On the first-order prefix hierarchy

Eric Rosen
RWTH Aachen, Germany
erosen@informatik.rwth-aachen.de

May, 1998

Abstract

We investigate the expressive power of fragments of first-order logic that are defined in terms of prefixes. The main result establishes a strict hierarchy among these fragments over the signature consisting of a single binary relation. It implies that for each prefix p, there is a sentence φ_p in prenex normal form with prefix p, over a single binary relation, such that for all sentences θ in prenex normal form, if θ is equivalent to φ_p, then p can be embedded in the prefix of θ. This strengthens a theorem of Walkoe.

In this paper we address the following question. Given two first-order prefixes p and q, is there a sentence with prefix p that is not equivalent to any sentence with prefix q? Walkoe [4] proved that if p and q are different prefixes of length n, then there is such a sentence, containing a single n-ary relation symbol. Keisler and Walkoe [3] then strengthened this result by showing that it also holds over the class of finite structures. Our main theorem improves on Walkoe's result. It implies that for each prefix p, there is a sentence φ_p in prenex normal form with prefix p, over a single binary relation, such that for all sentences θ in prenex normal form, if θ is equivalent to φ_p, then p can be embedded in the prefix of θ. (We leave its precise statement to Section 1.) This also resolves a conjecture of Grädel and McColm [1], explained below.

1 Background and statement of the main theorem

1.1 Terminology and definitions

We adopt the following terminology and conventions. We will consider (fragments of) first-order logic (FO) and infinitary logic ($L_{\infty\omega}$), which allows infinitary conjunctions and disjunctions. Throughout we assume that formulas are in negation normal form, that is, negation symbols only bind atomic formulas. Signatures are always purely relational and finite. A graph is a structure with signature $\{E\}$, E binary; it is simple if it is undirected and loop-free. In a (directed) graph, we say that vertex a is an *E-predecessor* of vertex b, if $Eab \wedge \neg Eba$. Throughout, we assume that any class of structures we are talking about is closed under isomorphism.

Both structures and their universes are denoted $A, B, \ldots$, etc. If A is a σ-structure, and $R \in \sigma$, then R^A denotes the interpretation of R in A. For $\tau \subseteq \sigma$, $A|\tau$ is the reduct of A to τ, a τ-structure. If $\psi(x)$ is a σ-formula with exactly one free variable, then

$A{\restriction}\psi(x)$ is the substructure of A with universe $\psi(A) = \{a \in A \mid A \models \psi[a]\}$. Given a tuple of elements $\overline{a}$ in A, $A{\restriction}\overline{a}$ is the substructure of A with universe $\overline{a}$. $A{\restriction}\overline{a} \cong B{\restriction}\overline{b}$ means moreover that the function from $A{\restriction}\overline{a}$ to $B{\restriction}\overline{b}$ that takes each $a_i \in \overline{a}$ to the corresponding $b_i \in \overline{b}$ is an isomorphism. The *positive diagram* of a σ-structure A is the set of all atomic sentences true over the signature σ_A, which contains, additionally, a constant for each $a \in A$. (In this case, we will not distinguish between an element and its name.)

A prefix p is a non-null finite string of $\exists$s and $\forall$s. The *dual* of p, denoted $\overline{p}$, is the prefix obtained by swapping occurences of $\exists$s and $\forall$s. A Σ-, resp. Π-prefix is a prefix beginning with $\exists$, resp. $\forall$.

We recall the following basic concepts from formal language theory. An *alphabet* Σ is a finite set of symbols and a *word* is a non-null finite string of symbols. $(\Sigma)^*$ denotes the set of words over the alphabet Σ. A *language* is a set of words in $(\Sigma)^*$. The concatenation of two words is written $p * q$. When p is a word and Q a language, we use $p * Q$ to denote the set $\{p * q \mid q \in Q\}$.

We view prefixes as words over the alphabet $\Sigma_1 = \{\exists, \forall\}$. We use $L_1 = (\{\exists, \forall\})^*$ to denote the set of all prefixes. For φ a first-order formula in *prenex normal form*, $pr(\varphi)$ is the prefix of φ. We define a partial order on prefixes as follows: $p \leqslant q$ iff p can be obtained from q by removing symbols from the latter. In this case, p is a (not necessarily initial) subword of q.

Given a prefix p, we define the *quantifier alternation number* of p, $alt(p)$, to be the number of quantifier blocks that p contains. Each such p, with $alt(p) = n$, can be written *succinctly*, in the obvious way, as $(s_1)^{i_1}(s_2)^{i_2}\dots(s_n)^{i_n}$, $i_m \in \omega$ and $s_m \in \{\exists, \forall\}$, $m \leqslant n$; $s_1 \neq s_2$; and for $l, m \leqslant n$, $s_l = s_m$ iff $|l - m|$ is even. For example, $\exists\exists\forall\forall\exists\forall\forall\forall$ is written $\exists^2\forall^2\exists\forall^3$.

We also consider words over $\Sigma_2 = \{\exists, \forall, \exists^*, \forall^*\}$, which we treat as regular expressions that denote regular languages. For example, $\exists\forall^*\exists\forall$ denotes the set $\{\exists\forall^n\exists\forall \mid n \in \omega\}$. Let $L_2 = (\{\exists, \forall, \exists^*, \forall^*\})^*$. A string v is in *reduced form* iff occurences of $\exists$, $\exists^*$ alternate strictly with occurences of $\forall$, $\forall^*$. Let $r(x) : \Sigma_2 \longrightarrow L_1$ be the map that takes a regular expression to the regular language it denotes, and define $r^-(x) : \Sigma_2 \longrightarrow L_1$ so that $r^-(v) = \{q \mid \text{there is an } q' \in r(v) \text{ and } q \leqslant q'\}$, the downward closure of $r(v)$.

The next lemma will be useful later.

Lemma 1.1 *For every prefix $p \in L_1$, there is a (unique) word $f(p) \in L_2$ in reduced form, so that $r^-(f(p)) = \{q \mid p \nleqslant q\}$.*

Proof. Define $f : L_1 \longrightarrow L_2$ as follows.

(i) If $p = \exists^n$, then $f(p) = a_1 \dots a_{2n-1}$, where $a_i = \forall^*$ for i odd, and $= \exists$ for i even.

(ii) If $p = \forall^n$, then $f(p) = a_1 \dots a_{2n-1}$, where $a_i = \exists^*$ for i odd, and $= \forall$ for i even.

(iii) If $p = (s_1)^{i_1} \dots (s_n)^{i_n}$, then $f(p) = f((s_1)^{i_1}) * \dots * f((s_n)^{i_n})$.

We argue by induction on the length of p. Observe that for all p, if p is a Σ-, resp. Π-prefix, then the first symbol of $f(p)$ is $\forall^*$, resp. $\exists^*$. For $p = \exists$, $f(\exists) = \forall^*$, and it is clear that $r^-(f(\exists)) = \{\forall^m \mid 0 \leqslant m\} = \{q \in L_1 \mid \exists \nleqslant q\}$, as desired. Likewise for $p = \forall$.

Now suppose that $p = \exists * p'$ is a Σ-prefix of length $n+1$, and the lemma holds for all prefixes of length $\leqslant n$. For each prefix $q = (\forall^m \exists) * q'$, $0 \leqslant m$, containing at least one $\exists$, $p \not\leqslant q$ iff $\exists * p' \not\leqslant (\forall^m \exists) * q'$ iff $p' \not\leqslant q'$. Therefore $\{q \mid p \not\leqslant q\} = \{\forall^m \mid 0 \leqslant m\} \cup \{(\forall^m \exists) * q' \mid 0 \leqslant m \text{ and } p' \not\leqslant q'\}$. Invoking the induction hypothesis, it is clear that this set is equal to $r^-((\forall^* \exists) * f(p'))$. If p' is a Σ-prefix, then $(\forall^* \exists) * f(p')$ is simply $f(p)$. In case p' is a Π-prefix, then the first symbol of $f(p')$ is an $\exists^*$, and clearly $r^-((\forall^* \exists) * f(p')) = r^-((\forall^*) * f(p'))$; again $\forall^* * f(p')$ is $f(p)$. The argument for Π-prefixes is dual. $\square$

We are interested in fragments of FO defined in terms of prefixes.

Definition 1.2 *For each prefix p, we define the* prefix class $\mathrm{FO}(p)$ *as the set* $\{\theta \mid \theta$ *is a FO formula in prenex normal form such that* $\mathrm{pr}(\theta) \leqslant p\}$.

More generally, we want to assign a 'quantifier structure' to every (FO and) $L_{\infty\omega}$ formula, which will be a set of prefixes. The following definition is from Grädel and McColm [1].

Definition 1.3 *The* quantifier structure *of a formula* $\varphi \in L_{\infty\omega}$, $qs(\varphi)$, *is defined inductively as follows.*

(i) If φ *is a literal, then* $qs(\varphi) = \emptyset$.

(ii) If $\varphi = \bigwedge_i \theta_i$ *or* $\bigvee_i \theta_i$, *then* $qs(\varphi) = \bigcup_i qs(\theta)$.

(iii) If $\varphi = \exists x\theta$, *resp.* $\forall x\theta$, *then* $qs(\varphi) = \exists * qs(\theta) \cup qs(\theta) \cup \{\exists\}$, *resp.* $qs(\varphi) = \forall * qs(\theta) \cup qs(\theta) \cup \{\forall\}$.

Observe that we have defined the quantifier structure of a formula so that it is always a set of prefixes closed under subwords. We will reserve the term *prefix set* for such sets of prefixes.

Quantifier classes are defined in analogy with prefix classes.

Definition 1.4 *Let* $\mathcal{L}$ *be either* FO *or* $L_{\infty\omega}$, *and let* P *be a prefix set. We define the quantifier class* $\mathcal{L}\{P\} = \{\theta \in \mathcal{L} \mid qs(\theta) \subseteq P\}$.

When v is a word in L_2, we will abuse the notation and write $L_{\infty\omega}\{v\}$ rather than the more cumbersome $L_{\infty\omega}\{r^-(v)\}$.

Observe that for all prefixes p, and all prefix sets P such that $p \in P$, we have that $\mathrm{FO}(p) \subseteq \mathrm{FO}\{p\} \subseteq L_{\infty\omega}\{p\} \subseteq L_{\infty\omega}\{P\}$. (In fact it is clear that, over any finite signature, every $L_{\infty\omega}\{p\}$ formula is equivalent to a $\mathrm{FO}\{p\}$ formula.) In particular, $\mathrm{FO}(p) \subseteq L_{\infty\omega}\{P\}$. The main theorem is a strong converse to this.

The main theorem

Theorem 1.5 *Let* p *be a prefix, and let* $P = \{q \in L_1 \mid p \not\leqslant q\}$. *There is a sentence* φ_p *in* $\mathrm{FO}(p)$, *over a single binary relation, such that* φ_p *is not equivalent to any sentence in* $L_{\infty\omega}\{P\}$.

(Observe that P is simply $r^{-}(f(p))$.) The following corollary is immediate.

Corollary 1.6 *Let P_1 and P_2 be prefix sets such that $P_1 \not\subseteq P_2$. Then there is a* (FO) *sentence $\varphi \in L_{\infty\omega}\{P_1\}$, over a single binary relation, such that φ is not equivalent to any $\theta \in L_{\infty\omega}\{P_2\}$.*

Theorem 1.5 clearly implies the following conjecture of Grädel and McColm [1].

> For all prefixes p and q, if p cannot be embedded in q, then there is a sentence containing a single binary relation with prefix p that is not equivalent to any sentence with prefix q.

Finite structures and an open question

It is an open question whether the Grädel-McColm conjecture holds over the class of finite structures. It is easy to show that the main theorem itself does not, as there is a collapse of the $L_{\infty\omega}$ quantifier class hierarchy to $L_{\infty\omega}\{r(\forall^*) \cup r(\exists^*)\}$. This is because every finite structure of size n can be described up to isomorphism by a sentence in FO$\{\exists^n, \forall^{n+1}\}$, so *every* class of finite structures is defined by a countable disjunction of such sentences, which is a sentence in $L_{\infty\omega}\{r(\forall^*) \cup r(\exists^*)\}$.

1.2 Definability and games

Definition 1.7 *Let $\mathcal{L}$ be a logical language. Given structures A and B, we write $A \Rightarrow_{\mathcal{L}} B$ iff for every $\theta \in \mathcal{L}$, if $A \models \theta$, then $B \models \theta$.*

The following definitions and result in this section are essentially from Grädel and McColm [1]. Below we restrict our attention to fragments of the form $L_{\infty\omega}\{q\}$, for q a word in L_2. as we will not need to consider the more general case, which is quite similar.

Definition 1.8 *Let $v = a_1 \dots a_n$ be a word in L_2, that is, $v \in \{\exists, \exists^*, \forall, \forall^*\}^*$. Given structures A and B, the $L_{\infty\omega}\{v\}$-game from A to B is an n-round game played between a Spoiler (S.) and a Duplicaotr (D.), with four types of rounds, depending on whether a_i is $\exists, \forall, \exists^*$, or $\forall^*$.*

[$\exists$ round] *The S. plays a (single) pebble on A. The D. then plays a pebble on B.*

[$\forall$ round] *The S. plays a pebble on B. The D. then plays a pebble on A.*

[$\exists^*$ round] *The S. plays a pebble on A. The D. then plays a pebble on B. The S. may repeat this as often as he wants to. That is, he is permitted to play arbitrarily many '$\exists$ moves' in a single $\exists^*$ round.*

[$\forall^*$ round] *Like an $\exists^*$ round, with the S. playing instead on B.*

The S. wins if at any point the pebbles do not determine a partial isomorphism from A to B.

An equivalent description of this game can be given as follows. The players play an ordinary (infinite) Ehrenfeucht-Fraïssé game, with the following additional restriction placed on the S.'s moves. We associate with each play of the game through n rounds a prefix p_n of length $n, p_n = s_1 \dots s_n$, such that for all $i \leqslant n$, $s_i = \exists$ if the S. played on A in the i^{th} round, and $= \forall$ otherwise. In each round n, the S. must choose a structure to play on so that the associated prefix p_n is in $r^{-}(f(v))$. (If this set is finite, then he is only premitted to play some fixed finite number of rounds.)

Proposition 1.9 (Grädel and McColm [1]) *Let A and B be structures, and let v be a word in L_2. D. has a winning strategy in the $L_{\infty\omega}\{v\}$-game from A to B iff $A \Rightarrow_{L_{\infty\omega}\{v\}} B$.*

1.3 Amalgamation classes and homogeneous structures

In this section, we present the model theoretic background used to construct structures in the proof of the main theorem. (For more information, see e.g. Hodges [2].) Recall that a structure A is *homogeneous* if every partial isomorphism between finite substructures of A can be extended to an automorphism of A.

Definition 1.10 *Let $\mathcal{K}$ be a class of finite structures over a finite relational language. $\mathcal{K}$ is an* amalgamation class *if it satisfies the following properties.*

(i) (downward closure) If $B \in \mathcal{K}$ and $A \subseteq B$, then $A \in \mathcal{K}$.

(ii) (joint embedding) If $A_1, A_2 \in \mathcal{K}$, then there is a $B \in \mathcal{K}$ with substructures $A_i' \cong A_i, i = 1, 2$.

(iii) (amalgamation) If $A_0, A_1, A_2 \in \mathcal{K}$ and $f_i : A_0 \longrightarrow A_i, i = 1, 2$, are embeddings, then there are a $B \in \mathcal{K}$ and embeddings $g_i : A_i \longrightarrow B$ such that $g_1 \circ f_1 = g_2 \circ f_2$.

If A is a structure, then $Sub(A)$ is the set of all finite structures B isomorphic to a substructure of A. The following result is due to Fraïssé.

Theorem 1.11 *Let A be a homogeneous structure. Then $Sub(A)$ is an amalgamation class. Conversely, if $\mathcal{K}$ is an amalgamation class, then there is a unique, up to isomorphism, finite or countable structure, such that $Sub(A) = \mathcal{K}$.*

In this case, A is called the *Fraïssé limit* of $\mathcal{K}$, denoted $Fr(\mathcal{K})$.

Definition 1.12 *Let $\mathcal{K}$ be a class of finite structures that is downwardly closed. Then A is a* constraint *of $\mathcal{K}$ iff $A \notin \mathcal{K}$, but for all proper substructures $B \subsetneq A, B \in \mathcal{K}$.*

Definition 1.13 *Let $\mathcal{J}$ be any class of finite structures. Define $Cl(\mathcal{J}) = \{A \mid A$ is a finite structure that has no substructure isomorphic to any $B \in \mathcal{J}\}$.*

Let $\mathcal{J}$ be any set of finite structures such that for all distinct $A, B \in \mathcal{J}$, A does not embed as a substructure in B. Observe that $Cl(\mathcal{J})$ is the unique downwardly closed class whose set of constraints is exactly $\mathcal{J}$. We now define a property of structures that can be used to show that $Cl(\mathcal{J})$ is, additionally, an amalgamation class.

Definition 1.14 *Let A be a σ-structure. Then A is* irreducible *iff for all $a_1, a_2 \in A, a_1 \neq a_2$, there are a k-ary relation $R \in \sigma$ and a k-tuple $\overline{a}$ in A containing a_1, a_2, such that $A \models R\overline{a}$.*

Observe that every A of cardinality 1 is, by default, irreducible.

The next lemma is straightforward.

Lemma 1.15 *Let $\mathcal{J}$ be a set of finite, irreducible σ-structures. Then $Cl(\mathcal{J})$ is a (strong) amalgamation class.*

Proof. It is clear that $Cl(\mathcal{J})$ is downwardly closed. We now argue that it has the joint embedding property. Let A_1, A_2 be two structures in $Cl(\mathcal{J})$ and assume that their universes are disjoint, $A_1 \cap A_2 = \emptyset$. Define B to be the disjoint union of A_1 and A_2, that is, the model with universe $A_1 \cup A_2$ such that for all relations $R \in \sigma$, $R^B = R^{A_1} \cup R^{A_2}$. Clearly any irreducible substructure of B is either a substructure of A_1 or of A_2. Therefore B is also in $Cl(\mathcal{J})$.

A similar argument shows that $Cl(\mathcal{J})$ has the amalgamation property. Again without loss of generality, let A_0, A_1, A_2 be structures in $Cl(\mathcal{J})$ such that $A_0 \subseteq A_i$ and $f_i : A_0 \to A_i$ is the identity map, for $i = 1, 2$, and $A_1 \cap A_2 = A_0$. Define B to be the model with universe $A_1 \cup A_2$ such that, as above, for all relations $R \in \sigma$, $R^B = R^{A_1} \cup R^{A_2}$. Arguing as before, we get that B is also in $Cl(\mathcal{J})$. $\square$

Combining this lemma with Fraïssé's theorem provides an easy way to (describe and) produce homogeneous structures.

Proposition 1.16 *Given any set of finite irreducible structures $\mathcal{J}$, $Fr(Cl(\mathcal{J}))$ is a homogenous structure.*

Observe that for any countable homogeneous structure A, there is a unique set of pairwise mutually non-embeddable finite structures $\mathcal{J}$ such that $A = Fr(Cl(\mathcal{J}))$. The following simple lemma will be important in Section 2.3.

Lemma 1.17 *Let A be a countable homogeneous structure, and let $\mathcal{J}$ be a set of finite structures such that $A = Fr(Cl(\mathcal{J}))$. Suppose that $\overline{a}$ is a k-tuple of distinct elements in A and $\psi'(\overline{x}, \overline{x}')$ is a quantifier-free formula, with $\overline{x} = x_1, \ldots, x_k$, and $\overline{x}' = x_{k+1}, \ldots, x_{k+l}$. Then there is an l-tuple $\overline{a}'$ of distinct elements in A, pairwise disjoint from $\overline{a}$, such that $A \models \psi'[\overline{a}, \overline{a}']$ iff there is a structure B of size $k+l$, with universe $\overline{b} \cup \overline{b}'$, $\overline{b}$ and $\overline{b}'$ tuples of length k and l, respectively, such that $A{\restriction}\overline{a} \cong B{\restriction}\overline{b}$, $B \models \psi'[\overline{b}, \overline{b}']$, and no substructure of B is isomorphic to any of the constraints $C \in \mathcal{J}$.*

2 Proof of the main theorem

Before presenting the technical details, we outline the structure of the proof. First we define, for each prefix p, countable structures A_p and B_p. These are constructed as reducts of homogeneous structures produced as Fraïssé limits of amalgamation classes. We then define sentences $\varphi_p \in \mathrm{FO}(p)$, for all prefixes p, and prove that $A_p \models \varphi_p$

and $B_p \not\models \varphi_p$. Finally we prove, using a game theoretic argument, that for each p, $A \Rightarrow_{L_{\infty\omega}\{P\}} B$, where $P = \{q \in L_1 \mid p \not\leqslant q\}$. Clearly, this shows that φ_p is not equivalent to any $L_{\infty\omega}\{P\}$ sentence, as desired. The last step is the longest and most difficult part of the proof.

The argument divides naturally into two cases, depending on whether or not $alt(p) = 1$. It is much simpler for $alt(p) = 1$, though already this case contains all the basic elements of the more general situation.

Observe that it suffices to prove the theorem for Σ-prefixes. For p a Π-prefix, it is clear that we can let $\varphi_p = \neg\varphi_{\overline{p}}$ (recall that $\overline{p}$ denotes the dual of p, a Σ-prefix).

2.1 The construction

Let $p = \exists^{n_1} \dots (s_k)^{n_k}$ be a Σ-prefix, $alt(p) = k$. We divide the construction into two cases, depending on whether or not $k = 1$.

It will be helpful to introduce the following defined predicates. Recall that the n-clique K_n is the simple, complete graph of size n. (The 1-clique has one vertex and no edges.)

Definition 2.1 *We define the following formulas.*

[Clique$_n$] *Let* $Kl_n(x_1, \dots x_n) = (\bigwedge_{i \leqslant n} \neg Ex_ix_i) \wedge (\bigwedge_{i<j\leqslant n}(Ex_ix_j \wedge Ex_jx_i))$.

[Arrow$_{m,n}$] *Let* $Ar_{m,n}(x_1, \dots, x_m, y_1, \dots, y_n) =$

$$Kl_m(\overline{x}) \wedge Kl_n(\overline{y}) \wedge (\bigwedge_{i\leqslant m, j\leqslant n} (Ex_iy_j \wedge \neg Ey_jx_i)).$$

Generally, when the length of the tuples is clear from the context, we will omit the subscripts and write simply, e.g., $Kl(\overline{x})$ and $Ar(\overline{x};\overline{y})$. We say that a j-tuple $\overline{a}$ in A is a j-clique if $A \models Kl_j[\overline{a}]$. Given tuples $\overline{a}, \overline{b}$ in A, $\overline{a}$ *arrows* $\overline{b}$ just in case $A \models Ar(\overline{a}; \overline{b})$. Observe that if $\overline{a}$ arrows $\overline{b}$, then $\overline{a}$ and $\overline{b}$ must be disjoint.

Let D_1 be the graph with one vertex and a loop, and let D_2 be the graph with two vertices, a, b, and a single directed edge from a to b. Observe that $\mathcal{G} = Cl(\{D_1, D_2\})$ is the set of all finite, simple graphs, and $Fr(\mathcal{G})$ is the *(countable) random*, or *Rado*, graph.

Case I $alt(p) = 1$

Let $p = \exists^n$. For $n = 1$, let $\varphi_p = \exists x Exx$. Let $A_\exists$ be the countable infinite graph with exactly one edge, which is a loop. Let $B_\exists$ be the empty, countable infinite graph, with no edges. Clearly $A_\exists \models \varphi_p$ and $B_\exists \not\models \varphi_p$.

Now suppose that $n \geqslant 2$. We choose φ_p to be $\exists x_1 \dots x_n Kl(\overline{x})$, which says that there is a substructure isomorphic to K_n. We define B_p directly as the universal countable homogeneous K_n-free graph. (Letting $\mathcal{J}_p^B = \{D_1, D_2, K_n\}$, B_p is simply $Fr(Cl(\mathcal{J}_p^B))$.) It is clear that $B_p \not\models \varphi_p$. We will define A_p as the E-reduct of a countable homogenous structure with signature $\sigma_p = \{E, S\}$, S unary. Let $\mathcal{J}_p^A$ be the set of constraints $\mathcal{J}_p^A = \{D_1, D_2\} \cup \{C \mid C|E \cong K_n$ and $C \models \exists x \neg Sx\}$, and let $A_p^+ = Fr(Cl(\mathcal{J}_p^A))$. Finally, let $A_p = A_p^+|E$. It is not hard to see that $A_p^+ \restriction S$ is isomorphic to the random graph. In particular, $A_p \models \varphi_p$, as desired.

The following easy fact will be needed later. For each n-tuple $\overline{a}$ in A_p, if $\overline{a}$ is an n-clique, then every $a_i \in \overline{a}$ is in S.

Case II $alt(p) \geqslant 2$

Given $p = \exists^{n_1} \dots (s_k)^{n_k}$, let $\sigma_p = \{E\} \cup \{P_i \mid 1 \leqslant i \leqslant k, P_i \text{ is unary}\} \cup \{R_i \mid 1 \leqslant i \leqslant k-1, R_i \text{ is } n_i+1\text{-ary}\} \cup \{S_{ij} \mid 1 \leqslant i < k, 1 \leqslant j \leqslant n_i, S_{ij} \text{ is } j\text{-ary}\} \cup \{S_{k1}\}$, S_{k1} unary. Using ideas from Section 1.3, we produce homogeneous σ_p-structures, A_p^+, B_p^+, and then define A_p and B_p to be $A_p^+|E$ and $B_p^+|E$, respectively. Below, we use $\overline{S}_i$, $1 \leqslant i \leqslant k-1$, for the set $\{S_{i1}, \dots, S_{in_i}\}$ and let $\overline{S} = \bigcup_{i<k} \overline{S}_i \cup \{S_{k1}\}$.

The interpretation of a j-ary relation symbol Q in a structure A is *simple* if for all j-tuples $\overline{a}$ in A, if $A \models Q\overline{a}$, then $\overline{a}$ consists of j distinct elements, and for every pair of j-tuples $\overline{a}, \overline{a}'$ consisting of the same j elements, $A \models Q\overline{a}$ iff $A \models Q\overline{a}'$. (For example, the interpretation of a binary relation symbol E is simple just in case it determines a simple graph.)

Definition 2.2 *Let $\mathcal{C}_p$ be the class of finite σ_p-structures A such that:*

(i) A is loop-free, i.e. $A \models \forall x \neg Exx$.

(ii) A is partitioned by the P_i's. For all $i \leqslant k$, $S_{i1}^A \subseteq P_i^A$; $S_{k1}^A = P_k^A$.

(iii) For $a, b \in A$, if $Eab \wedge Eba$, then there is an $i \leqslant k, P_i a \wedge P_i b$.

(iv) For $a, b \in A$, if $Eab \wedge \neg Eba$, then there is an $i < k, P_i a \wedge P_{i+1} b$.

(v) For each $i < k$, n_i-tuple $\overline{a} = a_1, \dots, a_{n_i}$ in A, and $b \in A$, if $A \models R_i(\overline{a}, b)$, then (i) $Kl(\overline{a})$; (ii) the elements in $\overline{a}$ are pairwise distinct and are each in P_i; (iii) $b \in P_{i+1}$; (iv) for each $a \in \overline{a}$, $A \models Eab$; (v) for each n_i-tuple $\overline{a}'$ containing exactly the same elements as $\overline{a}$, $R_i(\overline{a}', b)$.

(vi) For each relation symbol S_{ij}, $A^{S_{ij}}$ is simple.

(vii) For each relation symbol S_{ij}, and each j-tuple $\overline{a}$, if $A \models S_{ij}\overline{a}$, then (i) $\overline{a}$ is a j-clique of elements in P_i; (ii) for all subsequences $\overline{a}' \subseteq \overline{a}$ of length j', $S_{ij'}\overline{a}'$; (iii) $A \models \forall x(P_k x \leftrightarrow S_{k1} x)$.

(viii) For each $i < k$, A satisfies the sentence $\psi_i =$

$$\neg\exists x_1 \dots x_{n_i} y_1 \dots y_{n_{i+1}} \Big(\bigwedge_{m \leqslant n_i} P_i x_m \wedge \bigwedge_{m \leqslant n_{i+1}} P_{i+1} y_m \wedge \mathrm{Ar}(\overline{x}; \overline{y}) \wedge \bigvee_{m \leqslant n_{i+1}} \neg R_i(\overline{x}, y_m) \Big).$$

The following lemma is easy to establish.

Lemma 2.3 *$\mathcal{C}_p$ is an amalgamation class defined by a finite set of irreducible constraints.*

Let $\mathcal{J}_p$ be the unique minimal such set.

Sketch of proof. Each of the $(viii)$ conditions above on finite structures $A \in \mathcal{C}_p$ is equivalent to saying that A has no substructure isomorphic to one or more irreducible structures. For example, (i) holds of A iff it has no substructure isomorphic to D_1 (the graph with one vertex and a loop). It is an easy exercise to show that this is true also of the remaining of the conditions.

Let $\mathcal{J}_p$ be the minimal such set of irreducible structures.

We now define a number of irreducible σ_p-structures, which will serve as additional constraints.

Definition 2.4 *(a) For each $i < k$, define M_i^p to be the structure with universe $\{1, 2, \ldots, n_i + 1\}$ and positive diagram $\{P_i m \mid 2 \leqslant m \leqslant n_i + 1\} \cup \{P_{i+1} 1\} \cup \{S_{(i+1)1} 1\} \cup \{R_i(\overline{a}, 1) \mid \overline{a}$ is an n_i-tuple containing exactly $\{2, \ldots, n_i + 1\}\} \cup \bigcup_{j \leqslant n_i} \{S_{ij} \overline{m} \mid \overline{m}$ is a j-tuple of distinct elements in $\{2, \ldots, n_i + 1\}\}$.*

(b) Let N^p be the σ_p-structure with universe $\{1\}$ and positive diagram $\{P_1 1, S_{11} 1\}$.

We are now ready to define the structures A_p and B_p.

Definition 2.5 *(a) Let $\mathcal{J}_p^A = \mathcal{J}_p \cup \{M_i^p \mid i < k\}$ and $\mathcal{J}_p^B = \mathcal{J}_p \cup \{M_i^p \mid i < k\} \cup \{N^p\}$.*

(b) Let $A_p^+ = \mathrm{Fr}(Cl(\mathcal{J}_p^A))$ and $B_p^+ = \mathrm{Fr}(Cl(\mathcal{J}_p^B))$.

(c) Let $A_p = A_p^+ | E$ and $B_p = B_p^+ | E$.

By Corollary 1.16, A_p^+ and B_p^+ are indeed homogeneous structures. Observe, on the other hand, that the reducts A_p and B_p are not.

We make a number of observations about these structures. For any element a in A_p^+ or B_p^+, define *height*(a) to be the unique $i \leqslant k$ such that $P_i a$. When $\overline{a}$ is a tuple of elements that all have the same height, we sometimes write simply *height*$(\overline{a})$, instead of *height*(a_1), $a_1 \in \overline{a}$.

Observations

Observations (1),(3), and (2)(i) and (ii) hold equally for both A_p^+ and B_p^+; (2)(iii) does not.

(1) (i) For all a, b, if there is an undirected, resp. directed, edge from a to b, then *height*$(a) =$ *height*(b), resp. *height*$(b) =$ *height*$(a) + 1$. Furthermore, a has no E-predecessors iff *height*$(a) = 1$.

(ii) For all j-ary relations $T \in \sigma_p, j \geqslant 2$, and all j-tuples $\overline{a}$, if $T\overline{a}$, then for all $a, a' \in \overline{a}, |$*height*$(a) -$ *height*$(a')| \leqslant 1$.

(2) [The substructures $(A_p^+ \restriction P_j), j \leqslant k$]

(i) For all $j \leqslant k, (A_p^+ \restriction P_j) | E$ is the countable random graph.

(ii) Let $j < k$, $l < n_j$, and $\overline{a}$ be an l-tuple in $(A_p^+ \restriction P_j)$.
If $A_p^+ \models S_{jl} \overline{a}$, then $\overline{a}$ is an l-clique.
For all $m, m \leqslant n_j - l$, there is an m-tuple $\overline{a}'$ of elements in P_j such that $S_{j(l+m)}(\overline{a}, \overline{a}')$, iff, $S_{jl}\overline{a}$. [See Definition 2.2(vii).]

(iii) There is an n_1-tuple $\overline{a}$ in S_{1n_1} in A_p^+. For all $m \leqslant n_1$, the relations S_{1m} are empty in B_p^+. [See Definitions 2.2(vii) and 2.4(b).]

(3) [Arrowing]

(i) For all tuples $\overline{a}$ and $\overline{b}$, if $\overline{a}$ arrows $\overline{b}$, then $height(\overline{b}) = height(\overline{a}) + 1$.

(ii) For $i < k$, let $\overline{a}$ be an n_i-clique in P_i, b an element in P_{i+1}. Then, $R_i(\overline{a}, b)$, iff, there is an n_{i+1}-clique $\overline{b}, b \in \overline{b}$, such that $Ar(\overline{a}; \overline{b})$. [See Definition 2.2($viii$).]

(iii) For $i < k$, $\overline{a}$ an n_i-tuple, b an element, if $R_i(\overline{a}, b)$ and $S_{in_i}\overline{a}$, then b is in $P_{i+1} \wedge \neg S_{(i+1)1}$. [See Definition 2.4(a).]

(iv) For $i < k$, let $\overline{a}$ be an n_i-tuple in P_i, $\overline{b}$ an n_{i+1}-tuple in P_{i+1}. If $Ar(\overline{a}; \overline{b})$, then either $\neg S_{in_i}\overline{a}$ or every $b \in \overline{b}$ is in $\neg S_{(i+1)1}$. [See Definition 2.4(a).]

(v) For $i < k-1$ and each n_i-clique $\overline{a}$ of elements in P_i, there is an n_{i+1}-clique $\overline{b}$ in $\neg S_{(i+1)n_{i+1}}$, each $b \in \overline{b}$ in P_{i+1}, such that $\overline{a}$ arrows $\overline{b}$. Furthermore, there is an n_{i+1}-clique $\overline{c}$ in $S_{(i+1)n_{i+1}}$ such that $\overline{a}$ arrows $\overline{c}$, iff, $\overline{a}$ is in $\neg S_{in_i}$.
For $i = k-1$ and each n_i-clique $\overline{a}$ of elements in P_i, there is an n_k-clique $\overline{b}$ of elements in P_k such that $\overline{a}$ arrows $\overline{b}$, iff, $\overline{a}$ is in $\neg S_{in_i}$.

(4) (i) $A_p^+ \restriction \neg(P_1 \wedge S_{11}) \cong B_p^+$
(ii) $A_p^+ \restriction \neg P_1 \cong B_p^+ \restriction \neg P_1$

2.2 The sentences φ_p

We define, for each prefix p, a sentence $\varphi_p \in \mathrm{FO}(p)$, and prove that for the structures A_p and B_p defined above, $A_p \models \varphi_p$ and $B_p \not\models \varphi_p$. Recall that for each p with $alt(p) = 1$, φ_p has already been defined: $\varphi_\exists = \exists x Exx$ and for $n \geqslant 2$, $\varphi_{\exists^n} = \exists x_1 \ldots x_n Kl(x_1, \ldots, x_n)$.

We first explicitly define the sentences φ_p, for $alt(p) = 2$. We then give an inductive definition for prefixes p with $alt(p) \geqslant 3$. For $p = \exists^l \forall^m$, let $\varphi_p =$

$$\exists x_1 \ldots x_l \forall y_1 \ldots y_m \big(Kl(\overline{x}) \wedge \bigwedge_{i \leqslant l} \neg(Ey_1 x_i \wedge \neg Ex_i y_1) \wedge \neg Ar(\overline{x}; \overline{y})\big).$$

Now assume that $p = \exists^{n_1} \ldots s^{n_k}$ is a prefix with $alt(p) = k \geqslant 3$. We define a sequence of formulas $\theta_{k-1}, \theta_{k-2}, \ldots, \theta_1$; φ_p will be the sentence obtained by putting θ_1 into prenex normal form. Variables will be indexed so that x_{ij} is bound by the j^{th} quantifier in the i^{th} block: $\varphi_p = \exists x_{11} \ldots x_{1n_1} \forall x_{21} \ldots$, etc. We often use $\overline{x}_i$ for $x_{i1}, \ldots, x_{in_i}$.

For each $i \geqslant 2$, the free variables in θ_i will be exactly $x_{i1}, \ldots, x_{in_i}$. For k odd, we define $\theta_{k-1} = \exists x_{k1} \ldots x_{kn_k} Ar(\overline{x}_{k-1}; \overline{x}_k)$, $\overline{x}_i = (x_{i1}, \ldots, x_{in_i})$. For k even, $\theta_{k-1} = \forall x_{k1} \ldots x_{kn_k} \neg Ar(\overline{x}_{k-1}; \overline{x}_k)$.

Now suppose that $\theta_{k-1}, \ldots, \theta_i$ have been defined, and $i > 2$. For i odd, let $\theta_{i-1} = \exists \overline{x}_i (Ar(\overline{x}_{i-1}; \overline{x}_i) \wedge \theta_i(\overline{x}_i))$. For i even, let $\theta_{i-1} = \forall \overline{x}_i (Ar(\overline{x}_{i-1}; \overline{x}_i) \rightarrow \theta_i(\overline{x}_i))$.

Finally, let $\theta_1 =$

$$\exists \overline{x}_1 \Big(Kl(\overline{x}_1) \wedge \forall \overline{x}_2 \big(\bigwedge_{j \leqslant n_1} \neg(Ex_{21}x_{1j} \wedge \neg Ex_{1j}x_{21}) \wedge (Ar(\overline{x}_1; \overline{x}_2) \rightarrow \theta_2(\overline{x}_2))\big)\Big).$$

Observe that for all $i < k$, every quantifier in θ_i occurs positively, that is, in the scope of no negations. It is easy to see that, when θ_1 is put into prenex normal form, we obtain a sentence φ_p with $pr(\varphi_p) = p$.

The following idea is critical in the next two important lemmas. Though we consider the structures A_p and B_p, the description appeals constantly to the expanded structures A_p^+ and B_p^+. That is, because the universes of A_p and B_p are the same as those of A_p^+ and B_p^+, respectively, we can refer to the atomic σ_p-type of elements and tuples of the reduced structures. For example, we say that an element a in A_p is a P_i-element, or is in P_i, if $A_p^+ \models P_i a$. (We also say that an element $a \in A_p$ is an S-element if there is some $i \leqslant k$ such that $A_p^+ \models S_{i1} a$.) Homogeneous structures are particularly convenient to work with because the isomorphism type of a tuple is determined by its atomic type. In particular, Lemma 1.17 will be extremely helpful (both implicitly and explicitly).

Lemma 2.6 *For all prefixes p, $A_p \models \varphi_p$ and $B_p \not\models \varphi_p$.*

Proof. Suppose that $alt(p) = 2, p = \exists^l \forall^m$. φ_p says that there is an l-clique of elements, each having no E-predecessor, that does not arrow any m-tuple. Consider A_p. We claim that for any l-tuple $\overline{a}$ in A, if $\overline{a}$ is in S_{1n_1}, then it is such an l-clique. Such tuples exist, by Observation (2)(ii) and (iii). By Observation (1)(i), no a' in P_1, and hence no $a' \in \overline{a}$, has an E-predecessor. By Observations (3)(i) and (v), $\overline{a}$ arrows no m-tuple. Thus $A_p \models \varphi_p$, as desired. Next we show $B_p \not\models \varphi_p$. Suppose that $\overline{b}$ in B_p is an l-clique and no $b' \in \overline{b}$ has an E-predecessor. As above, $\overline{b}$ is in P_1 but, in contrast to A_p, $\overline{b}$ must be in $\neg S_{1n_1}$, also by Observation (2)(iii). By Observation (3)(v), there is an m-tuple $\overline{c}$ such that $\overline{b}$ arrows $\overline{c}$. Thus $B_p \not\models \varphi_p$.

We now suppose that $alt(p) = k \geqslant 3$. First we claim that for all even, resp. odd, $i, 2 \leqslant i \leqslant k-1$, and all n_i-tuples $\overline{a}$ in A_p such that $\overline{a}$ is an n_i-clique of P_i-elements, then $A_p \models \theta_i[\overline{a}]$ iff $\overline{a}$ is in $\neg S_{in_i}$, resp. $A_p \models \theta_i[\overline{a}]$ iff $\overline{a}$ is in S_{in_i}. Likewise for B_p (the argument is identical). We argue by downward induction on i, starting with $i = k-1$. If i is even, then k is odd, so $\theta_i = \exists \overline{x}_k Ar(\overline{x}_{k-1}; \overline{x}_k)$. The claim now follows immediately from Observation (3)(v). Likewise for i odd.

Assume now the claim holds for i, $i-1 \geqslant 2$, and i is odd. Recall that for i odd, $\theta_{i-1} = \exists \overline{x}_i (Ar(\overline{x}_{i-1}; \overline{x}_i) \wedge \theta_i(\overline{x}_i))$. Suppose that $\overline{a}$ is an n_{i-1}-clique of elements in P_{i-1}. By Observation (3)(i), for each n_{i+1}-clique $\overline{b}$, if $\overline{a}$ arrows $\overline{b}$ then $height(\overline{b}) = i$. Invoking the induction hypothesis, $A_p \models \theta_{i-1}[\overline{a}]$ iff there is an n_i-clique $\overline{b}$ of elements in P_i such that $\overline{a}$ arrows $\overline{b}$ and $\overline{b}$ is in S_{in_i}. By Observation (3)(v), $A_p \models \theta_{i-1}[\overline{a}]$ iff $\overline{a}$ is in $\neg S_{(i-1)n_{i-1}}$. The argument is similar for $i-1$ even.

It remains to show that $A_p \models \theta_1$ and $B_p \not\models \theta_1$ (recall $\theta_1 \equiv \varphi_p$). The sentence θ_1 says that there is an n_1-clique $\overline{a}$ of elements, each with no E-predecessor, such that for every n_2-tuple $\overline{b}$, if $\overline{a}$ arrows $\overline{b}$, then $\theta_2[\overline{b}]$. As in the case $alt(p) = 2$, we claim that any n_1-tuple $\overline{a}$ of P_1-elements in A_p such that $A_p \models S_{1n_1}\overline{a}$ witnesses that $A_p \models \theta_1$. We know already that such a tuple must form a clique, and that no $a \in \overline{a}$ has an E-predecessor. So it only remains to show that for any n_2-tuple $\overline{b}$, if $\overline{a}$ arrows $\overline{b}$, then $A_p \models \theta_2[\overline{b}]$. Suppose that $\overline{b}$ is an n_2-tuple such that $\overline{a}$ arrows $\overline{b}$. By the definition of arrowing, and Observation (3)(i), this implies that $\overline{b}$ is an n_2-clique of elements in P_2. Above, we showed that for any such tuple, $A_p \models \theta_2[\overline{b}]$ iff $\overline{b}$ is in $\neg S_{2n_2}$. Finally, note that Observation (3)(v) says that every n_2-tuple $\overline{b}$ arrowed by $\overline{a}$ is indeed in $\neg S_{2n_2}$, as desired.

To establish that $B_p \models \neg\theta_1$, using the above reasoning, it suffices to observe that every n_1-clique $\overline{a}$ of P_1-elements in B_p is in $\neg S_{1n_1}$ (by Observation (2)(iii)) and hence arrows an n_2-tuple $\overline{b}$ in S_{2n_2} (by Observation (3)(v)). Then $B_p \not\models \theta_2[\overline{b}]$, as shown above, so $B_p \models \neg\theta_1$. $\square$

2.3 The central argument

All that remains is prove the following lemma.

Lemma 2.7 *For each prefix p, $A_p \Rightarrow_{L_{\infty\omega}\{f(p)\}} B_p$.*

Proof. By Proposition 1.9, it suffices to show that the D. wins the $L_{\infty\omega}\{f(p)\}$-game from A_p to B_p. We again consider two cases, depending on whether or not $alt(p) = 1$. Without loss of generality, we assume that the S. always plays on a previously unpebbled element.

Case I $alt(p) = 1$

For $p = \exists$, recall that $f(\exists) = \forall^*$, $A_\exists$ is the countable graph with one loop, and $B_\exists$ is the countable graph with no edges at all. Observe that the $L_{\infty\omega}\{\forall^*\}$-game from $A_\exists$ to $B_\exists$ is an ordinary infinite Ehrenfeucht-Fraïssé game in which the S. may only play on $B_\exists$. The D.'s strategy is simply never to play on the element of $A_\exists$ with a loop. It is clear that she wins.

For $p = \exists^n, n \geqslant 2$, $f(p) = \forall^*\exists\ldots\exists\forall^*$, a word of length $2n-1$ containing n $\forall^*$s and $n-1$ $\exists$s. Thus, the $L_{\infty\omega}\{f(p)\}$-game from A_p to B_p is a $(2n-1)$-round game during which the S. is permitted to play $(n-1)$ pebbles on A_p. (In particular, he is not able to play, through the course of the entire game, n pebbles on an n-clique of S-elements in the relation S^{A_p}.) We claim that the D. can win by using the following strategy.

(i) In each round, play so as to maintain a partial isomorphism between the pebbled elements.

(ii) In each (even numbered) $\exists$ round, never play on an S-element of A_p.

Observe that this strategy implies that for all $m \leqslant n$, after round $2m-1$, there are at most $m-1$ pebbles on S-elements of A_p. We argue by induction on the length of the game (number of rounds and number of moves).

Each odd numbered round is an $\forall^*$ round, during which the S. may play pebbles on B_p an arbitary number of times. By the induction hypothesis, the active pebbles, on j-tuples $\overline{a}$ and $\overline{b}$, determine a partial isomorphism from A_p to B_p. Suppose that the S. now plays on some b' in B_p. Let $\psi(\overline{x}, x_{j+1})$ be the quantifier-free type of the $(j+1)$-tuple $(\overline{b}, b')$, that is, a maximally consistent quantifier-free formula such that $B_p \models \psi[\overline{b}, b']$. In order to maintain a partial isomorphism, it suffices for the D. to play on an element $a' \in A_p$ such that also $A_p \models \psi'[\overline{a}, a']$. Furthermore, to carry out part **(ii)** of the strategy described above, a' must also not be in S. Define $\psi'(\overline{x}, x_{j+1}) = \psi(\overline{x}, x_{j+1}) \wedge \neg S x_{j+1}$. Let C be the unique σ_p-structure of size $j+1$, with universe $(\overline{c}, c')$, such that $A\restriction\overline{a} \cong C\restriction\overline{c}$ and $C \models \psi'[\overline{c}, c']$. It is clear that C does not embed any of the constraints in $\mathcal{J}_p^A$ (e.g., C contains no n-clique because B_p does not). So by Lemma 1.17, there is an $a' \in A_p$, not in S, such that $A_p \models \psi'[\overline{a}, a']$. The D. places a pebble on this a'.

Each even numbered round is an $\exists$ round, during which the S. plays one pebble on A_p. By the induction hypothesis, because the D. never plays a pebble on A_p in S, there are less than n pebbles on A_p in S, so that the tuple $\overline{a}$ of elements in A_p on which there are pebbles cannot contain an n-clique. As B_p is the universal homogeneous K_n-free graph, it is clear that the D. can choose some $b' \in B_p$ so as to maintain a partial isomorphism. (To make the argument more explicit one can apply Lemma 1.17 as above.) Therefore, the D. does indeed have a winning strategy.

Case II $alt(p) \geqslant 2$

Let $p = \exists^{n_1} \ldots s^{n_k}, s \in \{\exists, \forall\}, alt(p) = k \geqslant 2$. Recall that $f(p) = f(\exists^{n_1}) * \ldots * f(s^{n_k})$ is a word of length $\sum_{i \leqslant k} 2n_i - 1$. It will be convenient to view the rounds of the $L_{\infty\omega}\{f(p)\}$-game as being divided into k *levels* by the subwords $f(s^{n_i}), s \in \{\exists, \forall\}$. Thus, the first $(2n_1 - 1)$ rounds are level 1, the next $(2n_2 - 1)$ rounds are level 2, etc. We write round $\langle j, m \rangle$ for the m^{th} round in level j. More precisely, $\langle j, m \rangle = (\sum_{i<j} 2n_i - 1) + m$.

We claim that the D. wins by playing according to the following strategy.

The D.'s strategy

(i)1. For all $a \in A_p, b \in B_p$, if there is a pair of pebbles on (a, b), then $height(a) = height(b)$.

Furthermore, let $\overline{a}$ in A_p, $\overline{b} \in B_p$, be corresponding m-tuples of pebbled elements, all of height $= i$, and $m \leqslant n_i$. Then either $S_{im}\overline{a}$ iff $S_{im}\overline{b}$ or, in case i is odd, resp. even, $S_{im}\overline{a}$ and $\neg S_{im}\overline{b}$, resp. $\neg S_{im}\overline{a}$ and $S_{im}\overline{b}$. Moreover, if $\overline{a}, \overline{b}$ were completely pebbled by the end of level j, and $j < i$, then in fact $S_{im}\overline{a}$ iff $S_{im}\overline{b}$.

(i)2. Suppose that round $\langle j, 2m \rangle$ or $\langle j, 2m+1 \rangle$ has just been completed, $j < k$ and odd, $0 \leqslant m < n_j$. Then for all $m', m < m' \leqslant n_j$, there is not a pair of correspondingly pebbled m'-tuples of P_j-elements $\overline{a}$ in A_p, $\overline{b}$ in B_p such that $A_p^+ \models S_{jm'}\overline{a}$ and $B_p^+ \models \neg S_{jm'}\overline{b}$.

Likewise, for $j < k$ and even, there is not a pair of correspondingly pebbled m'-tuples $\overline{a}$ in A_p, $\overline{b}$ in B_p such that $A_p^+ \models \neg S_{jm'}\overline{a}$ and $B_p^+ \models S_{jm'}\overline{b}$.

(ii) For each $j < k$, at the completion of level j the active pebbles, on $\overline{a}$ in A_p and $\overline{b}$ in B_p, induce a partial isomorphism between the $(\sigma_p \setminus \{\overline{S}_1, \ldots, \overline{S}_j\})$-reducts of A_p^+ and B_p^+. That is,

$$(A_p^+ \restriction \overline{a}) | (\sigma_p \setminus \{\overline{S}_1, \ldots, \overline{S}_j\}) \cong (B_p^+ \restriction \overline{b}) | (\sigma_p \setminus \{\overline{S}_1, \ldots, \overline{S}_j\}).$$

In particular, the D. plays so as to respect all the relations R_i.

(iii) The D.'s strategy during level k is somewhat different. We describe the case for k odd (k even is similar, essentially 'dual').

As in the case of previous levels, during round $\langle k, 1 \rangle$ (an $\forall^*$ round), the S. may perhaps be able to play so that there are (arbitrarily many) pairs of corresponding pebbles so that the A_p pebble is on an element in $P_{k-1} \wedge \neg S_{(k-1)1}$ and the B_p pebble is on an element in $P_{k-1} \wedge S_{(k-1)1}$. In particular, he may play so that there is an n_{k-1}-tuple of pebbles on an n_{k-1}-clique $\overline{b}$ in B_p of P_{k-1}-elements in $S_{(k-1)n_{k-1}}$ such that the corresponding pebbles are on an n_{k-1}-clique $\overline{a}$ in A_p of P_{k-1}-elements in $\neg S_{(k-1)n_{k-1}}$.

The significant difference between the above tuples $\overline{a}$ and $\overline{b}$ is the following. In A_p, there is an n_k-clique $\overline{a}'$ (in P_k) such that $\overline{a}$ arrows $\overline{a}'$, but there is no such matching n_k-clique in B_p (by Observation (3)(iv) again). Thus the S. would win if he was allowed to play n_k more pebbles on A_p. Fortunately, he may only play $n_k - 1$ more. Further, the S. can, in round $\langle k, 2\rangle$, and subsequent rounds $\langle k, 2m\rangle, m < n_k$, force the D. to 'break' R_{k-1}-types. That is, if he plays a pebble on some $a_0 \in A_p$ such that $R_{k-1}(\overline{a}, a_0)$, then there is no b_0 in B_p such that $R_{k-1}(\overline{b}, b_0)$ (by Observation (3)(iii)).

We now introduce some terminology. Call a pair of correspondingly pebbled n_{k-1}-tuples of P_{k-1}-elements $\overline{a}'$ in A_p, $\overline{b}'$ in B_p, *switched* just in case $\neg S_{(k-1)n_{k-1}}\overline{a}'$ and $S_{(k-1)n_{k-1}}\overline{b}'$. We will also say that each of the tuples $\overline{a}'$ and $\overline{b}'$ is itself switched. A pair (a, b) of correspondingly pebbled P_k-elements in A_p and B_p is *distinguished* if there are correspondingly pebbled switched n_{k-1}-tuples $\overline{a}', \overline{b}'$, such that $\neg S_{(k-1)n_{k-1}}\overline{a}' \wedge R_{k-1}(\overline{a}', a)$ and $S_{(k-1)n_{k-1}}\overline{b}' \wedge \neg R_{k-1}(\overline{b}', b)$.

We are now prepared to describe this part of the strategy.

($\exists$) **round.** In each ($\exists$) round $\langle k, 2m\rangle, 1 \leqslant m < n_k$, we modify the D.'s strategy so that it no longer requires that she always respects the relation R_{k-1}. (But she will still play so as to preserve a partial isomorphism on the $\{E, P_1, \ldots, P_k, R_1, \ldots, R_{k-2}\}$-reducts of the pebbled parts of A_p^+ and B_p^+.) In any such ($\exists$) round, if the S. plays on an element a' in P_k in A_p, then the D. plays on some b' in P_k in B_p in accordance with the following restrictions.

(a) She maintains the partial $\{E\}$-isomorphism between pebbled elements.

(b) If $\overline{a}', \overline{b}'$ are correspondingly pebbled tuples of P_{k-1}-elements that are *not* switched, then $R_{k-1}(\overline{a}', a')$ iff $R_{k-1}(\overline{b}', b')$.

(Observe that if $\overline{a}', \overline{b}'$ are switched, then necessarily $\neg R_{k-1}(\overline{b}', b')$, so the D. 'has no choice' here.) (b) is thus a weakened condition on 'respecting' R_{k-1}.

($\forall^*$) **round.** On the other hand, in each ($\forall^*$) round $\langle k, 2m+1\rangle, 0 \leqslant m < n_k$, in which the S. plays a pebble in P_k, the D. still plays so as to 'respect' R_{k-1}. More precisely, since the mapping between previously pebbled elements may not 'preserve' R_{k-1}, this means the following. Let $b' \in B_p$, $a' \in A_p$ be the P_k-elements on which the S. and the D. have played, respectively, in this round. Then for any pair $\overline{a}, \overline{b}$, in A_p, B_p, of $(n_{k-1}+1)$-tuples of correspondingly pebbled elements, such that $a' \in \overline{a}$ and $b' \in \overline{b}$, $R_{k-1}\overline{a}$ iff $R_{k-1}\overline{b}$.

Observe that the pair (a', b') are not distinguished, as defined above.

Outside of P_k. Finally, if the S. plays on some element not in P_k, then the D. plays as described above in parts **(i)** and **(ii)**. In particular, she preserves partial $(\{E\} \cup \overline{P} \cup \{R_1, \ldots, R_{k-2}\})$-isomorphism between the pebbled elements. Again, regarding R_{k-1}, she plays so as to respect R_{k-1} in the weaker sense explained above.

This completes the description of the strategy.

The D. does win if she can successfully carry out this strategy, since it entails that at the completion of level k, the end of the game, the pebbles determine a partial isomorphism between A_p and B_p, the E-reducts of A_p^+ and B_p^+, as desired. It remains to show that this is possible.

Proving that the D. wins.

We argue by induction on the number of rounds. The induction hypothesis is that the D. has, through round r, been able to carry out the above strategy successfully.

Before giving the argument, we make a general observation. By Observation (1)(ii), if $T \in \sigma_p$, $\overline{a} \in A_p$ or B_p, and $T\overline{a}$, then for all $a, a' \in \overline{a}, |height(a) - height(a')| \leqslant 1$. This implies the following point, which simplifies the D.'s choice of moves. Suppose that the S. has just played a pebble on an element of height $i, i \leqslant k$. By Condition **(i)**, the D. should play on an element of the same height in the other structure. We claim that, in making her choice, she only needs to consider the pebbles that are on elements of height $i-1, i$, or $i+1$. This is because the strategy is defined in terms of preserving partial isomorphisms between reducts of the σ_p-structures A_p^+ and B_p^+, and no tuple in any σ_p-relation contains an element of height i and one of height j, when $|j-i| > 1$.

Part I — Round r occurs in level j, $j < k$.

Case A Suppose the S. plays a pebble in P_i, where i is $< j$ and even. (The argument for i odd is similar.)
By the induction hypothesis, the D. has maintained Conditions **(i)**(1.) and **(ii)** of the strategy. By the preceding note, we can restrict our attention to (the pebbled elements in) the substructures of A_p and B_p with universe $P_{i-1} \vee P_i \vee P_{i+1}$. Let $\overline{a}, \overline{b}$ denote the (corresponding) tuples of currently pebbled elements in $A_p \upharpoonright (P_{i-1} \vee P_i \vee P_{i+1})$ and $B_p \upharpoonright (P_{i-1} \vee P_i \vee P_{i+1})$, respectively. We know that $\overline{a}$ and $\overline{b}$ realize the same $(\{E\} \cup \overline{P} \cup \overline{R})$-type. Also, for each $j \in \{i-1, 1, i+1\}$ and each pair of m-tuples $\overline{a}', \overline{b}'$, $m \leqslant n_j$, of correspondingly pebbled elements in A_p, B_p, such that every element in $\overline{a}'$ and in $\overline{b}'$ has height j, if $\overline{a}'$ and $\overline{b}'$ do not realize the same $\overline{S}$-type, then either $j = i$ and $\neg S_{im}\overline{a}'$ and $S_{im}\overline{b}'$; or $|j - i| = 1$ and $S_{jm}\overline{a}'$ and $\neg S_{im}\overline{b}'$.

There are various kinds of moves the S. can make.

A(i) The S. plays on an element b' in S_{i1} in B_p. We must show that the D. can choose an element a' in A_p while respecting the strategy described above. To this end, we use Lemma 1.17 to show that there is an a' in $P_i \wedge \neg S_{i1}$ such that $(B_p^+ \upharpoonright (\overline{b}, b'))|(\sigma_p \setminus \overline{S}) \cong (A_p^+ \upharpoonright (\overline{a}, a'))|(\sigma_p \setminus \overline{S})$. It is easy to see that this shows that the D. can make a move that respects the conditions of her strategy.

Let $\theta(\overline{x}, x')$ be the complete atomic $(\sigma_p \setminus \overline{S})$-type of $(\overline{b}, b')$, and let $\psi(\overline{x}, x')$ be a conjunction of every negated atomic formula of the form $S_{lm}\overline{y}$, $\overline{y} \subseteq \overline{x} \cup \{x'\}$ and $x' \in \overline{y}$. We define $\theta'(\overline{x}, x') = \theta(\overline{x}, x') \wedge \psi(\overline{x}, x')$. Then there is a unique σ_p-structure C, with universe $(\overline{c}, c')$, such that $C \upharpoonright \overline{c} \cong A_p \upharpoonright \overline{a}$ and $C \models \theta'[\overline{c}, c']$. Observe that there is a natural $(\sigma_p \setminus \overline{S})$-isomorphism from $(B_p^+ \upharpoonright (\overline{b}, b'))|(\sigma_p \setminus \overline{S})$ to $C|(\sigma_p \setminus \overline{S})$, taking each $b_l \in \overline{b}$ to the corresponding $c_l \in \overline{c}$ and b' to c'. By Lemma 1.17, it suffices to show that no constraint in $\mathcal{J}_p^A$ is isomorphic to a substructure of C.

Say that a constraint D in $\mathcal{J}_p^A \cup \mathcal{J}_p^B$ is $\overline{S}$-*symmetric* just in case for all σ_p-structures D', if $D|(\sigma_p \setminus \overline{S}) \cong D'|(\sigma_p \setminus \overline{S})$, then also D' in $\mathcal{J}_p^A$ *and* in $\mathcal{J}_p^B$. In particular, it is clear from the definitions of $\mathcal{J}_p^A$ and $\mathcal{J}_p^B$ that the only constraints that are not $\overline{S}$-symmetric are the structures $M_l^p, l < k$, and N^p from Definition 2.4.

Suppose, for the sake of contradiction, that there is a substructure $D \subseteq C$ that is isomorphic to a constraint in $\mathcal{J}_p^A$. We claim that D cannot be $\overline{S}$-symmetric. As $B_p^+ {\restriction} (\overline{b}, b')|(\sigma_p \setminus \overline{S}) \cong C|(\sigma_p \setminus \overline{S})$, there is a $D' \subseteq B_p^+ {\restriction} (\overline{b}, b')$, with $D|(\sigma_p \setminus \overline{S}) \cong D'|(\sigma_p \setminus \overline{S})$. Since $D' \subseteq B_p$, this implies that $D' \notin \mathcal{J}_p^B$, so that D is not $\overline{S}$-symmetric, as desired.

Consequently, if there is a substructure $D \subseteq C$ isomorphic to a constraint in $\mathcal{J}_p^A$ it must isomorphic to one of the structures M_l^p from Definition 2.4. Furthermore, this substructure must include the element c', since (the universe of C is $\overline{c} \cup \{c'\}$ and) $C {\restriction} \overline{c}$ is isomorphic to a substructure of A_P^+, so it cannot embed any such constraint. But this is not possible, as every element in any of the M_l^p, $l < k$, is an S-element (that is, in some relation $S_{j1}, 1 \leqslant j \leqslant k$), while c' is not. So we have shown that no substructure of C is isomorphic to a constraint of $\mathcal{J}_p^A$, as desired.

A(ii) The S. plays on an element a' in S_{i1} in A_p. We argue as above, using Lemma 1.17 to show that the D. can choose a b' in $P_i \wedge S_{i1}$ so as to satisfy Conditions **(i)** and **(ii)** of her strategy. Recall that $\overline{a}, \overline{b}$ are the (corresponding) tuples of currently pebbled elements in $A_p {\restriction} (P_{i-1} \vee P_i \vee P_{i+1})$ and $B_p {\restriction} (P_{i-1} \vee P_i \vee P_{i+1})$, respectively. Let g be the natural bijection from $\overline{a}$ to $\overline{b}$ that takes a pebbled element in $\overline{a}$ to the element in $\overline{b}$ on which the corresponding pebble is located.

Let $\theta(\overline{x}, x')$ be the complete atomic $(\sigma_p \setminus \overline{S})$-type of $(\overline{a}, a')$. Define $\psi_1(\overline{x}, x')$ to be the conjunction of (*positive* !) atomic $\overline{S}_i$ formulas satisfied by $(\overline{a}, a')$, and $\psi_2(\overline{x})$ to be the conjunction of atomic $\overline{S}$ formulas satisfied by $\overline{b}$. Finally, we let $\theta'(\overline{x}, x')$ be the unique σ_p-type extending $\theta \wedge \psi_1 \wedge \psi_2$ such that every atomic formula that occurs as a conjunct of θ' is a conjunct of θ, ψ_1, or ψ_2. One can easily check that it suffices for the D. to pebble an element $b' \in B_p$ such that $B_p \models \theta'[\overline{b}, b']$.

Again, there is a unique σ_p-structure C, with universe $(\overline{c}, c')$, such that $C {\restriction} \overline{c} \cong B_p {\restriction} \overline{b}$ and $C \models \theta'[\overline{c}, c']$. There is also a natural isomorphism h from $A_p^+ {\restriction} (\overline{a}, a')|(\sigma_p \setminus \overline{S})$ to $C|(\sigma_p \setminus \overline{S})$. By Lemma 1.17, it now suffices to show that no constraint in $\mathcal{J}_p^B$ embeds in C. If some constraint in $\mathcal{J}_p^B$ is isomorphic to a substructure of C, then, as in **A(i)**, it is not $\overline{S}$-symmetric and can only be one of the structures M_l^p. In fact, it must be either M_i^p or M_{i-1}^p, since these are the only such structures with an element in P_i. Here the element c' is in S_i, so we cannot argue quite as we did in **A(i)**.

We show that M_i^p does not embed in C. Suppose for contradiction that $D \subseteq C$ and $D \cong M_i^p$. We will show that this implies that M_i^p is also a substructure of $A_p {\restriction} (\overline{a}, a')$, which is not possible. Let $\{d_1, \ldots, d_{n_i+1}\}$ be the universe of D, $\overline{d} = (d_1, \ldots, d_{n_i})$. We can assume that d_{n_i+1} is in $S_{(i+1)1}$, so every $d' \in \overline{d}$ is in P_i and $S_{in_i}\overline{d}$ (because $D \cong M_i^p$). Let $a_l = h^{-1}(d_{n_i+1})$, for some $a_l \in \overline{a}$, and let $b_l = g(a_l)$. Observe that $a' \in h^{-1}(\overline{d})$ (where $h^{-1}(\overline{d}) = (h^{-1}(d_1), \ldots, h^{-1}(d_{n_i}))$.

First, we claim that $a_l \in S_{(i+1)1}$. Otherwise, if $a_l \notin S_{(i+1)1}$, then also $b_l \notin S_{(i+1)1}$, by Condition **(i)** of the strategy. Then clearly $\neg S_{(i+1)1} x_l$ is a conjunct of θ', which contradicts the fact that d_{n_i+1} is in $S_{(i+1)1}$. Furthermore $A_p \models R_i(h^{-1}(\overline{d}), a_l)$, because $C \models R_i(\overline{d}, d_{n_i+1})$ and h is a $(\sigma_p \setminus \overline{S})$-isomorphism from $(A_p {\restriction} (\overline{a}, a'))|(\sigma_p \setminus \overline{S})$ to $C|(\sigma_p \setminus \overline{S})$.

Finally, we want to show that in fact $A_p \models S_{in_i}(h^{-1}(\overline{d}))$. It is then easy to see that this implies that the substructure of A_p with universe $\{h^{-1}(d) \mid d \in D\}$ is isomorphic to M_i^p, yielding the desired contradiction. Because $C \models S_{in_i}\overline{d}$, there is a conjunct $S_{in_i}\overline{y}$

of θ', $\overline{y} \subseteq \overline{x} \cup \{x'\}$, with $x' \in \overline{y}$ (because a' must be in $h^{-1}(\overline{d})$). Clearly $S_{in_i}\overline{y}$ cannot be a conjunct of ψ_2, since x' does not occur in this formula, so we can conclude that $S_{in_i}\overline{y} \in \psi_1$ and, thus, that $A_p \models S_{in_i}(h^{-1}(\overline{d}))$.

The argument to show that M_{i-1}^p does not embed in C is a simple variant of the previous one, and straightforward.

There are a number of other cases that we do not treat explicitly, though we note the following facts. One, if the S. plays instead in $P_i \wedge \neg S_{i1}$, in either A_p or B_p, then it is easy to show that the D. can maintain her strategy by also playing on an element in $P_i \wedge \neg S_{i1}$. Two, the argument for i odd is essentially identical. When $i = 1$, things look slightly different, since the formula $P_1 x \wedge S_{11} x$ is not satisfied in B_p, but it is easy to see that this makes no real difference: whenever the S. plays on an element in $P_1 \wedge S_{11}$ in A_p, the D.'s strategy allows her to choose an element in $P_1 \wedge \neg S_{11}$ in B_p in response.

Case B S. plays in P_j.
Assume that j is even (odd is identical). We have the D. play precisely as she did in the preceding Case A, so it only remains to show that in doing so she also respects Condition **(ii)** of the strategy. It suffices to establish the following claim.

Claim 2.8 *Suppose that after round $\langle j, 2m+1\rangle, 0 \leqslant m \leqslant n_j - 1$, there are l-tuples, $l \leqslant n_j$, $\overline{a}$ in A_p, $\overline{b} \in B_p$, of correspondingly pebbled P_j-elements such that $\neg S_{jl}\overline{a}$ and $S_{jl}\overline{b}$. Then every $b \in \overline{b}$ was pebbled by the S.. In particular, $l \leqslant m$.*

To verify the claim, we examine the above construction in Case A. If the S. plays a pebble in $\neg S_{j1}$, in either structure, then the D. does so too, and neither element can be in any tuple that is in any relation $S_{jl}, l \leqslant k$. It only remains to consider the case when the S. plays a pebble on an element a' in S_{j1}. The D. then also plays on some $b' \in S_{j1}$, as described above. Suppose that there are correspondingly pebbled l-tuples $\overline{a}$ in A_p, $\overline{b} \in B_p$, of P_j-elements such that $\neg S_{jl}\overline{a}$ and $S_{jl}\overline{b}$. By the properties of $\overline{S}_j$, we know that $\neg S_{j(l+1)}(\overline{a}, a')$. So we must show that also $\neg S_{j(l+1)}(\overline{b}, b')$. But this is clear from the definition of $\theta'(\overline{x}, x')$ in Case A(ii).

Case C S. plays in P_i, and $i > j$.
This case is rather straightforward, and we give a less formal and more intuitive argument. Given that the D. is committed to following the strategy described above, we consider how it is possible for the S. to 'force' the D. to play a pebble so that there is a pair a, b of correspondingly pebbled elements in A_p, B_p such that $a \in S_{i1}$ iff $b \notin S_{i1}$, for some $i < k$. If the S. plays in P_1, this is easy, since there are no elements in $P_1 \wedge S_{11}$ in B_p. He simply plays on an element $a \in A_p$ in $P_1 \wedge S_{11}$. On the other hand, in order to accomplish this for elements of height $i+1 > 1$ (assume $i+1$ even), there must be a pair of correspondingly pebbled n_i-tuples $\overline{a}, \overline{b}$ in A_p, B_p such that $S_{in_i}\overline{a}$ and $\neg S_{in_i}\overline{b}$. In this case, the S. can pebble a $b' \in B_p$ such that $S_{(i+1)1} b \wedge R_i(\overline{b}, b')$. In A_p, there is no a' such that $S_{(i+1)1} a \wedge R_i(\overline{a}, a')$ (see Observation (3)(iii)), so the D. will choose some a' such that $\neg S_{(i+1)1} a \wedge R_i(\overline{a}, a')$. But by Conditions **(i)** and **(ii)** of the D.'s strategy (and Case B above), at no point during level $j \leqslant i$ of the game is there a pair of correspondingly pebbled n_i-tuples $\overline{a}, \overline{b}$ in A_p, B_p such that $S_{in_i}\overline{a}$ and $\neg S_{in_i}\overline{b}$. Thus, the D. cannot be compelled to 'break' $\overline{S}_i$-types, $i > j$, during level j of the game.

Part II – Round r occurs in level k.

This part is also divided into a number of cases. We assume, w.l.o.g., that k is odd. Note that Condition **(iii)** of the strategy ensures the following property of a play of the game. For each $m, 0 \leqslant m \leqslant n_k - 1$, after rounds $\langle k, 2m\rangle$ and $\langle k, 2m+1\rangle$, there are at most m distinguished pairs (a, b) of pebbled P_k-elements in A_p and B_p. (Recall the definition of a distinguished pair given in the description of Condition **(iii)** of the D.'s strategy.)

Let $\overline{a} = (a_0, \dots, a_t), \overline{b} = (b_0, \dots, b_t)$, be the (corresponding) tuples of currently pebbled elements in $A_p{\restriction}(P_{k-1} \vee P_k)$ and $B_p{\restriction}(P_{k-1} \vee P_k)$, respectively. Let g be the bijective mapping from $\overline{a}$ to $\overline{b}$ that takes a pebbled element $a_s, s \leqslant t$, in $\overline{a}$ to the element $b_s, s \leqslant t$, in $\overline{b}$ on which the corresponding pebble is located.

Case A S. plays in P_k.

A(i) S. plays on an element a' in A_p in an $(\exists)$ round $\langle k, 2m\rangle, m \leqslant n_k - 1$.
By the induction hypothesis, there are less than $n_k - 1$ distinguished pairs of pebbled elements in $\overline{a}, \overline{b}$ at the end of round $\langle k, 2m\rangle$. In particular, $\overline{a} \cup \{a'\}$ does not contain a switched n_{k-1}-tuple $\overline{a}_0$ and an n_k-tuple $\overline{a}_1$, each $a \in \overline{a}_1$ distinguished, such that $\overline{a}_0$ arrows $\overline{a}_1$. (This is the crucial point.)

Let $\theta(\overline{x}, x')$ be the complete atomic $\{E\} \cup \overline{P}$-type of $(\overline{a}, a')$ in A_p^+. Let h be the map that takes x' to a' and each $x_s \in \overline{x}$, $s \leqslant t$, to the naturally corresponding element $a_s \in \overline{a}$ which 'instantiates it'. Let $\psi_1(\overline{x}, x')$ be the conjunction of all atomic R_{k-1}-formulas $R_{k-1}(\overline{y}, x')$, $\overline{y} \subseteq \overline{x}$ such that the $n_{k-1}+1$-tuple $(\overline{a}_0, a')$ of A_p elements instantiated by $(\overline{y}, x')$, $(= h(\overline{y}) \cup \{a'\})$, satisfies $R_{k-1}(\overline{a}_0, \overline{a})$, and $\overline{a}_0$ is not a switched tuple. Let $\psi_2(\overline{x})$ be the complete σ_p-type of $\overline{b}$ in B_p. Finally, let $\theta'(\overline{x}, x')$ be the unique complete σ_p-type extending $\theta \wedge \psi_1 \wedge \psi_2 \wedge S_k x'$ such that every conjunct of θ' that is an atomic formula either is $S_k x'$ or occurs as a conjunct in θ, ψ_1, or ψ_2. (It is easy to check that this is well defined.)

It is easy to see that the D. will satisfy the conditions of her strategy if she can choose an element b' in B_p such that $\theta'[\overline{b}, b']$. Again, it suffices to show that the unique σ_p-structure C with universe $(\overline{c}, c')$ satisfying $\theta'[\overline{c}, c']$ does not embed any constraint in $\mathcal{J}_p^B$. By previous ideas, the only constraints for which it is non-trivial to verify this are M_{k-1}^p and the structure with universe $(\overline{d}, \overline{d}')$ such that $\overline{d}$ is a tuple of P_{k-1}-elements, $\overline{d}'$ is a tuple of P_k-elements, and $Ar(\overline{d}; \overline{d}')$. The argument that neither structure embeds in C is by now straightforward.

A(ii) S. plays on an element b' in B_p in an $(\forall^*)$ round.
Let $\theta(\overline{x}, x')$ be the complete atomic $\{E\} \cup \overline{P}$-type of $(\overline{b}, b')$ in B_p. Let $\psi_1(\overline{x}, x')$ be the conjunction of all atomic R_{k-1}-formulas $R_{k-1}(\overline{y}, x')$, $\overline{y} \subseteq \overline{x}$ such that the $(n_{k-1}+1)$-tuple $(\overline{b}_0, b')$ of B_p elements instantiated by $(\overline{y}, x')$ satisfies $R_{k-1}(\overline{b}_0, \overline{b})$. Let $\psi_2(\overline{x})$ be the complete σ_p-type of $\overline{a}$ in A_p. As before, we define $\theta'(\overline{x}, x')$ be the unique complete σ_p-type extending $\theta \wedge \psi_1 \wedge \psi_2$ such that every conjunct of θ' that is an atomic formula occurs as a conjunct in θ, ψ_1, or ψ_2.

The argument proceeds as in the previous cases. Once again, it suffices for the D. to choose an element a' in A_p such that $A_P^+ \models \theta'[\overline{a}, a']$. It is easy to show that this is possible. We only make the following observation. Suppose that $\overline{a}', \overline{b}'$ is a switched pair of corresponding n_{k-1}-tuples of pebbled P_{k-1}-elements; by definition, $\neg S_{(k-1)n_k}\overline{a}'$ and $S_{(k-1)n_k}\overline{b}'$. Recall that the relevant difference between $\overline{a}'$ and $\overline{b}'$ is that there is no b_0 in B_p such that $R_{k-1}(\overline{b}', b_0)$, though there is an a_0 in A_p, $R_{k-1}(\overline{a}', a_0)$. Informally, this means that, when the S. plays on b' in P_k, he in fact has a more limited choice of moves than the D., so the D. will have no trouble adequately answering him. (Again observe that (a', b') will not be a distinguished pair — this explains the bound on the number of distinguished pairs mentioned at the beginning of Part II.)

Case B S. plays in P_{k-1}.
The argument is straightforward, following the pattern of earlier cases. We only note that it is easy to check that if the S. does play in P_{k-1}, he cannot, by doing so, transform a previously pebbled non-distinguished pair of elements $a \in A_p, b \in B_p$, into a distinguished pair. (This is a consequence of the D.'s strategy of 'respecting' R_{k-1} when the S. plays in P_{k-1}.)

Case C S. plays in $P_j, j \leqslant k-2$.

In this case, the argument is exactly that of Part I.

This complete the proof of the theorem. □

References

[1] E. Grädel and G. McColm. Hierarchies in transitive closure logic, stratified Datalog and infinitary logic. *Annals of Pure and Applied Logic*, 77:166–199, 1996.

[2] W. Hodges. *Model Theory.* Cambridge University Press, 1993.

[3] H. J. Keisler and W. Walkoe, Jr. The diversity of quantifier prefixes. *Journal of Symbolic Logic*, 38:79–85, 1973.

[4] W. Walkoe, Jr. Finite partially-ordered quantification. *Journal of Symbolic Logic*, 35:535–555, 1970.