tilde{\varphi}_p$ iff $\tilde{\mathfrak{A}}_m^{-p} \models \tilde{\varphi}_{-p}$.

Proof. Even though $\tilde{\mathfrak{A}}_m^{-p}$ is defined as a structure which is obtained from $\tilde{\mathfrak{A}}_m^p$ by exchanging all the colours of the edges, we will prove that $\tilde{\mathfrak{A}}_m^p$ can be turned into $\tilde{\mathfrak{A}}_m^{-p}$ if we only exchange the colours of the edges that connect to root (from red to blue and from blue to red) in $\tilde{\mathfrak{A}}_m^p$, and vice versa:

- (1) The base cases when $|p| = 1$ are trivial;
- (2) Assume that it holds when $|p| = k$;
- (3) Assume that $p = \exists q$ where $|q| = k$. According to the definition, the structure obtained from $\mathfrak{D}_{q,m}^{A,B}$ by exchanging the edge colours is $\mathfrak{D}_{q,m}^{B,A}$. Likewise, the structure obtained from $\mathfrak{D}_{q,m}^{B,A}$ by exchanging the edge colours is $\mathfrak{D}_{q,m}^{A,B}$. For the reason of symmetry, the structure obtained from $\mathfrak{D}_{q,m}^{A,A}$ by exchanging the edge colours is the same as $\mathfrak{D}_{q,m}^{A,A}$. Therefore, the conclusion holds according to the definition of $\tilde{\mathfrak{A}}_m^{\exists q}$.
- (4) Similarly, we can prove it holds when $p = \forall q$ where $|q| = k$. Therefore, it holds when $|p| = k + 1$.

Also note that $\tilde{\varphi}_p$ is the same as $\tilde{\varphi}_{-p}$ except that $\tilde{\varphi}_p$ claims that the edges that connect to root are red and $\tilde{\varphi}_{-p}$ claims that the edges that connect to root are blue. Therefore, $\tilde{\mathfrak{A}}_m^p \models \tilde{\varphi}_p$ iff $\tilde{\mathfrak{A}}_m^{-p} \models \tilde{\varphi}_{-p}$. \square

Lemma 3.17. $\tilde{\mathfrak{A}}_m^p \models \tilde{\varphi}_p$ and $\tilde{\mathfrak{B}}_m^p \not\models \tilde{\varphi}_p$.

Proof. It is obvious when $p = \exists$ or \forall . Assume that it holds when $|p| = k$.

If $p = \exists q$, $\tilde{\mathfrak{A}}_m^p$ is composed of a node a_r , a copy of $\mathfrak{D}_{q,m}^{A,A}$, m copies of $\mathfrak{D}_{q,m}^{A,B}$ and m copies of $\mathfrak{D}_{q,m}^{B,A}$; $\tilde{\varphi}_p = \exists x_{k+1} (Rrx_{k+1} \wedge \tilde{\psi}_q(x_{k+1}) \wedge \tilde{\psi}_{-q}(x_{k+1}))$. Let a_r interpret r in $\tilde{\mathfrak{A}}_m^p$, and the junction point b_r of $\mathfrak{D}_{q,m}^{A,A}$ witness $\exists x_{k+1}$ in $\tilde{\varphi}_p$. Note that b_r divides $\mathfrak{D}_{q,m}^{A,A}$ into two parts: one is an isomorphic copy of $\tilde{\mathfrak{A}}_m^q$ and the other is an isomorphic copy of $\tilde{\mathfrak{A}}_m^{-q}$. For convenience, we still use $\tilde{\mathfrak{A}}_m^q$ and $\tilde{\mathfrak{A}}_m^{-q}$ to denote these copies. Therefore, when r is interpreted as b_r , $\tilde{\mathfrak{A}}_m^q \models \tilde{\psi}_q(r)$ and $\tilde{\mathfrak{A}}_m^{-q} \models \tilde{\psi}_{-q}(r)$ (according to the induction hypothesis and Lemma 3.16).

Moreover, all the quantifiers in $\tilde{\psi}_q$ are relativized by relations either Ryx or Byx , where x is the quantified variable. And $\tilde{\psi}_q$ expresses some property that has nothing to do with the elements outside the tree substructure $\tilde{\mathfrak{A}}_m^q$. More precisely, in Definition 3.15, the variable “ y ” does not occur free in the formulas $\tilde{\psi}_q(x_{d+1})$ and $\tilde{\psi}_{-q}(x_{d+1})$. As a consequence, $\tilde{\mathfrak{A}}_m^q \models \tilde{\psi}_q(r)$ implies $\tilde{\mathfrak{A}}_m^p \models \tilde{\psi}_q(b_r)$.

By the same argument, $\tilde{\mathfrak{A}}_m^{-q} \models \tilde{\psi}_{-q}(r)$ implies $\tilde{\mathfrak{A}}_m^p \models \tilde{\psi}_{-q}(b_r)$. Therefore, $\tilde{\mathfrak{A}}_m^p \models \tilde{\varphi}_p$.

Let c_r be a child of $r^{\mathfrak{B}_m^p}$ in $\tilde{\mathfrak{B}}_m^p$. By assumption, $\tilde{\mathfrak{B}}_m^q \not\models \tilde{\psi}_q(r)$ ($\tilde{\mathfrak{B}}_m^{-q} \not\models \tilde{\psi}_{-q}(r)$ resp.) where c_r interprets r in $\tilde{\mathfrak{B}}_m^q$ ($\tilde{\mathfrak{B}}_m^{-q}$ resp.). As explained before, this means $\tilde{\mathfrak{B}}_m^p \not\models \tilde{\psi}_q(c_r)$ ($\tilde{\mathfrak{B}}_m^p \not\models \tilde{\psi}_{-q}(c_r)$ resp.). So, c_r cannot be a witness of $\exists x_{k+1}$. Therefore, $\tilde{\mathfrak{B}}_m^p \not\models \tilde{\varphi}_p$.

Similarly, we can prove that it also holds when $p = \forall q$. Hence, it holds when $|p| = k + 1$. \square

Let S and S' be two finite Γ -labelled forests. We collect some simple facts below.

Lemma 3.18. $\tilde{\mathfrak{A}}_m^p \rightsquigarrow_S \tilde{\mathfrak{B}}_m^p$ iff $\tilde{\mathfrak{A}}_m^{-p} \rightsquigarrow_S \tilde{\mathfrak{B}}_m^{-p}$.

Proof. According to the definitions, $\tilde{\mathfrak{A}}_m^{-p}$ ($\tilde{\mathfrak{B}}_m^{-p}$ resp.) is similar to $\tilde{\mathfrak{A}}_m^p$ ($\tilde{\mathfrak{B}}_m^p$ resp.) except that the colours of their edges are exchanged. Therefore, the duplicator can mimic her winning strategy in one game when she plays in the other. \square

Definition 3.19. Let $\mathfrak{A} \oplus_a \mathfrak{B}$, a structure over signature σ , be the join of two disjoint σ -substructures \mathfrak{A} and \mathfrak{B} at its element a . That is,

- (1) $|\mathfrak{A}| \cap |\mathfrak{B}| = \{a\}$;
- (2) $|\mathfrak{A} \oplus_a \mathfrak{B}| = |\mathfrak{A}| \cup |\mathfrak{B}|$.
- (3) For any $R \in \sigma$, $R^{\mathfrak{A} \oplus_a \mathfrak{B}} := R^{\mathfrak{A}} \cup R^{\mathfrak{B}}$.

Lemma 3.20. Let S be a Γ -labelled forest. $\mathfrak{A} \oplus_a \mathfrak{B} \rightsquigarrow_S \mathfrak{A} \oplus_b \mathfrak{B}'$ if there is an automorphism h of \mathfrak{A} s.t. $b = h(a)$ and $\mathfrak{B} \rightsquigarrow_S \mathfrak{B}'$.

Proof. Assume that both \mathfrak{B} and \mathfrak{B}' are σ -structures. Let $\sigma^+ := \sigma \cup \{c_B\}$ where c_B is called a hook. Let $\check{\mathfrak{B}}$ ($\check{\mathfrak{B}}'$ resp.) be an expansion of \mathfrak{B} (\mathfrak{B}' resp.) to σ^+ . Let $\mathcal{K} = \{\mathcal{I}, \check{\mathfrak{B}}\}$ and $\mathcal{E}_{\mathcal{K}}^{\mathfrak{J}}(\mathfrak{A})$ be a point-expansion of \mathfrak{A} over \mathcal{K} defined by \mathfrak{J} such that the element a of \mathfrak{A} is expanded by $\check{\mathfrak{B}}$ and all the other elements are expanded by \mathcal{I} . Similarly, let $\mathcal{K}' = \{\mathcal{I}, \check{\mathfrak{B}}'\}$ and $\mathcal{E}_{\mathcal{K}'}^{\mathfrak{J}'}$ be a point-expansion of \mathfrak{A} over \mathcal{K}' defined by \mathfrak{J}' such that the element b of \mathfrak{A} is expanded by $\check{\mathfrak{B}}'$ and all the other elements are expanded by \mathcal{I} . If there is an automorphism h of \mathfrak{A} s.t. $b = h(a)$, then by Lemma 3.11 the following holds:

$$\mathcal{E}_{\mathcal{K}}^{\mathfrak{J}}(\mathfrak{A}) \rightsquigarrow_S \mathcal{E}_{\mathcal{K}'}^{\mathfrak{J}'}(\mathfrak{A}) \text{ if } \mathfrak{B} \rightsquigarrow_S \mathfrak{B}'.$$

Observe that $\mathcal{E}_{\mathcal{K}}^{\mathfrak{J}}(\mathfrak{A})$ is exactly $\mathfrak{A} \oplus_a \mathfrak{B}$ and $\mathcal{E}_{\mathcal{K}'}^{\mathfrak{J}'}(\mathfrak{A})$ is exactly $\mathfrak{A} \oplus_b \mathfrak{B}'$. Hence, the lemma holds. \square




Figure 2: $(\mathfrak{D}_{p,m}^{A,A}, \bar{a})$ and $(\mathfrak{D}_{p,m}^{A,B}, \bar{b})$ where \bar{a} includes $r^1 \tilde{\mathfrak{A}}_m^p$ and \bar{b} includes $r^2 \tilde{\mathfrak{A}}_m^p$.

We use ${}^1 \tilde{\mathfrak{A}}_m^p$ and ${}^2 \tilde{\mathfrak{A}}_m^p$ to denote two isomorphic copies of $\tilde{\mathfrak{A}}_m^p$ in $\mathfrak{D}_{p,m}^{A,A}$ and $\mathfrak{D}_{p,m}^{A,B}$ respectively (see Figure 2. Note that in the picture we use “ ${}^1 A_p$ ” to denote ${}^1 \tilde{\mathfrak{A}}_m^p$). The superscripts “1” and “2” are used to distinguish these two copies. Let $\bar{a} \in |{}^1 \tilde{\mathfrak{A}}_m^p|^k$ and $\bar{b} \in |{}^2 \tilde{\mathfrak{A}}_m^p|^k$. The following lemma is a special case of Lemma 3.20.

Lemma 3.21. *Assume that $({}^1\tilde{\mathfrak{A}}_m^p, \bar{a}) \cong ({}^2\tilde{\mathfrak{A}}_m^p, \bar{b})$. Then,*

$$(\mathfrak{D}_{p,m}^{A,A}, \bar{a}) \rightsquigarrow_S (\mathfrak{D}_{p,m}^{A,B}, \bar{b}) \text{ if } \tilde{\mathfrak{A}}_m^{-p} \rightsquigarrow_S \tilde{\mathfrak{B}}_m^{-p}$$

Clearly, for two Γ -labelled forests S_1 and S_2 , $\mathscr{W}(S_1) \subseteq \mathscr{W}(S_2)$ if $S_1 \preceq_e S_2$. Let $f_m^p := \{q \in \gamma^-(f(p)) \mid |q| \leq m\}$.

Lemma 3.22. *Let $p \in \Gamma^*$ and $m \in \mathbb{N}^+$, then the following holds:*

- (1) *The duplicator has a winning strategy in $G_{\mathcal{F}(f_m^p)}(\tilde{\mathfrak{A}}_m^p, \tilde{\mathfrak{B}}_m^p)$;*
- (2) *For any ψ such that $\mathscr{W}(qs(\psi)) \subseteq f_m^p$, we have that*

$$\tilde{\mathfrak{A}}_m^p \models \psi \Rightarrow \tilde{\mathfrak{B}}_m^p \models \psi.$$

Proof. By Lemma 2.13, $qs(\psi) \preceq_e \mathcal{F}(f_m^p)$. Hence by Theorem 3.8 and Lemma 3.6, (2) is implied by (1) and we need only prove (1).

Note that $\tilde{\varphi}_p$ is obtained from $\neg\tilde{\varphi}_{\bar{p}}$ by adding a “ \neg ” before all the occurrences of the unary predicate U . And $\tilde{\varphi}_p$ is equivalent to a sentence in $\text{FO}\{f(p)\}$ iff $\neg\tilde{\varphi}_{\bar{p}}$ is equivalent to a sentence in $\text{FO}\{f(p)\}$ iff $\tilde{\varphi}_{\bar{p}}$ is equivalent to a sentence in $\text{FO}\{f(\bar{p})\}$ (the second “iff” is due to Lemma 2.3). That is, the duplicator has a winning strategy in $G_{\mathcal{F}(f_m^p)}(\tilde{\mathfrak{A}}_m^p, \tilde{\mathfrak{B}}_m^p)$ iff she has a winning strategy in $G_{\mathcal{F}(f_m^{\bar{p}})}(\tilde{\mathfrak{A}}_m^{\bar{p}}, \tilde{\mathfrak{B}}_m^{\bar{p}})$. Therefore, we need only consider the case when $p[1] = \exists$ since the lemma holds when $p[1] = \exists$ iff it holds when $p[1] = \forall$. Soon we shall see that the case when $p[2] = \forall$ is different from the case when $p[2] = \exists$, provided that $p[1] = \exists$. Therefore, we discuss them separately.

Let $|p| = d$. When $d = 1$ i.e. $p = \exists$, $f(p) = \forall^*$. In the game, the spoiler can only pick at most m distinct elements in $\tilde{\mathfrak{B}}_m^p$. And $\tilde{\mathfrak{A}}_m^p$ has $2m$ distinct elements that are not black. Hence, the duplicator is able to mimic the spoiler’s picking as follows. She picks the root of $\tilde{\mathfrak{A}}_m^p$ if the spoiler picks the root of $\tilde{\mathfrak{B}}_m^p$; she picks a leaf that is not black in $\tilde{\mathfrak{A}}_m^p$ if the spoiler picks a leaf of $\tilde{\mathfrak{B}}_m^p$. That is, the duplicator has a winning strategy in this game. Similarly, when $p = \forall$, the duplicator also win the game.

Assume that it holds when $d \leq k$ for some $k \geq 1$. That is, the duplicator has a winning strategy in the game $G_{\mathcal{F}(f_m^p)}(\tilde{\mathfrak{A}}_m^p, \tilde{\mathfrak{B}}_m^p)$ for any $|p| \leq k$ and any $m \in \mathbb{N}^+$.

Assume that $p = \exists\forall q$ where $|q| = k - 1$. Then $f(p) = \forall^* * f(\forall q)$. This case is relatively easy to explain.

The strategy of the duplicator in the game $G_{\mathcal{F}(f_m^{\forall q})}(\mathcal{T}_{\exists,m}^{\mathfrak{A}}, \mathcal{T}_{\exists,m}^{\mathfrak{B}})$ is very simple:

- (I) If the spoiler picks the junction point of $\mathfrak{D}_{\forall q,m}^{A,A}$ that is a leaf of $\mathcal{T}_{\exists,m}^{\mathfrak{A}}$, the duplicator replies with a junction point of one copy of $\mathfrak{D}_{\forall q,m}^{A,B}$, called \mathfrak{D}_A , that is a leaf of $\mathcal{T}_{\exists,m}^{\mathfrak{B}}$.
- (II) If the spoiler picks a junction point of $\mathfrak{D}_{\forall q,m}^{A,B}$ or $\mathfrak{D}_{\forall q,m}^{B,A}$ that is a leaf of one tree and that is not \mathfrak{D}_A , the duplicator replies with a junction point of an isomorphic copy that is a leaf of the other tree.
- (III) If the spoiler picks an element, say a , which has been picked before, the duplicator picks b , which was picked in the same round when a was picked.

We can regard the words that can be read off $\mathcal{F}(f_m^p)$ as the “resource” that the spoiler can use to detect the difference between the structures. Note that the first universal quantifier block in the words that can possibly be read off $\mathcal{F}(f_m^p)$ is useless for the spoiler: no matter how he picks in $\tilde{\mathfrak{B}}_m^p$, the duplicator always picks an isomorphic substructure and mimics the spoiler’s picks in the isomorphic substructure. By Lemma 3.20, if the spoiler can win in the end, he can also win if the players do not pick in these isomorphic substructures.

Let $Q := \{s \in f_j^{\forall q} \mid 0 \leq j \leq m\}$. Clearly, $Q \preceq f_m^{\forall q}$.

Observe that the strategy described before is a winning strategy for the duplicator in the game $G_{\mathcal{F}(f_m^{\forall q})}(\mathcal{T}_{\exists, m}^{\mathfrak{A}}, \mathcal{T}_{\exists, m}^{\mathfrak{B}})$. By induction hypothesis, she also has a winning strategy in the game $G_{\mathcal{F}(f_m^{\forall q})}(\tilde{\mathfrak{A}}_m^{\forall q}, \tilde{\mathfrak{B}}_m^{\forall q})$. Hence, by Lemma 3.21 and Lemma 3.18, $\mathfrak{D}_{\forall q, m}^{A, A} \rightsquigarrow_{f_m^{\forall q}} \mathfrak{D}_{\forall q, m}^{A, B}$. Moreover, the winning strategies of the duplicator in the games $G_{\mathcal{F}(f_m^{\forall q})}(\mathcal{T}_{\exists, m}^{\mathfrak{A}}, \mathcal{T}_{\exists, m}^{\mathfrak{B}})$ and $G_{\mathcal{F}(f_m^{\forall q})}(\mathfrak{D}_{\forall q, m}^{A, A}, \mathfrak{D}_{\forall q, m}^{A, B})$ can be combined together: if the spoiler picks a leaf of $\mathcal{T}_{\exists, m}^{\mathfrak{B}}$ ($\mathcal{T}_{\exists, m}^{\mathfrak{A}}$ resp.) that is the junction point of $\mathfrak{D}_{\forall q, m}^{A, B}$ or $\mathfrak{D}_{\forall q, m}^{B, A}$ (not \mathfrak{D}_A), the duplicator picks a leaf of $\mathcal{T}_{\exists, m}^{\mathfrak{A}}$ ($\mathcal{T}_{\exists, m}^{\mathfrak{B}}$ resp.) that is the junction point of an isomorphic structure; if the spoiler picks a leaf of $\mathcal{T}_{\exists, m}^{\mathfrak{A}}$ that is the junction point of $\mathfrak{D}_{\forall q, m}^{A, A}$, the duplicator picks a leaf of $\mathcal{T}_{\exists, m}^{\mathfrak{B}}$ that is the junction point of \mathfrak{D}_A . Therefore, she has a combined winning strategy in the game $G_{\mathcal{F}(f_m^{\forall q})}(\tilde{\mathfrak{A}}_m^p, \tilde{\mathfrak{B}}_m^p)$, by Lemma 3.11. By Lemma 3.6, she also has a winning strategy in the game $G_{\mathcal{F}(Q)}(\tilde{\mathfrak{A}}_m^p, \tilde{\mathfrak{B}}_m^p)$. Note that $f_m^p = \{s \in \Gamma^m \mid s = \forall^i * s' \text{ where } i \leq m \text{ and } s' \in Q\}$. In other words, if we remove the path initiating from a root in $\mathcal{F}(f_m^p)$ where the word that can be read off the path is \forall^m , we can turn $\mathcal{F}(f_m^p)$ into $\mathcal{F}(Q)$. Recall that the first universal quantifier block in the words that can be read off $\mathcal{F}(f_m^p)$ is useless for the spoiler. Therefore, the duplicator has a winning strategy in the game $G_{\mathcal{F}(f_m^p)}(\tilde{\mathfrak{A}}_m^p, \tilde{\mathfrak{B}}_m^p)$.

Assume that $p = \exists \exists q$ where $|q| = k - 1$ (see Figure 3). Then $f(p) = \forall^* \exists * f(\exists q)$. As we have explained before, the first universal quantifier block in the words that can possibly be read off $\mathcal{F}(f_m^p)$ is useless for the spoiler.

The strategy of the duplicator in the game $G_{\mathcal{F}(f_m^{\exists q})}(\mathcal{T}_{\exists, m}^{\mathfrak{A}}, \mathcal{T}_{\exists, m}^{\mathfrak{B}})$ is the same as her strategy in the game $G_{\mathcal{F}(f_m^{\forall q})}(\mathcal{T}_{\exists, m}^{\mathfrak{A}}, \mathcal{T}_{\exists, m}^{\mathfrak{B}})$.

For the first existential quantifier that can possibly be read off $\mathcal{F}(f_m^p)$, there are several choices for the spoiler:

- *Picking the root* (or *picking a junction point* resp.).

The game is reduced to a composition of the main game $G_{\mathcal{F}(f_m^{\exists q})}(\mathcal{T}_{\exists, m}^{\mathfrak{A}}, \mathcal{T}_{\exists, m}^{\mathfrak{B}})$ and the subgames in which the duplicator has a winning strategy, that is, in the subgames

$$G_{\mathcal{F}(f_m^{\exists q})}(\mathfrak{D}_{\exists q, m}^{A, A}, \mathfrak{D}_{\exists q, m}^{A, B}) \text{ or } G_{\mathcal{F}(f_m^{\exists q})}((\mathfrak{D}_{\exists q, m}^{A, A}, r^1 \tilde{\mathfrak{A}}_m^p), (\mathfrak{D}_{\exists q, m}^{A, B}, r^2 \tilde{\mathfrak{A}}_m^p)),$$

or subgames between isomorphic structures, according to Lemma 3.18 and Lemma 3.21. And by Lemma 3.11, the duplicator has a winning strategy.

- *Picking inside the structure*, either $\mathfrak{D}_{\exists q, m}^{A, A}$ or $\mathfrak{D}_{\exists q, m}^{A, B}$ or $\mathfrak{D}_{\exists q, m}^{B, A}$, except their junction points.

When the spoiler picks inside $\mathfrak{D}_{\exists q, m}^{A, B}$ or $\mathfrak{D}_{\exists q, m}^{B, A}$, the duplicator can mimic it in an isomorphic copy. Hence, only one new case need be taken care of: the spoiler picks the element inside $\mathfrak{D}_{\exists q, m}^{A, A}$. In this case, the duplicator has a strategy as follows:

- If the spoiler picks the element inside $\tilde{\mathfrak{A}}_m^{\exists q}$, the duplicator mimics his picking in the isomorphic copy that is a part of $\mathfrak{D}_{\exists q, m}^{A, B}$.
- If the spoiler picks the element inside $\tilde{\mathfrak{A}}_m^{-\exists q}$, the duplicator mimics his picking in the isomorphic copy that is a part of $\mathfrak{D}_{\exists q, m}^{B, A}$.

According to Lemma 3.20, picking in these isomorphic substructures doesn't influence the outcome. In either case, the game is reduced to a composition of the main game $G_{\mathcal{F}(f_m^{\exists q})}(\mathcal{T}_{\exists, m}^{\mathfrak{A}}, \mathcal{T}_{\exists, m}^{\mathfrak{B}})$ and the subgame $G_{\mathcal{F}(f_m^{\exists q})}(\tilde{\mathfrak{A}}_m^{\exists q}, \tilde{\mathfrak{B}}_m^{\exists q})$ or $G_{\mathcal{F}(f_m^{\exists q})}(\tilde{\mathfrak{A}}_m^{-\exists q}, \tilde{\mathfrak{B}}_m^{-\exists q})$ or a

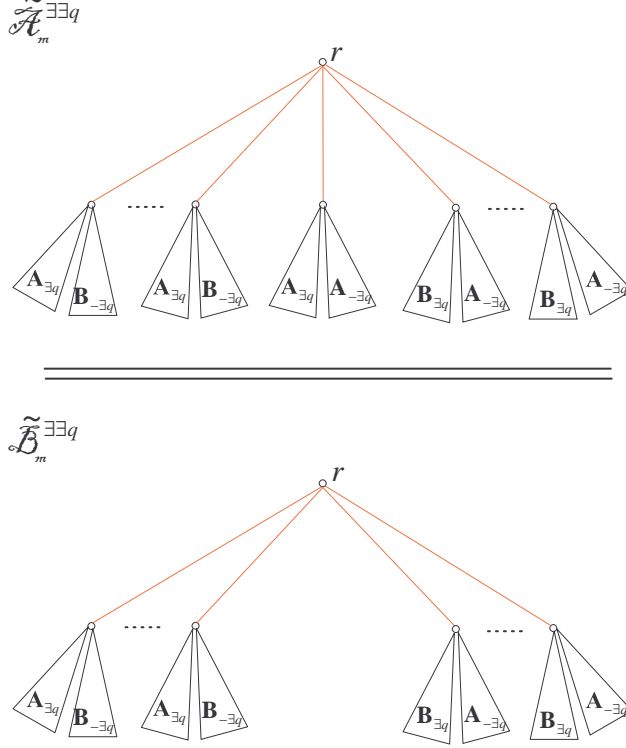


Figure 3: The structures $\tilde{\mathfrak{A}}_m^{\exists\exists q}$ and $\tilde{\mathfrak{B}}_m^{\exists\exists q}$.

subgame between two isomorphic structures. By the assumption we know that the duplicator has a winning strategy in $G_{\mathcal{F}(f_m^{\exists q})}(\tilde{\mathfrak{A}}_m^{\exists q}, \tilde{\mathfrak{B}}_m^{\exists q})$. Furthermore, by Lemma 3.18, the duplicator has a winning strategy in $G_{\mathcal{F}(f_m^{\exists q})}(\tilde{\mathfrak{A}}_m^{-\exists q}, \tilde{\mathfrak{B}}_m^{-\exists q})$. Finally, by Lemma 3.21, Lemma 3.6 and Lemma 3.11, the duplicator has a winning strategy in $G_{\mathcal{F}(f_m^p)}(\tilde{\mathfrak{A}}_m^p, \tilde{\mathfrak{B}}_m^p)$.

When $p[1] = \exists$, we have proved that $\tilde{\mathfrak{A}}_m^p \rightsquigarrow_{f_m^p} \tilde{\mathfrak{B}}_m^p$. Therefore, this theorem follows when $p[1] = \exists$. As a consequence it also holds when $p[1] = \forall$ by previous analysis. \square

Theorem 3.23. *Let S_1 and S_2 be two finite Γ -labelled forests. Over the class of all finite τ^+ -structures,*

$$\text{if } \mathcal{W}(S_1) \not\subseteq \mathcal{W}(S_2), \text{ then } \text{FO}\{S_1\} \not\subseteq \text{FO}\{S_2\}.$$

Proof. Let $p \in \mathcal{W}(S_1) \setminus \mathcal{W}(S_2)$. According to Lemma 2.4, $\gamma^-(f(p))$ is the set of all prefixes that p is not a subsequence of. Hence, $\mathcal{W}(S_2) \subseteq \gamma^-(f(p))$. Clearly, $\tilde{\varphi}_p \in \text{FO}\{S_1\}$. By Lemma 3.17, for any positive natural number n , $\tilde{\mathfrak{A}}_n^p \models \tilde{\varphi}_p$ but $\tilde{\mathfrak{B}}_n^p \not\models \tilde{\varphi}_p$. Assume for the purpose of a contradiction that there is a formula ψ such that $qs(\psi) \preceq_e S_2$, $qr(\psi) = m$ and ψ defines the same property as $\tilde{\varphi}_p$. Clearly, $\mathcal{W}(qs(\psi)) \preceq f_m^p$. By Lemma 3.22, $\tilde{\mathfrak{A}}_m^p \models \psi \Rightarrow \tilde{\mathfrak{B}}_m^p \models \psi$.

Together with Lemma 3.17, and Corollary 3.9, the property defined by $\tilde{\varphi}_p$ is not definable by any first-order sentence whose quantifier structure is \mathcal{F} s.t. $\mathcal{W}(\mathcal{F}) \subseteq f_m^p$, which is in contradiction with the assumption that ψ defines the same property as $\tilde{\varphi}_p$ does. \square

If we want to prove something similar to Theorem 3.23, but over a restricted signature $\langle r, E \rangle$, then we need to adapt the constructions and formulas a bit. Note that we can use forward arrows and backward arrows to replace the red edges and blue edges. And we can use bi-directional edges to indicate where the black leaves are. More precisely, the new structure \mathfrak{A}_m^p (\mathfrak{B}_m^p resp.) is obtained from $\tilde{\mathfrak{A}}_m^p$ ($\tilde{\mathfrak{B}}_m^p$ resp.) by the following process:

- (1) Use an arc from b to c to represent that the edge between the vertices b and c is red;
- (2) Use an arc from c to b to represent that the edge between b and c is blue;
- (3) Add an edge from every black leaf to the junction point in the same connected component and from the junction point to every black leaf in the same connected component;
- (4) When $|p| = 2$, add a self loop to every leaf which is an endpoint of a red edge.

Correspondingly, the new formula, called φ'_p , is obtained from $\tilde{\varphi}_p$ by the following process:

- (1*) Rxy is replaced by Exy ;
- (2*) Bxy is replaced by Eyx ;
- (3*) When $|p| \neq 2$, Uy is replaced by $Ex_1y \wedge Eyx_1$;
- (4*) When $|p| = 2$, Ux_1 in $\tilde{\psi}_{\exists}(y)$ and $\tilde{\psi}_{\forall}(y)$ is replaced by $Ex_1y \wedge Eyx_1 \wedge Ex_1x_1$; Ux_1 in $\tilde{\psi}_{-\exists}(y)$ and $\tilde{\psi}_{-\forall}(y)$ is replaced by $Ex_1y \wedge Eyx_1 \wedge \neg Ex_1x_1$.

More precisely, we inductively define φ'_p as follows. Let $q \in \Gamma^*$ and assume $|q| = d \geq 0$.

- (1) $\psi'_\epsilon(x, y) = \psi'_{-\epsilon}(x, y) := Exy \wedge Eyx$;
- (2) $\psi'_{\exists q}(x, y) := \exists x_{d+1}(Eyx_{d+1} \wedge x_{d+1} \neq x \wedge \psi'_q(y, x_{d+1}) \wedge \psi'_{-q}(y, x_{d+1}))$;
 $\psi'_{\forall q}(x, y) := \forall x_{d+1}(Eyx_{d+1} \wedge x_{d+1} \neq x \rightarrow \psi'_q(y, x_{d+1}) \vee \psi'_{-q}(y, x_{d+1}))$;
- (3) $\psi'_{-\exists q}(x, y) := \exists x_{d+1}(Ex_{d+1}y \wedge x_{d+1} \neq x \wedge \psi'_q(y, x_{d+1}) \wedge \psi'_{-q}(y, x_{d+1}))$;
 $\psi'_{-\forall q}(x, y) := \forall x_{d+1}(Ex_{d+1}y \wedge x_{d+1} \neq x \rightarrow \psi'_q(y, x_{d+1}) \vee \psi'_{-q}(y, x_{d+1}))$;
- (4) Now we define a sentence φ'_p over the signature τ :
 - When $|p| \neq 2$,

$$\varphi'_p := \psi'_p(r, r).$$
 (note that $\varphi'_{-p} := \psi'_{-p}(r, r)$).
 - When $|p| = 2$,

$$\varphi'_{\exists\exists} := \exists x_2(Ex_2x_2 \wedge \exists x_1(Ex_2x_1 \wedge Ex_1x_2 \wedge Ex_1x_1) \wedge \exists x_1(Ex_1x_2 \wedge Ex_2x_1 \wedge \neg Ex_1x_1))$$

$$\varphi'_{\exists\forall} := \exists x_2(Ex_2x_2 \wedge \forall x_1(Ex_2x_1 \rightarrow Ex_1x_2 \wedge Ex_1x_1) \wedge \forall x_1(Ex_1x_2 \rightarrow Ex_2x_1 \wedge \neg Ex_1x_1))$$

$$\varphi'_{\forall\exists} := \forall x_2(Ex_2x_2 \rightarrow \exists x_1(Ex_2x_1 \wedge Ex_1x_2 \wedge Ex_1x_1) \vee \exists x_1(Ex_1x_2 \wedge Ex_2x_1 \wedge \neg Ex_1x_1))$$

$$\varphi'_{\forall\forall} := \forall x_2(Ex_2x_2 \rightarrow \forall x_1(Ex_2x_1 \rightarrow Ex_1x_2 \wedge Ex_1x_1) \vee \forall x_1(Ex_1x_2 \rightarrow Ex_2x_1 \wedge \neg Ex_1x_1)).$$

Let $\tau := \langle E \rangle$ where E is a binary relation symbol. Could we prove something similar to Theorem 3.23, but over τ ? To achieve it, in addition to the adaption we have introduced above, we need to find a way to get rid of the root, which is obvious. The τ -structure \mathfrak{A}_m^p is obtained from \mathfrak{A}_m^p by the following process: removing $r^{\mathfrak{A}_m^p}$ and for any junction point v , using a self loop at node v to replace the edge from $r^{\mathfrak{A}_m^p}$ to v .

\mathfrak{B}_m^p is obtained from \mathfrak{B}_m^p in exactly the same way.

Correspondingly, φ_p is obtained from φ'_p by substituting Rrx with Exx and removing the atoms $x \neq r$.

Remark 3.24. It is a special case when $|p| = 2$ because using bi-directional edges as a scheme of colouring does not work in this case. See Figures 4 for example, when $p = \exists\exists$ and $m = 1$.




Figure 4: The structures $\mathfrak{A}_1^{\exists\exists}$ and $\mathfrak{B}_1^{\exists\exists}$.

Call the above reductions between structures and formulas “reductions from τ^+ to τ ”.

In a similar way to Lemma 3.17, we can prove the following lemma:

Lemma 3.25. $\mathfrak{A}_m^p \models \varphi_p$ and $\mathfrak{B}_m^p \not\models \varphi_p$.

Lemma 3.26. For any first-order sentence ζ over τ , there is a first-order sentence ξ over τ^+ , with the same quantifier structure, such that the following hold:

- (1) $\mathfrak{A}_m^p \models \zeta$ iff $\tilde{\mathfrak{A}}_m^p \models \xi$;
- (2) $\mathfrak{B}_m^p \models \zeta$ iff $\tilde{\mathfrak{B}}_m^p \models \xi$.

Proof. Let ξ be obtained from ζ by

- (a) relativising all quantifiers in ζ by $x \neq r$;
- (b) replacing all occurrences of Exy by

$$Rxy \vee Byx \vee (Rrx \wedge x = y) \vee (Rrx \wedge U(y)) \vee (Rry \wedge U(x)).$$

Note that the quantifier structure of ξ is the same as that of ζ .

Because \mathfrak{A}_m^p is obtained from $\tilde{\mathfrak{A}}_m^p$ by:

- deleting the root (corresponding to relativising quantifiers);
- substituting red edges with forward arcs (corresponding to the disjunct Rxy in the replacing of Exy);

- substituting blue edges with backward arcs (corresponding to the disjunct Bxy in the replacing of Exy);
- adding self-loops at the junction points that are connected to the root (corresponding to the disjunct $Rrx \wedge x = y$ in the replacing of Exy);
- adding bi-directional edges between the black leaves and junction points that are in the same connected component (corresponding to the disjuncts $(Rrx \wedge U(y))$ and $(Rry \wedge U(x))$ in the replacing of Exy),

it gives a reduction from the property defined by ξ to the property defined by ζ . In other words, $\mathfrak{A}_m^p \models \zeta$ iff $\tilde{\mathfrak{A}}_m^p \models \xi$. For the same reason, (2) also holds. \square

Now, we prove the main result in this section, i.e. the theorem 3.1.

Proof. Let $p \in \mathscr{W}(S_1) \setminus \mathscr{W}(S_2)$. Clearly, the property definable by φ_p is in $\text{FO}\{S_1\}$. We try to prove that this property is not in $\text{FO}\{S_2\}$. Assume on the contrary that there is a formula ψ such that $qs(\psi) \preceq_e S_2$ and ψ defines the same property as φ_p does. And let $qr(\psi)$ be m . According to Lemma 3.25, $\mathfrak{A}_m^p \models \psi$ and $\mathfrak{B}_m^p \not\models \psi$. According to Lemma 3.26, there is ξ , with the same quantifier structure as ψ , such that $\tilde{\mathfrak{A}}_m^p \models \xi$ and $\tilde{\mathfrak{B}}_m^p \not\models \xi$. Note that $\mathscr{W}(qs(\xi)) \subseteq \gamma^-(f(p))$ since $qs(\xi) \preceq_e S_2$ and $\mathscr{W}(S_2) \subseteq \gamma^-(f(p))$. But this is in contradiction to Lemma 3.22. \square

Theorem 3.1 tells us that the distinctive collections of quantifier classes form a strict hierarchy, which we call quantifier structure hierarchy.

4. STRICTNESS OF QUANTIFIER HIERARCHY OVER ORDERED FINITE STRUCTURES

Up to now, using logics to characterize complexity classes inside NP requires the structures to be ordered, i.e. there is a linear order over the universe of the structures. Therefore, it is interesting to extend the main result in the last section to ordered structures: the first-order quantifier structure hierarchy is strict over ordered finite structures. However, separating the expressive power of logics over ordered structures is often difficult, because the spoiler may detect the difference between the structures using a given linear order. But we will see in this section that the structures will be constructed in such a way that the power of linear order that the spoiler can use is quite limited: it is equivalent to the power that the spoiler can use in a game over a pair of linear orders, and a well-known result tells us that the duplicator has a winning strategy over a game between two linear orders that are sufficiently long. In this section we sketch the main ideas that conquer the order problem and omit most details that resemble those in the formal proof of Theorem 3.1.

4.1. The constructions and separating property.

Definition 4.1. Recall that $\tau^{+\text{ORD}}$ is $\langle R, B, U, r, \leq \rangle$ where \leq is interpreted as a linear order over the universe. Let $\tau'_o := \tau^{+\text{ORD}} \setminus \{U\}$.

Let the structure $\vec{\mathfrak{A}}_m^p$ be a $\tau^{+\text{ORD}}$ -structure defined as follows:

- $\vec{\mathfrak{A}}_m^p|_{\tau^+}$ and $\vec{\mathfrak{B}}_m^p|_{\tau^+}$ are trees, whose roots interpret r , and whose edges are coloured either red or blue.
- If $p = \exists$,

- (1) $\vec{\mathfrak{A}}_m^p|\tau^+$ is a depth 1 tree that has $2^{m+2} + 1$ leaves. To construct $\vec{\mathfrak{A}}_m^p$ from $\vec{\mathfrak{A}}_m^p|\tau^+$, give these leaves some order such that the $(2^{m+1} + 1)$ -th leaf is black. All the other leaves are not black.
- (2) $\vec{\mathfrak{B}}_m^p|\tau^+$ is a depth 1 tree that has 2^{m+1} leaves. *None* of them is black. Give these leaves an arbitrary order to construct $\vec{\mathfrak{B}}_m^p$ from $\vec{\mathfrak{B}}_m^p|\tau^+$.

Let $\vec{\mathcal{T}}_{\exists,m}^{\mathfrak{A}} := \vec{\mathfrak{A}}_m^{\exists}|\tau'_o$; $\vec{\mathcal{T}}_{\exists,m}^{\mathfrak{B}} := \vec{\mathfrak{B}}_m^{\exists}|\tau'_o$.

- If $p = \forall$,
 - (1) $\vec{\mathfrak{A}}_m^p|\tau^+$ is a depth 1 tree that has 2^{m+1} leaves. All of them are black.
 - (2) $\vec{\mathfrak{B}}_m^p|\tau^+$ is a depth 1 tree that has $(2^{m+2} + 1)$ leaves. All are black except one. Give these leaves some order such that the $(2^{m+1} + 1)$ -th leaf is not black.

In the above definition, we let a node be earlier in the linear order \leq than its children. Moreover, we define all the colours of the edges to be red when $|p| = 1$.

- When $|p| > 1$:

Let $\vec{\mathfrak{A}}_m^{-q}$ be the same as $\vec{\mathfrak{A}}_m^q$ except that the colours of all the edges are exchanged, i.e. red is changed to blue and vice versa. Let $\vec{\mathfrak{D}}_{q,m}^{A,B}$, a $(\tau^{+\text{ORD}} \cup \{\epsilon\} \setminus \{r\})$ -structure, be built from a copy of $\vec{\mathfrak{A}}_m^q$ and a copy of $\vec{\mathfrak{B}}_m^{-q}$ where they are joined together at their roots. Their shared root is called junction point, which is used to interpret the constant ϵ . Similarly define $\vec{\mathfrak{D}}_{q,m}^{A,A}$ as the join of $\vec{\mathfrak{A}}_m^q$ and $\vec{\mathfrak{A}}_m^{-q}$, define $\vec{\mathfrak{D}}_{q,m}^{B,A}$ as the join of $\vec{\mathfrak{B}}_m^q$ and $\vec{\mathfrak{A}}_m^{-q}$, and define $\vec{\mathfrak{D}}_{q,m}^{B,B}$ as the join of $\vec{\mathfrak{B}}_m^q$ and $\vec{\mathfrak{B}}_m^{-q}$. In $\vec{\mathfrak{D}}_{q,m}^{A,B}$, we let the elements in $\vec{\mathfrak{A}}_m^q$ be later in the linear order than the elements in $\vec{\mathfrak{B}}_m^{-q}$. Similarly, in $\vec{\mathfrak{D}}_{q,m}^{A,A}$, we let the elements in $\vec{\mathfrak{A}}_m^q$ be later in the order than the elements in $\vec{\mathfrak{A}}_m^{-q}$; in $\vec{\mathfrak{D}}_{q,m}^{B,A}$, the elements in $\vec{\mathfrak{B}}_m^q$ are later in the order than the elements in $\vec{\mathfrak{A}}_m^{-q}$; in $\vec{\mathfrak{D}}_{q,m}^{B,B}$, the elements in $\vec{\mathfrak{B}}_m^q$ are later in the order than the elements in $\vec{\mathfrak{B}}_m^{-q}$.

In the following, we assume that any node in $\vec{\mathfrak{D}}_{q,m}^{X_1,Y_1}$ is later in the order than all nodes in $\vec{\mathfrak{D}}_{q,m}^{X_2,Y_2}$ if the junction point of $\vec{\mathfrak{D}}_{q,m}^{X_1,Y_1}$ is later in the order than the junction point of $\vec{\mathfrak{D}}_{q,m}^{X_2,Y_2}$.

- If $p = \exists q$,

(1 \exists) $\vec{\mathfrak{A}}_m^p$ is a point-expansion of $\vec{\mathcal{T}}_{\exists,m}^{\mathfrak{A}}$ over $\{\mathcal{I}, \vec{\mathfrak{D}}_{q,m}^{A,A}, \vec{\mathfrak{D}}_{q,m}^{A,B}, \vec{\mathfrak{D}}_{q,m}^{B,A}\}$. Its root is expanded by \mathcal{I} . Recall that \mathcal{I} is a structure over the signature $\langle c_I \rangle$, whose universe has exactly one element, and this element is used to interpret the constant c_I . When $i \leq 2^{m+1}$, the i -th leaf is expanded by a copy of $\vec{\mathfrak{D}}_{q,m}^{A,B}$ if i is odd, by a copy of $\vec{\mathfrak{D}}_{q,m}^{B,A}$ if i is even; the i -th leaf is expanded by a copy of $\vec{\mathfrak{D}}_{q,m}^{A,A}$ if $i = 2^{m+1} + 1$; when $i > 2^{m+1} + 1$, the i -th leaf is expanded by a copy of $\vec{\mathfrak{D}}_{q,m}^{A,B}$ if $i - (2^{m+1} + 1)$ is odd, by a copy of $\vec{\mathfrak{D}}_{q,m}^{B,A}$ if $i - (2^{m+1} + 1)$ is even.

(2 \exists) $\vec{\mathfrak{B}}_m^p$ is a point-expansion of $\vec{\mathcal{T}}_{\exists,m}^{\mathfrak{B}}$ over $\{\mathcal{I}, \vec{\mathfrak{D}}_{q,m}^{A,B}, \vec{\mathfrak{D}}_{q,m}^{B,A}\}$. Its root is expanded by \mathcal{I} . The i -th leaf is expanded by a copy of $\vec{\mathfrak{D}}_{q,m}^{A,B}$ if i is odd, by a copy of $\vec{\mathfrak{D}}_{q,m}^{B,A}$ if i is even.

- If $p = \forall q$,

11 \forall) The *lem* of $\vec{\mathfrak{A}}_m^p$ is the same as the above *lem* (2 \exists) of $\vec{\mathfrak{B}}_m^{\exists q}$.

21 \forall) The *lem* of $\vec{\mathfrak{B}}_m^p$ is the same as the *lem* (1 \exists) of $\vec{\mathfrak{A}}_m^{\exists q}$ except that $\vec{\mathfrak{D}}_{q,m}^{A,A}$ is replaced by a copy of $\vec{\mathfrak{D}}_{q,m}^{B,B}$.

Note that every structure under consideration is just an ordered coloured tree and there is a path from its root to any node. We can read off the string of colours of edges along this path without a skip. Define $lab(x)$ to be a string in $\{\text{red}, \text{blue}\}^*$ associated with the path. $lab(r) := \epsilon$, where ϵ is the empty string. We let “red” be later in the order than “blue” in the lexicographic order. Let x, y be two nodes. We use x^f (y^f resp.) to denote the father of x (y resp.) in this section.

In the above, we have defined the structures that will be used in the games. And the sentence that we use to define the separating property is the same as Definition 3.15. That is, we actually use the same separating property to achieve the goal.

In the last section, we use point-expansions to realize strategy compositions. In some special cases, such compositions can be simplified: some of the substructures collapse to “points”, i.e. the details of the substructures are omitted, and we use “colours” to distinguish different substructures, which are now *regarded* temporarily as elements. We call such a method a kind of “structural abstraction”, which is used to omit unrelated details of structures and simplify game arguments, and the games played on the simplified structures are images of the original games.

4.2. The duplicator’s winning strategy. Note that, even in the unordered case, when the spoiler picks x , the duplicator’s strategy in the games is always picking y s.t. $lab(x) = lab(y)$ and picking a child of an element which is picked in some previous round, say the i -th round, when the spoiler picks a child of the other element which is picked in the i -th round. This lays the crucial basis for the previous inductive proof of Theorem 3.23 to extend to classes of linearly ordered finite structures, because now we can use something similar to the following well-known result [4]:

Let $k \geq 1$, and let L_1, L_2 be linear orders of length no less than 2^k , then

$$L_1 \equiv_k L_2. \quad (*)$$

Here, $L_1 \equiv_k L_2$ means that, for any $\varphi \in \text{FO}$ with $qr(\varphi) \leq k$, $L_1 \models \varphi$ iff $L_2 \models \varphi$.

Let $\tilde{\varphi}_p$ be given by the definition 3.15. The proof of lemma 3.17 also shows that $\vec{\mathfrak{A}}_m^p \models \tilde{\varphi}_p$ and $\vec{\mathfrak{B}}_m^p \not\models \tilde{\varphi}_p$, since “ \leq ” does not appear in $\tilde{\varphi}_p$.

To prove a version of Theorem 3.23 over ordered structures, the main idea is almost the same. Here, we just need to take care of the linear order. As we have explained before, the players will always pick a pair of elements that have the same label, and if the spoiler picks more than one child of a node in one structure, so does the duplicator in the other structure. Hence, we can use structural abstraction to conceal the details of the substructures like $\vec{\mathfrak{D}}_{q,m}^{A,B}$, and regard the problem to be a game over two linear orders with three “colours”, which represent three “colours” $\vec{\mathfrak{D}}_{q,m}^{A,B}$, $\vec{\mathfrak{D}}_{q,m}^{B,A}$, and $\vec{\mathfrak{D}}_{q,m}^{A,A}$ respectively. See Figure 5 for example. Here, $\vec{\mathfrak{D}}_{q,m}^{A,B}$ is identified with light red; $\vec{\mathfrak{D}}_{q,m}^{B,A}$ is identified with light blue; $\vec{\mathfrak{D}}_{q,m}^{A,A}$ is identified with green.

As explained before, we may safely assume that $p[1] = \exists$.

Assume that $p = \exists q$. We can pair the children of $r^{\vec{\mathfrak{A}}_m^p}$ as:

$$(\vec{\mathfrak{D}}_{q,m}^{A,B}, \vec{\mathfrak{D}}_{q,m}^{B,A}), \dots, (\vec{\mathfrak{D}}_{q,m}^{A,A}, \times), \dots, (\vec{\mathfrak{D}}_{q,m}^{A,B}, \vec{\mathfrak{D}}_{q,m}^{B,A}). \quad (4.1)$$

We can regard two $(\tau^{+\text{ORD}} \cup \{\epsilon\} \setminus \{r\})$ -structures in the brackets as a whole. More precisely, an s -2-tuple is a pair of structures from $\{\vec{\mathfrak{D}}_{s,m}^{A,B}, \vec{\mathfrak{D}}_{s,m}^{B,A}, \vec{\mathfrak{D}}_{s,m}^{A,A}, \vec{\mathfrak{D}}_{s,m}^{B,B}\}$ ($s \leq p$), which is regarded as a single “super-element”. From now on, we omit “ s ” in “ s -2-tuple” when it

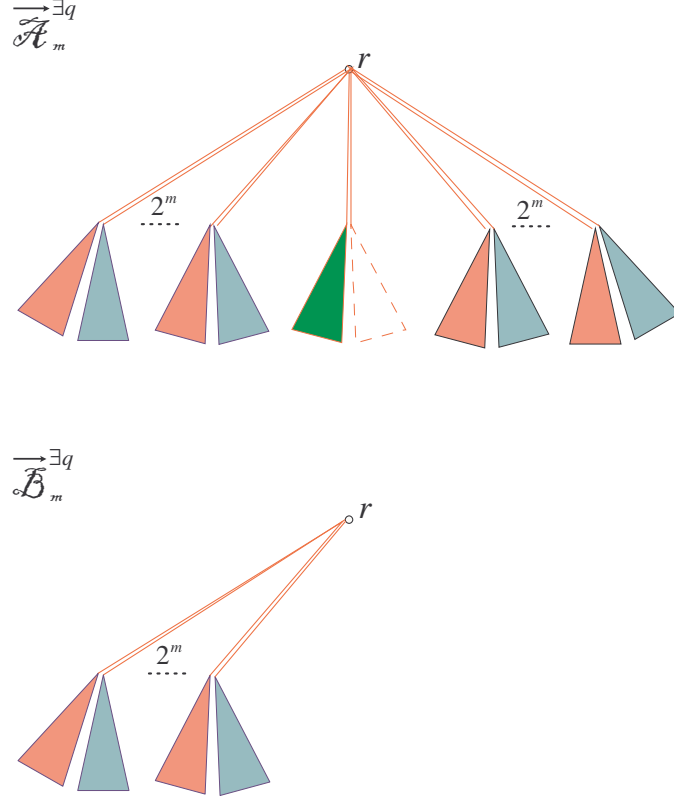


Figure 5: The structures where subgraphs are identified with “colours”.

will not cause confusion from the context. In (4.1), “ \times ” in the 2-tuple $(\vec{\mathcal{D}}_{q,m}^{A,A}, \times)$ represents an empty, or imaginary, structure, which is just used to make up a 2-tuple. See Figure 6. But, since we are regarding a 2-tuple as one single object now, we use one dashed line to represent two arrows. Note that, by definition, the elements of $\vec{\mathcal{D}}_{s,m}^{A,B}$ are earlier than those of $\vec{\mathcal{D}}_{s,m}^{B,A}$, in a 2-tuple $(\vec{\mathcal{D}}_{s,m}^{A,B}, \vec{\mathcal{D}}_{s,m}^{B,A})$. These 2-tuples have a natural linear order \leq_2 inherited from $\leq_{\vec{\mathcal{A}}_m^p}$: for any two 2-tuples $(\mathcal{X}_1, \mathcal{Y}_1)$ and $(\mathcal{X}_2, \mathcal{Y}_2)$, $(\mathcal{X}_1, \mathcal{Y}_1) \leq_2 (\mathcal{X}_2, \mathcal{Y}_2)$ iff $r^{\mathcal{Y}_1} \leq_{\vec{\mathcal{A}}_m^p} r^{\mathcal{X}_2}$ where $r^{\mathcal{Y}_1}$ and $r^{\mathcal{X}_2}$ are the junction points of \mathcal{Y}_1 and \mathcal{X}_2 respectively. Similarly, we can pair the children of $r^{\vec{\mathcal{B}}_m^p}$. Note that there are $(2^{m+1} + 1)$ 2-tuples in $\vec{\mathcal{A}}_m^p$, which form a linear order L_A , and 2^m 2-tuples in $\vec{\mathcal{B}}_m^p$, which form another linear order L_B . See Figure 7: a yellow node represents a 2-tuple, which corresponds to a pair of light red and light blue structures in Figure 5; the green node represents the 2-tuple that is green in Figure 5.

Let S be a Γ -labelled forest and $rk(S) = m$. In the game $G_S(\vec{\mathcal{A}}_m^p, \vec{\mathcal{B}}_m^p)$, when the spoiler picks an element x , which is not the root of the structure, in some round, x determines a q -2-tuple, where $p = \exists q$, s.t. x is in the universe of this q -2-tuple. And we may also say that this 2-tuple is picked in this round. Hence, $G_S(\vec{\mathcal{A}}_m^p, \vec{\mathcal{B}}_m^p)$ induces a new game in which the

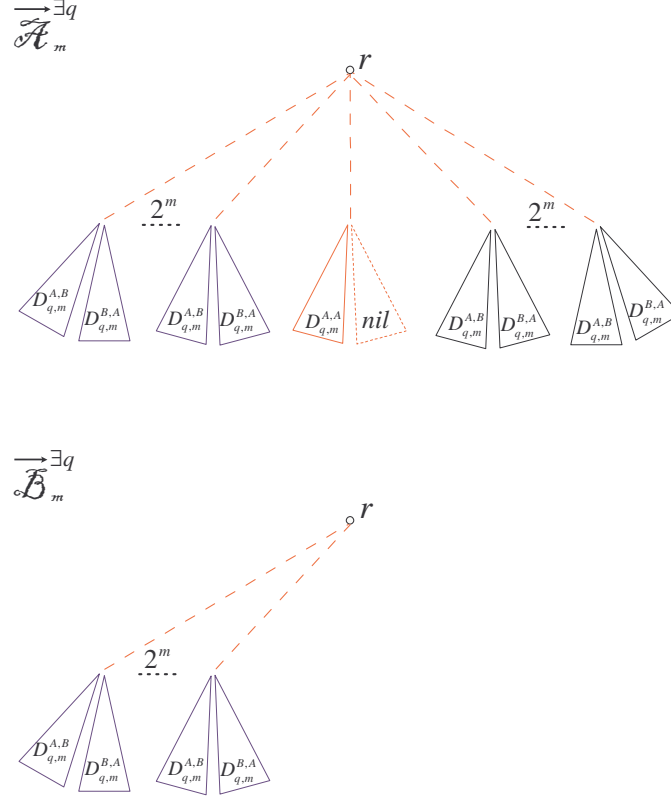


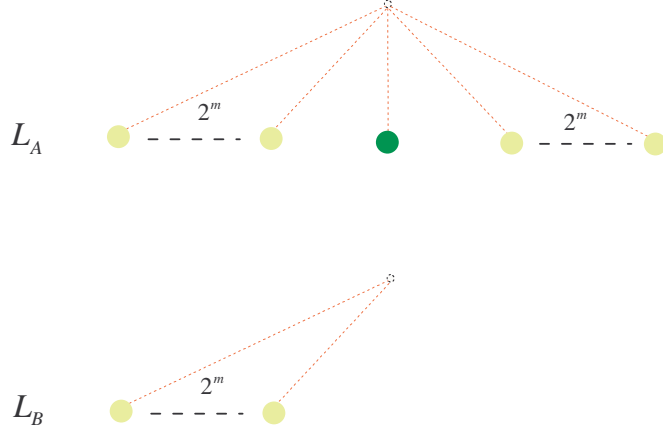
Figure 6: The structures and 2-tuples.

players pick 2-tuples instead of elements in the universe. Moreover, these 2-tuples form a pair of linear orders L_A and L_B , as shown in Figure 7. Call this new game “coloured linear order game” $G_S(L_A, L_B)$.

From now on, we use a natural number to denote a 2-tuple, in order to omit the details of 2-tuples that are not related to our concern but at the same time retain the order relation between 2-tuples. Therefore, we can subtract one 2-tuple from another 2-tuple in this context. Note that in this viewpoint linear orders can also be regarded as intervals. Moreover, every set of elements that have the same labels form an interval.

The duplicator’s strategy in $G_S(L_A, L_B)$ is as follows. Assume that $\mathcal{B}_i, \mathcal{B}_j$ are already picked. Let \mathcal{A}_i (\mathcal{A}_j resp.) be the element picked in L_A in the same round as \mathcal{B}_i (\mathcal{B}_j resp.) was picked. Recall that these 2-tuples can be compared by the induced order.

- If in the current round the spoiler picks a 2-tuple \mathcal{B} in the interval $[\mathcal{B}_i, \mathcal{B}_j]$ and $\mathcal{B} - \mathcal{B}_i \leq \mathcal{B}_j - \mathcal{B}$ then the duplicator picks \mathcal{A} s.t. $\mathcal{A} - \mathcal{A}_i = \mathcal{B} - \mathcal{B}_i$; otherwise she picks \mathcal{A} s.t. $\mathcal{A}_j - \mathcal{A} = \mathcal{B}_j - \mathcal{B}$.
- If in the current round the spoiler picks a 2-tuple \mathcal{A} in the interval $[\mathcal{A}_i, \mathcal{A}_j]$, the duplicator’s strategy is as follows:

Figure 7: L_A and L_B .

- (1) if $\mathcal{A} - \mathcal{A}_i \leq \frac{1}{2}(\mathcal{B}_j - \mathcal{B}_i)$ then the duplicator picks \mathcal{B} in $[\mathcal{B}_i, \mathcal{B}_j]$ s.t. $\mathcal{B} - \mathcal{B}_i = \mathcal{A} - \mathcal{A}_i$; otherwise
- (2) if $\mathcal{A}_j - \mathcal{A} \leq \frac{1}{2}(\mathcal{B}_j - \mathcal{B}_i)$ then the duplicator picks \mathcal{B} in $[\mathcal{B}_i, \mathcal{B}_j]$ s.t. $\mathcal{B}_j - \mathcal{B} = \mathcal{A}_j - \mathcal{A}$; otherwise
- (3) the duplicator picks the middle element in $[\mathcal{B}_i, \mathcal{B}_j]$: if the special 2-tuple $(\vec{\mathcal{D}}_{q,m}^{A,A}, \times)$ (the green node in Figure 7) is in the region $[\mathcal{A}, \mathcal{A}_j]$ then $\mathcal{B} - \mathcal{B}_i = \lfloor \frac{1}{2}(\mathcal{B}_j - \mathcal{B}_i) \rfloor$, otherwise $\mathcal{B}_j - \mathcal{B} = \lfloor \frac{1}{2}(\mathcal{B}_j - \mathcal{B}_i) \rfloor$.

Assume that in the first k rounds the spoiler is restricted to pick in $\vec{\mathfrak{B}}_m^p$.

Because a 2-tuple may contain smaller 2-tuples when looking inside it, this strategy can be applied recursively until the duplicator finds an element to pick in the game $G_S(\vec{\mathfrak{A}}_m^p, \vec{\mathfrak{B}}_m^p)$: she always picks an element that has the same label as the element picked by the spoiler in the same round. Using this strategy together with the strategy that is used in the game between two unordered structures, the duplicator can ensure that $(\vec{\mathcal{D}}_{q,m}^{A,A}, \times)$ will never be picked in the first k rounds, in which the spoiler picks 2-tuples in L_B , and order itself will not cause a problem throughout the game $G_S(\vec{\mathfrak{A}}_m^p, \vec{\mathfrak{B}}_m^p)$. More precisely, the following two lemmas are true: (Recall that S is a Γ -labelled forest, and $m = rk(S)$)

Lemma 4.2. *Assume that $k \in [0, m]$, and in the first k rounds of the game $G_S(L_A, L_B)$, the spoiler picks 2-tuples in L_B .*

- (i) *At the end of the i -th round of $G_S(L_A, L_B)$, for each $i \leq k$, there is only one interval in L_A that is not isomorphic to the corresponding interval in L_B , both of which are no shorter than 2^{m-i} , and $(\vec{\mathcal{D}}_{q,m}^{A,A}, \times)$ always lies in this interval of L_A .*
- (ii) *In the game $G_S(L_A, L_B)$, for any $i, j \leq k$,*

$$\mathcal{A}_i \leq_2 \mathcal{A}_j \text{ iff } \mathcal{B}_i \leq_2 \mathcal{B}_j.$$

Proof. Before the game $G_S(L_A, L_B)$, there is only one interval in L_A and L_B . Namely L_A and L_B themselves. Let $\mathcal{A}_f, \mathcal{A}_l$ be the first and the last 2-tuples of L_A respectively. And \mathcal{B}_f and \mathcal{B}_l be the first and the last 2-tuples of L_B respectively.

- (1) In the first round, there are two possibilities:
- (a) The players pick the first 2-tuple or the last 2-tuple in the respective intervals and the intervals remain unchanged;
 - (b) The picked 2-tuples split L_A and L_B into two intervals. In this case, according to the duplicator's strategy, one interval in L_A and L_B has the same length, which is no larger than 2^{m-1} , and the other interval in L_A includes $(\vec{\mathcal{D}}_{q,m}^{A,A}, \times)$ because the length between the first 2-tuple (the last 2-tuple resp.) and $(\vec{\mathcal{D}}_{q,m}^{A,A}, \times)$ is the length of L_B plus one which means that this length is larger than any interval in L_B . It also means that the interval which includes $(\vec{\mathcal{D}}_{q,m}^{A,A}, \times)$ is longer than the other one in L_A . Hence, it is larger than 2^{m-1} , like the corresponding interval in L_B .
- Therefore, (i) holds when the game is at the end of the first round. Assume that (i) holds when the game is at the end of the s -th round, where $1 \leq s \leq k$.

When the game is in the $(s+1)$ -th round, there are only two cases:

- (a) The spoiler picks a 2-tuple that is in the interval, say I_b in L_B . Let the corresponding interval in L_A be I_a , which includes the 2-tuple $(\vec{\mathcal{D}}_{q,m}^{A,A}, \times)$. By assumption, both of them are longer than 2^{m-s} . In this round, the unique pair of intervals I_a and I_b are split into two pairs: one is isomorphic; the other one is not and includes $(\vec{\mathcal{D}}_{q,m}^{A,A}, \times)$ because $(\vec{\mathcal{D}}_{q,m}^{A,A}, \times)$ splits I_a into two pieces and both of them are longer than I_b . Note that the pair of isomorphic intervals are no longer than 2^{m-s-1} . Hence, the other pair of intervals are no shorter than 2^{m-s-1} , by the duplicator's strategy.
 - (b) The spoiler picks a 2-tuple in other intervals. Note that all the other pairs of intervals have the same length. According to the duplicator's strategy, splitting such a pair of intervals only produces pairs of intervals of the same length. And the pair of intervals which are not isomorphic is unchanged. By the inductive assumption, (i) still holds.
- (2) We prove an equivalent conclusion, i.e. (ii) holds if we add two rounds before the game in which $\mathcal{A}_f, \mathcal{A}_l, \mathcal{B}_f$ and \mathcal{B}_l are picked. Clearly, $\mathcal{A}_f \leq_2 \mathcal{A}_l$ iff $\mathcal{B}_f \leq_2 \mathcal{B}_l$. In other words, it holds when these two elements are picked.

Suppose (ii) holds when the game is at the end of the s -th round, where $s \geq 0$.

Now assume that the game is in the $(s+1)$ -th round of the game. If the spoiler picks a 2-tuple that was picked before, then it still holds. If the spoiler picks a 2-tuple \mathcal{B} that splits some interval in L_B , say $[\mathcal{B}_h, \mathcal{B}_t]$, the duplicator also picks a 2-tuple \mathcal{A} that splits the corresponding interval $[\mathcal{A}_h, \mathcal{A}_t]$ by her strategy. If $[\mathcal{B}_h, \mathcal{B}_t]$ and $[\mathcal{A}_h, \mathcal{A}_t]$ are isomorphic, then obviously (ii) holds. If it is not the case, due to (i), we know that $[\mathcal{A}_h, \mathcal{A}_t]$ is sufficiently long such that it allows such splitting. Then for any 2-tuple \mathcal{B}_x that is picked before, $\mathcal{B}_x \leq_2 \mathcal{B}_h$ if $\mathcal{B}_x \leq_2 \mathcal{B}$, which, by assumption, implies that $\mathcal{A}_x \leq_2 \mathcal{A}_h$. Hence, $\mathcal{A}_x \leq_2 \mathcal{A}$. Likewise, $\mathcal{A}_x \geq_2 \mathcal{A}$ if $\mathcal{B}_x \geq_2 \mathcal{B}$. Therefore, for any i, j , $\mathcal{A}_i \leq_2 \mathcal{A}_j$ iff $\mathcal{B}_i \leq_2 \mathcal{B}_j$. The "only if" part can be proven similarly. \square

Part (i) of Lemma 4.2 tells us that the spoiler cannot use the linear order to force the duplicator to violate her winning strategy in $G_S(\mathfrak{A}_m^p, \mathfrak{B}_m^p)$, no matter how he picks in L_B . That is, if the spoiler picks a yellow node in L_B (see Figure 7), the duplicator can also pick a yellow node. A similar thing can be proved when $p[1] = \forall$. Because the collection of 2-tuples

whose roots are children (with the same label) of an element (node) in $\vec{\mathfrak{A}}_m^p$ form an interval, we can generalize (ii) of Lemma 4.2 such that it applies to any pair of intervals that are split in the same round. Recall that a 2-tuple is composed of two trees. A 2-tuple is at depth i if the roots of trees of this 2-tuple are at depth i . Hence, we can call these intervals *intervals at depth i* if the elements (2-tuples) of these intervals are at depth i of the structures (trees).

Lemma 4.3. *Let S be a Γ -labelled forest and $m = rk(S)$. Let p be a prefix. In the game $G_S(\vec{\mathfrak{A}}_m^p, \vec{\mathfrak{B}}_m^p)$ where $p \notin \mathcal{W}(S)$, the duplicator can play in such a way that:*

- (i) *she follows her winning strategy in $G_S(\mathfrak{A}_m^p, \mathfrak{B}_m^p)$ that has been described in Lemma 3.22;*
- (ii) *for any $a_i, a_j \in |\vec{\mathfrak{A}}_m^p|$, $b_i, b_j \in |\vec{\mathfrak{B}}_m^p|$, which are picked by the players in i -th and j -th rounds ($i, j \leq m$),*

$$a_i \leq \vec{\mathfrak{A}}_m^p a_j \text{ iff } b_i \leq \vec{\mathfrak{B}}_m^p b_j.$$

Proof. We assume that, before the first round of the game, both the first child and last child of any inner node are already picked. It means that we are trying to prove an equivalent result.

In each round, when the spoiler picks an element a in a structure, say $\vec{\mathfrak{A}}_m^p$, there is a path P from the root of the tree $\vec{\mathfrak{A}}_m^p$ to the node a . For any node v in the path P there is a 2-tuple $(\mathcal{X}, \mathcal{Y})$ such that $(\mathcal{X}, \mathcal{Y})$ is the 2-tuple that include a and v is the root of either \mathcal{X} or \mathcal{Y} . Let a be at the depth r of the tree $\vec{\mathfrak{A}}_m^p$ and let the path P be (a_0, \dots, a_r) where a_0 is the root of $\vec{\mathfrak{A}}_m^p$ and $a_r = a$. Assume that the intervals of 2-tuples that are split by the path is (L_1^A, \dots, L_r^A) . That is, L_i^A is an interval at the depth i of the tree $\vec{\mathfrak{A}}_m^p$, which includes a_i . Let (L_1^B, \dots, L_r^B) be the collection of intervals of 2-tuples where L_i^B is at depth i and is formed in the same rounds as L_i^A . The duplicator first picks all the nodes in the path (a_0, \dots, a_r) (Looking at it in another way, the spoiler implicitly picks all the nodes in the path). Then she uses her strategy in coloured linear order games recursively as follows: she first plays the one round coloured linear order game at the depth 1, i.e. over the pair of intervals (L_1^A, L_1^B) , picking a 2-tuple at the depth 1 of $\vec{\mathfrak{B}}_m^p$, which splits the interval L_1^B as a consequence; then she picks the root of a tree in the 2-tuple (recall that a 2-tuple is composed of two trees) which respects her winning strategy in $G_S(\mathfrak{A}_m^p, \mathfrak{B}_m^p)$ (recall Definition 3.13 for the definition of \mathfrak{A}_m^p and \mathfrak{B}_m^p): whenever possible she tries to choose the tree, whose root is b_1 , that is isomorphic to the one implicitly picked by the spoiler in L_1^A whose root is a_1 , and pick b_1 . Then she goes on to play the game over the pair of intervals (L_2^A, L_2^B) . For $1 \leq i \leq r-1$, once she picked a 2-tuple at depth i , she will pick the root, say c , of a tree from the 2-tuple, afterwards she picks a 2-tuple at depth $i+1$, whose trees are the children of c . At last, she picks a 2-tuple in the interval L_r^B and picks the root of a tree from this 2-tuple that respects her winning strategy in $G_S(\mathfrak{A}_m^p, \mathfrak{B}_m^p)$.

(i) will be violated only when the lengths of the pair of non-isomorphic intervals are not long enough such that the spoiler can force the duplicator to pick a different type of 2-tuple when he picks repeatedly in the shorter interval. Note that any pair of isomorphic intervals have the same type of 2-tuples. Hence, to prove (i), we need only show that (i') at the end of the i -th round, if a pair of intervals is not isomorphic, then the length of them are no less than 2^{m-i} . However, (i') is obvious because the strategy of the duplicator ensures that after the i -th round ($i \geq 1$), the lengths of any pair of non-isomorphic intervals reduce at most 2^{m-i} , while they are at least 2^{m-i+1} before the i -th round, which can be proved inductively

as in Lemma 4.2. In other words, this strategy is able to incorporate her winning strategy in $G_S(\mathfrak{A}_m^p, \mathfrak{B}_m^p)$ - it can guide her to win the game if order is not taken into account. It remains to show that this is a strategy for the duplicator to avoid the order problem, i.e. to show (ii).

For the sake of convenience, we call the stage before the players play the game as the 0-th round. In the first round, the order will not be a problem since other than those nodes at the ends of intervals, there is no other node that can violate (ii).

Assume that (ii) holds when it is at the end of the s -th round.

Now suppose that the game is in the $(s + 1)$ -th round. Let a_i and a_j be two elements picked in the i -th and j -th rounds ($i, j \leq s$) in $\vec{\mathfrak{A}}_m^p$. We may further assume that either i or j equals $s + 1$. When r, a_i, a_j are in a path, then $a_i < a_j$ implies $b_i < b_j$ because the duplicator's strategy ensures that r, b_i, b_j are also in a path.

Now assume that r, a_i, a_j are not in a path.

If a_i and a_j have the same father and the same label, which means they are in the same interval, then we can apply the same argument of Lemma 4.2, simply by regarding an element as a 2-tuple. If a_i and a_j have the same father but have different labels, then by definition their order is determined by their labels. So are b_i and b_j . Note that the label of a_i and b_i (a_j and b_j resp.) are the same, by the duplicator's strategy. Therefore, (ii) holds.

Assume that a_i and a_j have different fathers. Note that a_i, a_j always share at least one ancestor, i.e. the root r . Let c be such an shared ancestor, and for all other shared ancestors, c is later in the order. Let a'_i be the ancestor of a_i (or a_i itself) and the child of c . Let a'_j be the ancestor of a_j (or a_j itself) and the child of c . Let b'_i, b'_j be defined in a similar way in $\vec{\mathfrak{B}}_m^p$. Then by definition the order between a_i and a_j is determined by the order between a'_i and a'_j . And the duplicator can ensure that $a'_i < a'_j$ iff $b'_i < b'_j$, according to the same argument as Lemma 4.2. Therefore, (ii) holds. \square

Lemma 4.3 tell us that linear order does not cause a problem to the duplicator, and together with the arguments in Theorem 3.1 the following holds.

Theorem 4.4. *Let S_1 and S_2 be two finite Γ -labelled forests. Over the class of all finite τ^{+ORD} -structures,*

$$\text{if } \mathcal{W}(S_1) \not\subseteq \mathcal{W}(S_2), \text{ then } \text{FO}\{S_1\} \not\subseteq \text{FO}\{S_2\}.$$

Using similar arguments as in the last section, in particular the same reduction as in Lemma 3.26, we can prove the following Theorem.

Theorem 4.5. *Let S_1 and S_2 be two finite Γ -labelled forests. Over the class of all ordered finite digraphs,*

$$\text{if } \mathcal{W}(S_1) \not\subseteq \mathcal{W}(S_2), \text{ then } \text{FO}\{S_1\} \not\subseteq \text{FO}\{S_2\}.$$

Here we call $\tau \cup \{\leq\}$ -structures (linearly) ordered digraphs.

The following corollary is a special case of Theorem 4.5, where S_1 and S_2 are degenerate trees (or directed paths).

Corollary 4.6. *Let $p, q \in \Gamma^*$. Over the class of all ordered finite digraphs,*

$$\text{if } p \not\leq q, \text{ then } \text{FO}\{p\} \not\subseteq \text{FO}\{q\}.$$

Note that it is different from Grädel and McColm's conjecture [2].

A natural question is whether something similar to Theorem 3.1 holds, but over finite digraphs with built-in BIT. Here, BIT is the binary relation for the bit operator: $\text{BIT}(x, y) = 1$ if the y -th bit of the binary representation of x is 1. The operator BIT seems very powerful. It is known that first-order logic equipped with BIT can define arbitrary algorithmic operators, including $\leq, \times, +, \text{Exp}$, and Squares (Schweikardt, [7]). Supprisingly, Schweikardt and Schwentick [8] showed that BIT is similar to linear orders in terms of expressive power in first-order logic. Based on their constructions, it is not difficult to show that the quantifier structure hierarchy is strict in FO, even in the presence of BIT.

5. A REFINED QUANTIFIER STRUCTURE HIERARCHY

5.1. The structures and separating property. It is possible that two Γ -labelled forests cannot embed to each other but the set of words that can be read off them are the same. It is natural to conjecture that they represent different logical resources. However, the hierarchy we defined in the last section cannot tell us about it. In the following, we are going to show a refined strict hierarchy, which confirms this intuition:

Theorem 5.1. *Let S_1 and S_2 be two Γ -labelled forests. Over the class of all digraphs,*

$$\text{if } S_1 \not\leq_e S_2, \text{ then } \text{FO}\{S_1\} \not\subseteq \text{FO}\{S_2\}.$$

As in the last section, we let $\tau^+ := \langle R, B, r, U \rangle$, and let $\tau_0^+ := \tau^+ \setminus \{U\}$.

Definition 5.2. A Γ -labelled tree \mathcal{T} is an irreducible tree if for any inner node a the following holds:

Let b_1, \dots, b_k be the children of a and $\mathcal{T}_1, \dots, \mathcal{T}_k$ be the maximal subtrees of \mathcal{T} that are rooted at b_1, \dots, b_k respectively, then \mathcal{T}_i cannot embed in \mathcal{T}_j for any $i, j \in [1, k]$ where $i \neq j$.

Definition 5.3. Let \mathcal{T} be a Γ -labelled irreducible k -ary tree.

- $\tilde{\mathfrak{A}}_m^T$ and $\tilde{\mathfrak{B}}_m^T$ are coloured trees.² The constant r is interpreted as the root of the respective trees. As in Definition 3.13, we say an element a is black if $a \in U$.

- (1) $\tilde{\mathfrak{A}}_{m,l}^\exists$ is a depth 1 tree that has $lm + 1$ leaves. One of the leaves is black and all the other leaves are not black.

- (2) $\tilde{\mathfrak{B}}_{m,l}^\exists$ is a depth 1 tree that has lm leaves. None of them is black.

All edges in $\tilde{\mathfrak{A}}_{m,l}^\exists$ and $\tilde{\mathfrak{B}}_{m,l}^\exists$ are red. Recall that we say an edge (a, b) is red if, and only if, $(a, b) \in R$.

$$\mathcal{T}_{m,l}^{\mathfrak{A},\exists} := \tilde{\mathfrak{A}}_{m,l}^\exists | \tau_0^+; \quad \mathcal{T}_{m,l}^{\mathfrak{B},\exists} := \tilde{\mathfrak{B}}_{m,l}^\exists | \tau_0^+.$$

- (1) $\tilde{\mathfrak{A}}_{m,l}^\forall$ is a depth 1 tree that has lm leaves. All of the leaves are black.

- (2) $\tilde{\mathfrak{B}}_{m,l}^\forall$ is a depth 1 tree that has $lm + 1$ leaves. One of them is not black and all the other leaves are black.

All edges in $\tilde{\mathfrak{A}}_{m,l}^\forall$ and $\tilde{\mathfrak{B}}_{m,l}^\forall$ are red.

$$\mathcal{T}_{m,l}^{\mathfrak{A},\forall} := \tilde{\mathfrak{A}}_{m,l}^\forall | \tau_0^+; \quad \mathcal{T}_{m,l}^{\mathfrak{B},\forall} := \tilde{\mathfrak{B}}_{m,l}^\forall | \tau_0^+.$$

- If \mathcal{T} contains a single \mathcal{E} node,

$$\tilde{\mathfrak{A}}_m^T \text{ is } \tilde{\mathfrak{A}}_{m,1}^\exists; \quad \tilde{\mathfrak{B}}_m^T \text{ is } \tilde{\mathfrak{B}}_{m,1}^\exists.$$

²Here, the “ T ” in $\tilde{\mathfrak{A}}_m^T$ and $\tilde{\mathfrak{B}}_m^T$ refers to the tree \mathcal{T} . Also cf. Figure 8.

- If \mathcal{T} contains a single \mathcal{A} node,

$$\tilde{\mathfrak{A}}_m^T \text{ is } \tilde{\mathfrak{A}}_{m,1}^\vee; \tilde{\mathfrak{B}}_m^T \text{ is } \tilde{\mathfrak{B}}_{m,1}^\vee.$$

- When $|rk(\mathcal{T})| > 1$:

Assume that the root r of \mathcal{T} has k children, and the maximal subtrees that are rooted at these children are $\mathcal{T}_1, \dots, \mathcal{T}_k$ respectively. Recall that \mathcal{T} is an irreducible tree, which implies \mathcal{T}_i and \mathcal{T}_j are not isomorphic if $i \neq j$.

For any Γ -labelled tree \mathcal{T}' , let $\tilde{\mathfrak{A}}_m^{-T'}$ be the same as $\tilde{\mathfrak{A}}_m^{T'}$ except that the colours of all the edges are exchanged, i.e. red is exchanged with blue. Let P_i , which is a member of the set $\{(A, A), (A, B), (B, A), (B, B), A, B\}$, and $\mathfrak{D}_{T_1, \dots, T_k, m}^{P_1, \dots, P_k}$, which is a $(\tau^+ \cup \{\epsilon\} \setminus \{r\})$ -structure where ϵ is a hook constant, be built by the following process:

step 1: Let $M := \emptyset$;

step 2: For $i = 1$ to k do the following:

- if $P_i = (A, A)$ then $M := M \cup \{\tilde{\mathfrak{A}}_m^{T_i}, \tilde{\mathfrak{A}}_m^{-T_i}\}$;
- if $P_i = (A, B)$ then $M := M \cup \{\tilde{\mathfrak{A}}_m^{T_i}, \tilde{\mathfrak{B}}_m^{-T_i}\}$;
- if $P_i = (B, A)$ then $M := M \cup \{\tilde{\mathfrak{B}}_m^{T_i}, \tilde{\mathfrak{A}}_m^{-T_i}\}$;
- if $P_i = (B, B)$ then $M := M \cup \{\tilde{\mathfrak{B}}_m^{T_i}, \tilde{\mathfrak{B}}_m^{-T_i}\}$;
- if $P_i = A$ then $M := M \cup \{\tilde{\mathfrak{A}}_m^{T_i}\}$;
- if $P_i = B$ then $M := M \cup \{\tilde{\mathfrak{B}}_m^{T_i}\}$;

step 3: join all the trees in M at their roots. Call their shared root a *junction point*, which interprets the hook constant ϵ .

We use O_i to denote the label of the root of the tree \mathcal{T}_i . Let $H := \{i \in [1, k] \mid \text{the root of } \mathcal{T}_i \text{ has the same label as that of the root of } \mathcal{T}\}$.

In the following, we define some substructure for the constructions:

- For any $j \in H$, let $\mathfrak{C}_{\exists, j, m}^{A, B, T} := \mathfrak{D}_{T_1, \dots, T_k, m}^{P_1, \dots, P_k}$ where $P_j = (A, B)$ and [if $i \neq j$ then ($P_i = (A, A)$ if $O_i = \exists$) and ($P_i = A$ if $O_i = \forall$)];
- For any $j \in H$, let $\mathfrak{C}_{\exists, j, m}^{B, A, T} := \mathfrak{D}_{T_1, \dots, T_k, m}^{P_1, \dots, P_k}$ where $P_j = (B, A)$ and [if $i \neq j$ then ($P_i = (A, A)$ if $O_i = \exists$) and ($P_i = A$ if $O_i = \forall$)];
- For any $j \in [1, k]$ and $j \notin H$, let $\mathfrak{C}_{\exists, j, m}^{B, B, T} := \mathfrak{D}_{T_1, \dots, T_k, m}^{P_1, \dots, P_k}$ where $P_j = B$ and [if $i \neq j$ then ($P_i = (A, A)$ if $O_i = \exists$) and ($P_i = A$ if $O_i = \forall$)];
- We use $\mathfrak{C}_{\exists, 0, m}^{A, A, T}$ to denote the structure $\mathfrak{D}_{T_1, \dots, T_k, m}^{P_1, \dots, P_k}$ where $P_i = (A, A)$ if $O_i = \exists$; $P_i = A$ if $O_i = \forall$.

Dually, we can define $\mathfrak{C}_{\forall, j, m}^{B, A, T}, \mathfrak{C}_{\forall, j, m}^{A, B, T}, \mathfrak{C}_{\forall, j, m}^{A, A, T}, \mathfrak{C}_{\forall, 0, m}^{B, B, T}$ as follows:

- For any $j \in H$, let $\mathfrak{C}_{\forall, j, m}^{A, B, T} := \mathfrak{D}_{T_1, \dots, T_k, m}^{P_1, \dots, P_k}$ where $P_j = (A, B)$ and [if $i \neq j$ then ($P_i = (B, B)$ if $O_i = \forall$) and ($P_i = B$ if $O_i = \exists$)];
- For any $j \in H$, let $\mathfrak{C}_{\forall, j, m}^{B, A, T} := \mathfrak{D}_{T_1, \dots, T_k, m}^{P_1, \dots, P_k}$ where $P_j = (B, A)$ and [if $i \neq j$ then ($P_i = (B, B)$ if $O_i = \forall$) and ($P_i = B$ if $O_i = \exists$)];
- For $j \notin H$, let $\mathfrak{C}_{\forall, j, m}^{A, A, T} := \mathfrak{D}_{T_1, \dots, T_k, m}^{P_1, \dots, P_k}$ where $P_j = A$ and [if $i \neq j$ then ($P_i = (B, B)$ if $O_i = \forall$) and ($P_i = B$ if $O_i = \exists$)];
- We use $\mathfrak{C}_{\forall, 0, m}^{B, B, T}$ to denote the structure $\mathfrak{D}_{T_1, \dots, T_k, m}^{P_1, \dots, P_k}$ where $P_i = (B, B)$ if $O_i = \forall$; $P_i = B$ if $O_i = \exists$.

As usual, we define the main structures based on point-expansions over some sets of structures. Now we define such sets.

- $\mathcal{K}_0^{\exists, T} := \{\mathfrak{C}_{\exists, i, m}^{A, B, T} \mid i \in H\} \cup \{\mathfrak{C}_{\exists, i, m}^{B, A, T} \mid i \in H\} \cup \{\mathfrak{C}_{\exists, i, m}^{B, B, T} \mid i \in [1, k] \text{ and } i \notin H\}$;

- $\mathcal{K}_A^{\exists,T} := \{\mathfrak{c}_{\exists,0,m}^{A,A,T}\} \cup \{\mathcal{I}\} \cup \mathcal{K}_0^{\exists,T}$;
- $\mathcal{K}_B^{\exists,T} := \{\mathcal{I}\} \cup \mathcal{K}_0^{\exists,T}$.
- $\mathcal{K}_0^{\forall,T} := \{\mathfrak{c}_{\forall,i,m}^{A,B,T} \mid i \in H\} \cup \{\mathfrak{c}_{\forall,i,m}^{B,A,T} \mid i \in H\} \cup \{\mathfrak{c}_{\forall,i,m}^{A,A,T} \mid i \in [1,k] \text{ and } i \notin H\}$;
- $\mathcal{K}_A^{\forall,T} := \{\mathcal{I}\} \cup \mathcal{K}_0^{\forall,T}$;
- $\mathcal{K}_B^{\forall,T} := \{\mathfrak{c}_{\forall,0,m}^{B,B,T}\} \cup \{\mathcal{I}\} \cup \mathcal{K}_0^{\forall,T}$.

Now we define $\tilde{\mathfrak{A}}_m^T$ and $\tilde{\mathfrak{B}}_m^T$ as follows:

- If the root of \mathcal{T} is an \mathcal{E} node, which is connected to the roots of k trees $\mathcal{T}_1, \dots, \mathcal{T}_k$, and assume that the number of \exists in the tuple (O_1, \dots, O_k) is j , then $\tilde{\mathfrak{A}}_m^T$ and $\tilde{\mathfrak{B}}_m^T$ are defined as follows:
 - * $\tilde{\mathfrak{A}}_m^T$ is a point-expansion of $\mathcal{T}_{m,j+k}^{\mathfrak{A},\exists}$ over $\mathcal{K}_A^{\exists,T}$: The root of $\mathcal{T}_{m,j+k}^{\mathfrak{A},\exists}$ is expanded by \mathcal{I} . It is also called the “root” of $\tilde{\mathfrak{A}}_m^T$ that interprets r . One leaf is expanded by a copy of $\mathfrak{c}_{\exists,0,m}^{A,A,T}$. For each element of $\mathcal{K}_0^{\exists,T}$ there are exactly m distinct leaves which are expanded by it.
 - * $\tilde{\mathfrak{B}}_m^T$ is a point-expansion of $\mathcal{T}_{m,j+k}^{\mathfrak{B},\exists}$ over $\mathcal{K}_B^{\exists,T}$, similar to $\tilde{\mathfrak{A}}_m^T$: The root of $\mathcal{T}_{m,j+k}^{\mathfrak{B},\exists}$ is expanded by \mathcal{I} . For each element of $\mathcal{K}_0^{\exists,T}$ there are exactly m distinct leaves which are expanded by it.
- If the root of \mathcal{T} is an \mathcal{A} node, which is connected to the roots of k trees $\mathcal{T}_1, \dots, \mathcal{T}_k$, and assume that the number of \forall in the tuple (O_1, \dots, O_k) is j , then $\tilde{\mathfrak{A}}_m^T$ and $\tilde{\mathfrak{B}}_m^T$ are defined as follows:
 - * $\tilde{\mathfrak{A}}_m^T$ is a point-expansion of $\mathcal{T}_{m,j+k}^{\mathfrak{A},\forall}$ over $\mathcal{K}_A^{\forall,T}$: The root of $\mathcal{T}_{m,j+k}^{\mathfrak{A},\forall}$ is expanded by \mathcal{I} . For each element of $\mathcal{K}_0^{\forall,T}$ there are exactly m distinct leaves which are expanded by it.
 - * $\tilde{\mathfrak{B}}_m^T$ is a point-expansion of $\mathcal{T}_{m,j+k}^{\mathfrak{B},\forall}$ over $\mathcal{K}_B^{\forall,T}$: The root of $\mathcal{T}_{m,j+k}^{\mathfrak{B},\forall}$ is expanded by \mathcal{I} . One leaf is expanded by a copy of $\mathfrak{c}_{\forall,0,m}^{B,B,T}$. For each element of $\mathcal{K}_0^{\forall,T}$ there are exactly m distinct leaves which are expanded by it.

Example 5.4. Let \mathcal{T} be a Γ -labelled irreducible binary tree. Assume that its root is an \mathcal{E} node and is connected to two subtrees \mathcal{T}_1 and \mathcal{T}_2 , where the root of \mathcal{T}_1 is an \mathcal{E} node and the root of \mathcal{T}_2 is an \mathcal{A} node. See Figure 8 for the illustration of the structures $\tilde{\mathfrak{A}}_m^T$ and $\tilde{\mathfrak{B}}_m^T$. A “*” at the root of a subtree \mathcal{T}' means that we have m disjoint isomorphic copies of this tree \mathcal{T}' and for each copy we add an edge between the root, r , of the whole structure and the root of this copy.

Now we define a sentence $\tilde{\phi}_{\mathcal{T}}$ over the signature τ^+ such that, for any m , $\tilde{\mathfrak{A}}_m^T \models \tilde{\phi}_{\mathcal{T}}$ but $\tilde{\mathfrak{B}}_m^T \not\models \tilde{\phi}_{\mathcal{T}}$.

Definition 5.5. Let \mathcal{T} be a Γ -labelled irreducible k -ary tree, which is connected to the roots of k trees $\mathcal{T}_1, \dots, \mathcal{T}_k$. Assume that $rk(\mathcal{T}) = d$. Recall that $H = \{i \in [k] \mid \text{the root of } \mathcal{T}_i \text{ has the same label as that of the root of } \mathcal{T}\}$. We define τ^+ -sentences $\tilde{\phi}_{\mathcal{T}}$ and $\tilde{\phi}_{-\mathcal{T}}$ based on the tree \mathcal{T} as follows:

- (1) If \mathcal{T} is empty, then

$$\tilde{\xi}_{\mathcal{T}}(x) = \tilde{\xi}_{-\mathcal{T}}(x) := U(x);$$

- (2) If the root of \mathcal{T} is an \mathcal{E} node, then

$$\tilde{\xi}_{\mathcal{T}}(y) := \exists x_d (Ryx_d \wedge \bigwedge_{i \in [k]} \tilde{\xi}_{\mathcal{T}_i}(x_d) \wedge \bigwedge_{j \in H} \tilde{\xi}_{-\mathcal{T}_j}(x_d));$$




Figure 8: The structures $\tilde{\mathfrak{A}}_m^T$ and $\tilde{\mathfrak{B}}_m^T$, where \mathcal{T} is an irreducible binary tree, which is connected to \mathcal{T}_1 and \mathcal{T}_2 . The root of \mathcal{T} is an \mathcal{E} node; The root of \mathcal{T}_1 is an \mathcal{E} node; and the root of \mathcal{T}_2 is an \mathcal{A} node. In the figure, A_{T_i} and B_{T_i} represent $\tilde{\mathfrak{A}}_m^{T_i}$ and $\tilde{\mathfrak{B}}_m^{T_i}$ respectively.

$$\tilde{\xi}_{-\mathcal{T}}(y) := \exists x_d (B y x_d \wedge \bigwedge_{i \in [k]} \tilde{\xi}_{-T_i}(x_d) \wedge \bigwedge_{j \in H} \tilde{\xi}_{T_j}(x_d));$$

(3) If the root of \mathcal{T} is an \mathcal{A} node, then

$$\tilde{\xi}_{\mathcal{T}}(y) := \forall x_d (R y x_d \rightarrow \bigvee_{i \in [k]} \tilde{\xi}_{T_i}(x_d) \vee \bigvee_{j \in H} \tilde{\xi}_{-T_j}(x_d));$$

$$\tilde{\xi}_{-\mathcal{T}}(y) := \forall x_d (B y x_d \rightarrow \bigvee_{i \in [k]} \tilde{\xi}_{-T_i}(x_d) \vee \bigvee_{j \in H} \tilde{\xi}_{T_j}(x_d));$$

Now, $\tilde{\phi}_{\mathcal{T}}$ and $\tilde{\phi}_{-\mathcal{T}}$ are defined as:

$$\tilde{\phi}_{\mathcal{T}} := \tilde{\xi}_{\mathcal{T}}(r); \tag{5.1}$$

$$\tilde{\phi}_{-\mathcal{T}} := \tilde{\xi}_{-\mathcal{T}}(r). \tag{5.2}$$

5.2. The duplicator's winning strategy. In the following, we show that a refined strict quantifier hierarchy exists by proving that the duplicator has a winning strategy in the games that we introduced before. Since the proof resembles that in Section 3, we sketch the main ideas and omit similar details.

Lemma 5.6. *Let S_1 and S_2 be two finite Γ -labelled forests. If $S_1 \not\leq_e S_2$, then there is an irreducible subtree \mathcal{T} in S_1 such that the duplicator has a winning strategy in $G_{S_2}(\tilde{\mathfrak{A}}_m^T, \tilde{\mathfrak{B}}_m^T)$ for any $m \geq rk(S_2)$.*

Proof. Observe that when $rk(S_1) = 1$, i.e. the rank of S_1 is 1, $S_1 \not\leq_e S_2$ implies $\mathscr{W}(S_1) \not\subseteq \mathscr{W}(S_2)$. According to Lemma 3.22, the duplicator has a winning strategy.

Assume that the lemma holds when $rk(S_1) \leq h$.

Now assume that $rk(S_1) = h + 1$.

Because $S_1 \not\leq_e S_2$, for some $k \in \mathbb{N}^+$ there exists a k -ary subtree \mathcal{T} in S_1 such that $\mathcal{T} \not\leq_e \mathcal{T}'$ for any subtree \mathcal{T}' in S_2 . Moreover, we assume that \mathcal{T} is the minimal subtree in S_1 that is not embeddable in S_2 , i.e. any subtree of \mathcal{T} , which is not \mathcal{T} itself, is embeddable in S_2 . Note that such a tree is an irreducible tree, whose rank is no larger than $h + 1$. Here we use $r(\mathcal{T})$ to denote the root of \mathcal{T} . And we assume that $r(\mathcal{T})$ is an \mathcal{E} node. We assume that there are $1 \leq k' \leq k$ subtrees, $\mathcal{T}_1, \dots, \mathcal{T}_{k'}$, whose roots are connected to $r(\mathcal{T})$. Let Mg be the set of \mathcal{E} nodes in S_2 such that for any $a \in \text{Mg}$ no other \mathcal{E} node appears in the path from the root of S_2 to a . By assumption \mathcal{T} is not embeddable in any subtree of S_2 . For any $a \in \text{Mg}$, let \mathcal{T}' be a tree rooted at a and \mathcal{F} be the forest obtained from \mathcal{T}' by removing the root of \mathcal{T}' . Hence, at least one of the trees $\mathcal{T}_1, \dots, \mathcal{T}_{k'}$ cannot be embedded in \mathcal{F} . Observe that $rk(\mathcal{T}_i) \leq h$ for any $i \in [k']$.

Note that, in any play of the game, the moving track of the token in S_2 is a directed path. If the first place where the token lies is an \mathcal{A} node, then the path is initiated with a block of universal quantifiers. However, in the rounds based on this first block of universal quantifiers the spoiler has to pick in $\tilde{\mathfrak{B}}_m^T$ and no matter how he picks the duplicator can mimic his picking in the isomorphic subtrees. By Lemma 3.20, if the spoiler can win the game by picking these elements, he can also win the game without picking these elements.

When the token is on a node of Mg , say a , we assume that the spoiler picks in $\mathfrak{C}_{\exists,0,m}^{A,A,T}$ (i.e. in a copy of $\tilde{\mathfrak{A}}_m^T$) because if he picks in other places the duplicator can mimic his picking in an isomorphic tree substructure, hence by Lemma 3.20 the spoiler can win the game without picking in these places if he has at least one winning strategy. In this round, the duplicator's strategy is as follows (i.e. picking in a copy of $\tilde{\mathfrak{B}}_m^T$):

- If \mathcal{T}_i is not embeddable in the forest \mathcal{F} and $r(\mathcal{T}_i)$ is an \mathcal{E} node, then
 - if the spoiler picks inside $\tilde{\mathfrak{A}}_m^{T_i}$, which is a part of $\mathfrak{C}_{\exists,0,m}^{A,A,T}$, then the duplicator mimics it in an isomorphic copy of $\tilde{\mathfrak{A}}_m^{T_i}$, which is a part of $\mathfrak{C}_{\exists,i,m}^{A,B,T}$.
 - if the spoiler picks inside $\tilde{\mathfrak{A}}_m^{-T_i}$, which is a part of $\mathfrak{C}_{\exists,0,m}^{A,A,T}$, then the duplicator mimics it in an isomorphic copy of $\tilde{\mathfrak{A}}_m^{-T_i}$, which is a part of $\mathfrak{C}_{\exists,i,m}^{B,A,T}$.
 - if the spoiler picks inside $\tilde{\mathfrak{A}}_m^{T_j}$ or $\tilde{\mathfrak{A}}_m^{-T_j}$ where $j \neq i$, then the duplicator mimics it in an isomorphic copy of $\tilde{\mathfrak{A}}_m^{T_j}$ or $\tilde{\mathfrak{A}}_m^{-T_j}$, which is a part of $\mathfrak{C}_{\exists,i,m}^{A,B,T}$.
 - if the spoiler picks the junction point, i.e. the root, of $\mathfrak{C}_{\exists,0,m}^{A,A,T}$, then the duplicator picks the junction point of $\mathfrak{C}_{\exists,i,m}^{A,B,T}$.
- If \mathcal{T}_i is not embeddable in \mathcal{F} and $r(\mathcal{T}_i)$ is an \mathcal{A} node, then

- if the spoiler picks inside $\tilde{\mathfrak{A}}_m^{T_i}$, which is a part of $\mathfrak{C}_{\exists,0,m}^{A,A,T}$, the duplicator is able to mimic the spoiler's picks inside $\tilde{\mathfrak{B}}_m^{T_i}$, which is a part of $\mathfrak{C}_{\exists,i,m}^{B,B,T}$. It is because the root of \mathcal{T}_i is an \mathcal{A} node, which means that the number of different types of subtrees, whose roots are connected to the root of $\tilde{\mathfrak{A}}_m^{T_i}$ (in $\tilde{\mathfrak{A}}_m^T$), is less than the number of different types of subtrees, whose roots are connected to the root of $\tilde{\mathfrak{B}}_m^{T_i}$ (in $\tilde{\mathfrak{B}}_m^T$).
- if the spoiler picks inside $\tilde{\mathfrak{A}}_m^{T_j}$ or $\tilde{\mathfrak{A}}_m^{-T_j}$ where $j \neq i$, then the duplicator mimics it in an isomorphic copy of $\tilde{\mathfrak{A}}_m^{T_j}$ or $\tilde{\mathfrak{A}}_m^{-T_j}$, which is a part of $\mathfrak{C}_{\exists,i,m}^{B,B,T}$.
- if the spoiler picks the junction point, i.e. the root, of $\mathfrak{C}_{\exists,0,m}^{A,A,T}$, then the duplicator picks the junction point of $\mathfrak{C}_{\exists,i,m}^{B,B,T}$.

Note that such a strategy has exploited the feature of the structures on the game board: (recall that the root of \mathcal{T} is an \mathcal{E} node) the structures are constructed in such a way that if the spoiler does not pick the junction point of $\mathfrak{C}_{\exists,0,m}^{A,A,T}$ in $\tilde{\mathfrak{A}}_m^T$ then the duplicator is able to mimic his picking in an isomorphic tree substructure of $\tilde{\mathfrak{B}}_m^T$ - such isomorphic subtrees always exist. And by Lemma 3.20 the spoiler can also win without such picking if he can win in any way. If the spoiler pick the junction point of $\mathfrak{C}_{\exists,0,m}^{A,A,T}$, by her strategy, the duplicator picks the junction point of $\mathfrak{C}_{\exists,i,m}^{A,B,T}$ where $\mathcal{T}_i \not\leq_e \mathcal{F}$. Observe that, the difference between $\mathfrak{C}_{\exists,0,m}^{A,A,T}$ and $\mathfrak{C}_{\exists,i,m}^{A,B,T}$ is the difference between $\tilde{\mathfrak{A}}_m^{-T_i}$ and $\tilde{\mathfrak{B}}_m^{-T_i}$ (or equivalently the difference between $\tilde{\mathfrak{A}}_m^{T_i}$ and $\tilde{\mathfrak{B}}_m^{T_i}$). Therefore, using this strategy, no matter how the spoiler picks, the duplicator can reply properly such that in the end the spoiler can win the game $G_{\mathcal{F}}(\tilde{\mathfrak{A}}_m^{-T_i}, \tilde{\mathfrak{B}}_m^{-T_i})$ (or $G_{\mathcal{F}}(\tilde{\mathfrak{A}}_m^{T_i}, \tilde{\mathfrak{B}}_m^{T_i})$), if the spoiler can win in any way. Recall that $rk(\mathcal{T}_i) \leq h$. By induction assumption, the duplicator can win the game using this strategy. Therefore, the spoiler cannot win the game if the duplicator plays according to this strategy. In other words, this strategy is a winning strategy for the duplicator in the game $G_{S_2}(\tilde{\mathfrak{A}}_m^T, \tilde{\mathfrak{B}}_m^T)$.

When $r(\mathcal{T})$ is an \mathcal{A} node, the analysis is similar. □

Note that the duplicator's winning strategy in Lemma 5.6 depends on how the spoiler moves the token. That is, she has to keep an eye on the token track before she make the choices.

The following corollary is directly due to Lemma 3.6 and Theorem 3.8.

Corollary 5.7. *Let S_1 and S_2 be two finite Γ -labelled forests. If $S_1 \not\leq_e S_2$, then for any ψ such that $qs(\psi) \leq_e S_2$, and any $m \geq rk(S_2)$, we have:*

$$\tilde{\mathfrak{A}}_m^{S_1} \models \psi \Rightarrow \tilde{\mathfrak{B}}_m^{S_1} \models \psi.$$

The following lemma is similar to Lemma 3.17. Actually, the proof resembles that of Lemma 3.17.

Lemma 5.8. *Let \mathcal{T} be Γ -labelled irreducible tree. Then*

$$\tilde{\mathfrak{A}}_m^T \models \tilde{\phi}_{\mathcal{T}} \text{ but } \tilde{\mathfrak{B}}_m^T \not\models \tilde{\phi}_{\mathcal{T}}.$$

Proof. It is obvious when \mathcal{T} is a single node. Assume that it holds when $rk(\mathcal{T}) = h$.

Now assume that \mathcal{T} is a k -ary Γ -labelled tree and its rank is $h + 1$. Suppose that the root of \mathcal{T} is an \mathcal{E} node, and its subtrees, whose roots are children of $r(\mathcal{T})$, are $\mathcal{T}_1, \dots, \mathcal{T}_k$. Recall that $H = \{i \in [k] \mid \text{the root of } \mathcal{T}_i \text{ has the same label as that of the root of } \mathcal{T}\}$.

Let a_r be the node that interprets r in $\tilde{\mathfrak{A}}_m^T$. And let b_r be the node that interprets the hook constant in $\mathfrak{C}_{\exists,0,m}^{A,A,T}$, i.e. b_r is the junction point of $\mathfrak{C}_{\exists,0,m}^{A,A,T}$. By assumption and an observation similar to Lemma 3.16, $\tilde{\mathfrak{A}}_m^{T_i} \models \tilde{\xi}_{\mathcal{T}_i}(r)$ ($i \in [k]$), and $\tilde{\mathfrak{A}}_m^{-T_j} \models \tilde{\xi}_{-\mathcal{T}_j}(r)$ ($j \in H$).

Moreover, all the quantifiers in $\tilde{\xi}_{\mathcal{T}_i}(r)$ are relativized by relations either Ryx or Byx , where x is the quantified variable. It means that $\tilde{\xi}_{\mathcal{T}_i}(r)$ expresses some property that is nothing to do with the elements of $\tilde{\mathfrak{A}}_m^T$ outside the tree substructure $\tilde{\mathfrak{A}}_m^{T_i}$. As a consequence, $\tilde{\mathfrak{A}}_m^{T_i} \models \tilde{\xi}_{\mathcal{T}_i}(r)$ implies $\tilde{\mathfrak{A}}_m^T \models \tilde{\xi}_{\mathcal{T}_i}(b_r)$.

By the same argument, $\tilde{\mathfrak{A}}_m^{-T_i} \models \tilde{\xi}_{-\mathcal{T}_i}(r)$ implies $\tilde{\mathfrak{A}}_m^T \models \tilde{\xi}_{-\mathcal{T}_i}(b_r)$.

Therefore, in $\tilde{\mathfrak{A}}_m^T$, b_r is the witness of the quantifier $\exists x_d$ in the formula $\tilde{\xi}_{\mathcal{T}}(r)$. That is, $\tilde{\mathfrak{A}}_m^T \models \tilde{\xi}_{\mathcal{T}}(r)$, or $\tilde{\mathfrak{A}}_m^T \models \tilde{\phi}_{\mathcal{T}}$.

Likewise, by assumption and similar analysis, in $\tilde{\mathfrak{B}}_m^T$ the subtrees $\mathfrak{C}_{\exists,i,m}^{A,B,T}$, or $\mathfrak{C}_{\exists,i,m}^{B,A,T}$, or $\mathfrak{C}_{\exists,i,m}^{B,B,T}$, has exactly one subtree $\tilde{\mathfrak{B}}_m^{T_i}$, or $\tilde{\mathfrak{B}}_m^{-T_i}$, which does not satisfy some subformula of $\tilde{\phi}_{\mathcal{T}}$ (i.e. one conjunct of $\bigwedge_{i \in [k]} \tilde{\xi}_{\mathcal{T}_i}(x_d) \wedge \bigwedge_{j \in H} \tilde{\xi}_{-\mathcal{T}_j}(x_d)$), when its root interprets x_d . In other

words, b_r cannot be a witness of the quantifier $\exists x_d$. Therefore, $\tilde{\mathfrak{B}}_m^T \not\models \tilde{\phi}_{\mathcal{T}}$.

When the root of \mathcal{T} is an \mathcal{A} node, the analysis is similar. \square

Theorem 5.9. *Let S_1 and S_2 be two Γ -labelled forests. Over the class of all finite τ^+ -structures,*

$$\text{if } S_1 \not\leq_e S_2, \text{ then } \text{FO}\{S_1\} \not\subseteq \text{FO}\{S_2\}.$$

Proof. By Lemma 5.6, we know that the duplicator has a winning strategy in the game $G_{S_2}(\tilde{\mathfrak{A}}_m^{S_1}, \tilde{\mathfrak{B}}_m^{S_1})$. In other words, the property defined by $\tilde{\phi}_{\mathcal{T}}$ is not expressible in $\text{FO}\{S_2\}$, by Lemma 3.9. Observe that $\tilde{\phi}_{\mathcal{T}} \in \text{FO}\{S_1\}$. Therefore, $\text{FO}\{S_1\} \not\subseteq \text{FO}\{S_2\}$. \square

Recall that $\tau = \langle E \rangle$. Based on the same transformations as the “reductions from τ^+ to τ ” (p. 20), the diagram \mathfrak{A}_m^T (\mathfrak{B}_m^T resp.) is obtained from $\tilde{\mathfrak{A}}_m^T$ ($\tilde{\mathfrak{B}}_m^T$ resp.), as \mathfrak{A}_m^p (\mathfrak{B}_m^p resp.) is obtained from $\tilde{\mathfrak{A}}_m^p$ ($\tilde{\mathfrak{B}}_m^p$ resp.); $\phi_{\mathcal{T}}$ is obtained from $\tilde{\phi}_{\mathcal{T}}$ as φ_p is obtained from $\tilde{\varphi}_p$. Recall that the reductions are mainly doing three things:

- change red edges to forward edges; change blue edges to backward edges;
- use self-loops to indicate the positions of junction points;
- use bi-directional edges to indicate the positions of black leaves.

We can use the same reductions, prove a new version of Lemma 3.26 as the following:

For any first-order sentence ζ over τ , there is a first-order sentence ξ over τ^+ , with the same quantifier structure, such that

- (1) $\mathfrak{A}_m^T \models \zeta$ iff $\tilde{\mathfrak{A}}_m^T \models \xi$;
- (2) $\mathfrak{B}_m^T \models \zeta$ iff $\tilde{\mathfrak{B}}_m^T \models \xi$.

And Theorem 5.1 is immediate, by the same argument in the proof of Theorem 3.1.

The following is a direct corollary of Theorem 5.1:

Corollary 5.10. *Let S_1 and S_2 be two Γ -labelled forests and σ_g includes a k -ary relation symbol where $k \geq 2$. Over the class of all finite σ_g -structures,*

$$\text{if } S_1 \not\leq_e S_2, \text{ then } \text{FO}\{S_1\} \not\subseteq \text{FO}\{S_2\}.$$

Proof. We can use the k -ary relation ($k \geq 2$) to encode the binary relations we defined in the structures for Theorem 5.1 and let all the other relations be empty. \square

6. SUMMARY

It is natural to classify fragments of first-order logic based on quantifier structures. So far, most related works are focused on those fragments based on quantifier prefixes, a special kind of quantifier structure, and contribute to our understanding of their different expressiveness. Grädel-McColm’s conjecture, which claims the strictness of the first-order prefix hierarchy, generalizes the results of Walkoe [10], Keisler & Walkoe [3] and Chandra & Harel [1]. Rosen proved this conjecture by showing that it holds over infinite structures [6] and raised the question of whether it also holds over finite structures. We define games that characterize the quantifier classes, which generalize the standard Ehrenfeucht-Fraïssé games, and prove that a similar and natural hierarchy, i.e. the first-order quantifier structure hierarchy, is strict over finite structures. Although these two hierarchies are similar, they are independent in that none implies the other directly. Nevertheless, our constructions do provide justifications for some special cases of Grädel-McColm’s conjecture over finite structures. For example, from the constructions introduced in this paper we can see that there is a property that is expressible in $\text{FO}(\exists\forall\forall)$, but not in $\text{FO}(\forall\forall\forall)$, $\text{FO}(\forall\forall\exists)$, $\text{FO}(\forall\exists\exists)$ and $\text{FO}(\exists\exists\exists)$. But we don’t know whether it is expressible in $\text{FO}(\forall\exists\forall)$ or not.

Recall that the mapping ι between quantifier structures is not necessarily injective when we define the embedding relation \preceq_e (Definition 2.7). Now, we change the definition of quantifier structure embedding a little bit, i.e. require ι to be injective. As a consequence, the quantifier classes are also changed correspondingly. And a finite version of Rosen’s main theorem [6] can be stated as the following, based on such change. Recall that we assume that all the structures are finite.

Let $p \in \Gamma^*$ and S be a Γ -labelled forest, $\text{FO}(p) \not\preceq_e \text{FO}\{S\}$ if $p \notin \mathcal{W}(S)$.

As has been mentioned in the introduction, a proof of this theorem will solve Grädel-McColm’s conjecture over finite structures, which is still open at present. Another possible way to resolve this conjecture is to define Ehrenfeucht-Fraïssé style games for the prefix classes. In any cases, we need new techniques and insights.

Let \mathcal{T} be a Γ -labelled tree and \mathcal{T}^i be a forest composed of i disjoint copies of \mathcal{T} . Let μ_t be a map from a Γ -labelled tree to a natural number such that $\mu_t(\mathcal{T})$ is the minimum number for m such that $\text{FO}\{\mathcal{T}^m\} = \text{FO}\{\mathcal{T}^{m+1}\}$.³ Then, the following is conceivable, under the new definition of quantifier structure embedding:

Let \mathcal{T} be a Γ -labelled tree and S be an arbitrary Γ -labelled forests. Over the class of all finite digraphs, for any $m \leq \mu_t(\mathcal{T})$,

$$\text{if } \mathcal{T}^m \not\preceq_e S, \text{ then } \text{FO}\{\mathcal{T}^m\} \not\preceq_e \text{FO}\{S\}.$$

Note that it generalizes not only Grädel-McColm’s conjecture over finite structures but also the finite version of Rosen’s main theorem. Furthermore, it is also interesting to study $\mu_t(\mathcal{T})$. For example, how fast will it grow w.r.t. the size of \mathcal{T} ?

Another question is whether we can prove similar hierarchies in other logics. One candidate is so called independence-friendly logic (IFL). What makes it interesting is that IFL has the form of first-order logic, while has the expressive power equals existential

³The existence of such a number m is determined by Lemma 3.3.

second-order logic (ESO). It is well-known that ESO captures the complexity class NP over finite structures. Clearly, establishing a natural and strict hierarchy for the NP problems would be very interesting.

ACKNOWLEDGEMENT

I am grateful to Anuj Dawar for giving me a lot of suggestions and ideas which improve the paper significantly and to Bjarki Holm for useful discussions and proofreading. I am also grateful to the anonymous referees for their valuable comments.

REFERENCES

- [1] A. Chandra and D. Harel. Structure and complexity of relational queries. *Journal of Computer and System Science*, 25:99–128, 1982.
- [2] E. Grädel and G. McColm. Hierarchies in transitive closure logic, stratified datalog and infinitary logic. *Annals of Pure and Applied Logic*, 77(2):169–199, 1996.
- [3] H. Jerome Keisler and JR. Wilbur Walkoe. The diversity of quantifier prefixes. *The Journal of Symbolic Logic*, 38:79–85, 1973.
- [4] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [5] E. Rosen. Some aspects of model theory and finite structures. *Bulletin of Symbolic Logic*, 8(3):380–403, 2002.
- [6] E. Rosen. On the first-order prefix hierarchy. *Notre Dame J. Formal Logic*, (online draft, 1998), 46:147–164, 2005.
- [7] N. Schweikardt. Arithmetic, first-order logic, and counting quantifiers. *ACM Trans. Comput. Logic*, 6(3):634–671, 2005.
- [8] N. Schweikardt and T. Schwentick, “A note on the expressive power of linear orders,” *Logical Methods in Computer Science*, 7(4), pp.1–13, 2011.
- [9] M. Sipser. Borel sets and circuit complexity. In *Proceedings of the fifteenth annual ACM symposium on Theory of computing*, pages 61–69. ACM, 1983.
- [10] W. Walkoe. Finite partially-ordered quantification. *Journal of Symbolic Logic*, 35:535–555, 1970.