Arity Hierarchies

Martin Grohe

Abstract

Many logics considered in finite model theory have a natural notion of an arity. The purpose of this article is to study the hierarchies which are formed by the fragments of such logics whose formulae are of bounded arity.

Based on a construction of finite graphs with a certain property of homogeneity, we develop a method that allows us to prove that the arity hierarchies are strict for several logics, including fixed–point logics, transitive cloure logic and its deterministic version, variants of the database language Datalog, and extensions of first–order logic by implicit definitions.

Furthermore, we show that all our results already hold on the class of finite graphs.

Contents

Appeared in *Annals of Pure and Applied Logic*, 82:103–163, 1996.

1 Introduction

For some well–known reasons, first–order logic loses the dominating role it has always played in classical model theory when it comes to the study of finite models. Instead, various extensions of first–order logic are investigated with the same rights. In many of such extensions, we can naturally associate a non–negative integer, the *arity*, with each formula. For example, the arity of a second–order formula is the maximal arity of a relation variable occuring in this formula.

The restriction of such a logic $\mathbf{L}$ to formulae of arity $\leq k$ forms the *k–ary fragment* $\mathbf{L}^k$ of $\mathbf{L}$. We call the sequence $(\mathbf{L}^k)_{k\geq 0}$ the *arity hierarchy of* $\mathbf{L}$. This article is concerned with the question whether this hierarchy is strict for several logics playing prominent roles in finite model theory.

This question has received a certain amount of attention throughout the development of finite model theory. For least fixed–point logic, it was posed by Chandra and Harel [CH82] in 1982. After some partial solutions by Gaifman (unpublished) and Dublish and Maheshwari [DM89], it was solved in [Gro93].

The arity problem for variants of transitive closure logic was investigated in [CM92], as a first step the unary was separated from the binary fragment there. Whereas the results of [Gro93] imply the hierarchy of ordinary transitive closure logic being strict, the analogous statement for deterministic transitive closure logic could not be proved by the methods used there.

In [Kol90], Kolaitis conjectured that the arity hierachy was strict for implicitly definable queries.

Arity questions have also been studied in database theory. Afrati and Cosmadakis [AC89] showed that the hierarchy is strict for the database language (pure) datalog.

Although we are not dealing with second–order logic, let us mention a theorem of Ajtai [Ajt83] which implies that the arity hierarchies of Σ^1_1 and Π^1_1 are strict. Furthermore, generalized quantifiers also have a natural notion of an arity. See [Hel92, GH95] for arity results for extensions of first–order logic by fixed–point operators and generalized quantifiers.

So some results have been obtained so far, but all of them have a severe shortcoming — the proofs do not work with a uniform signature. More precisely, to separate $\mathbf{L}^k$ from $\mathbf{L}^{k-1}$ they use queries of a signature which contains a k–ary or $2k$–ary relation symbol. The main contribution of this article is to develop a method which allows to establish arity hierarchy results on finite structures of a fixed signature, actually on the class of finite graphs.

The main (technical) result reads:

1.1 Main Lemma *For each $k \geq 1$, there is a formula $\textsc{Edge}_k(\overset{k}{x}, \overset{k}{y})$ (which is a conjunction of atomic formulae in the signature $\{E\}$ of graphs), such that the transitive closure of $\textsc{Edge}_k(\overset{k}{x}, \overset{k}{y})$ cannot be defined in the $(k-1)$–ary fragment $\mathbf{s}$–$\mathbf{PFP}^{k-1}$ of simultaneous partial fixed–point logic on the class of finite graphs.*

Let us explain this first: Simultaneous partial fixed–point logic is a variant of partial fixed–point logic allowing simultaneous inductions. We use it here because it is the most expressive of the common fixed–point logics, thus our lemma becomes most general. However, the reader who is not familiar with this logic may replace $\mathbf{s\text{–}PFP}^{k-1}$, for example, by the $(k-1)$–ary fragment of least fixed–point logic (and the statement remains true).

The *transitive closure* of a $2k$–ary relation R on a set A is the $2k$–ary relation defined by

$$\mathrm{TC}(R) = \{\overset{k}{a}\overset{k}{b} \mid \exists n \geq 2, \overset{k}{a}_1 = \overset{k}{a}, \overset{k}{a}_2, \ldots, \overset{k}{a}_n = \overset{k}{b} \forall i < n : \ R\overset{k}{a}_i\overset{k}{a}_{i+1}\}.$$

The *transitive cloure* of a formula $\varphi(\overset{k}{x}, \overset{k}{y})$ with $2k$ free variables $x_1, \ldots, x_k,\ y_1, \ldots, y_k$ is the $2k$–ary query which associates $\mathrm{TC}(\{\overset{k}{a}\overset{k}{b} \mid \mathfrak{A} \models \varphi[\overset{k}{a}, \overset{k}{b}]\})$ with a structure $\mathfrak{A}$.

So essentially, the main lemma says that we cannot tell in $\mathbf{s\text{–}PFP}^{k-1}$ whether or not there exists an $\textsc{Edge}_k$–path in a graph $\mathfrak{G}$ from a tuple $\overset{k}{c}$ of vertices to a tuple $\overset{k}{d}$.

However, the transitive closure of $\textsc{Edge}_k$ can be defined in the k–ary fragment of all common fixed–point logics and also in the k–ary fragment of transitive closure logic. Since the $(k-1)$–ary fragments of these logics are contained in $\mathbf{s\text{–}PFP}^{k-1}$, the main lemma implies that their arity hierarchies are strict (on the class of graphs) (cf. Corollary 3.1).

By slightly modifying the query we consider, we can also prove that the arity hierarchies of pure Datalog and several extensions are strict (cf. Corollary 4.2), and establish Kolaitis' conjecture that the hierarchy is strict for implicitly defined queries (cf. Theorem 5.3).

Unfortunately, for reasons explained in Subsection 6.1, it turns out that slight modifications do not suffice to prove that the arity hierarchy of deterministic transitive closure logic is strict. We have to improve our technique considerably to eventually obtain this result (cf. Theorem 6.7).

The proof of the main lemma proceeds as follows: We first construct certain families of structures which are somehow "homogeneous in the $(k-1)$–tuples" (thus $(k-1)$–ary fixed–point formulae become ineffective). To achieve this we use a method that was introduced by Hrushovski [Hru92] (in a completely different context) to extend partial isomorphisms of finite graphs.

Then we establish the non–expressibility of our query by $(k-1)$–ary simultaneous fixed–point formulae on these structures employing an Ehrenfeucht-Fraïssé game (which is based on the well–known pebble game for infinitary logic with only finitely many variables).

The article is organized as follows: The construction of the structures mentioned above will take place in Section 2. We first introduce Hrushovski's general method and then apply it to our situation. The crucial properties of the resulting family $\mathfrak{G}(k,n)$ $(n, k \geq 1)$ of graphs are described in Theorem 2.17.

Fixed–point logics are introduced in Section 3, and after defining a suitable game and proving its correctness in Subsection 3.1 we can complete the proof of the main lemma in Subsection 3.2.

The main problem in extending the results to Datalog (introduced in Section 4) is to avoid the use of constants (which would lead us out of the pure graph signature). In Subsection 4.1

we handle this by extending our graphs in a way that all constants needed are definable by a positive existential first–order formula (and thus by a Datalog program).

Kolaitis' problem on arities of implicit definitions will be considered in Section 5. Again, we have to extend our structures slightly to prove that we have a strict arity hierarchy here.

As we have mentioned, the situation is different for deterministic transitive closure logic. Actually, in Subsection 6.1 we will see that deterministic transitive closure logic is already contained in the binary fragment of inductive fixed-point logic. Thus the main lemma cannot just be extended taking deterministic transitive closure instead of transitive closure. However, we can show that, for any $k \geq 1$, the k–ary fragment of deterministic transitive closure logic is not contained in the $(k-1)$–ary fragment of transitive closure logic, again on the class of finite graphs. The proof of this result (carried out in the rest of Section 6) requires to go back to Section 2 and employ Hrushovski's method again to build a new class of graphs. The technical difficulty occurring here is that one application of the construction does not suffice, we have to iterate it to be able to handle nested transitive closure operators.

In the concluding Section 7 we shortly suggest some directions for further research on arity problems.

Since the logics we consider are mainly of interest in finite model theory, let us agree that *all structures and all signatures we are going to consider are finite.* But note that, since our results separate logics, they remain true in the infinite (but may be easier to prove there). An exception is the hierarchy theorem for implicitly definable queries (because of Beth's theorem).

Our notation is standard, as it can be found for instance in [EFT94]. The universe of a structure $\mathfrak{A}$ is always denoted by A. We often abbreviate tuples $x_1 \dots x_k$ by $\overset{k}{x}$ or $\bar{x}$. For convenience, we assume that signatures do not contain any function symbols.

2 The Structures

2.1 Hrushovski's Method

For the proofs of our hierarchy theorems we need structures which have a certain property of homogeneity. It is not too difficult to get infinite homogeneous structures. However, we are working with finite structures, and there it is a real problem. Here we use a method that was introduced by Hrushovski [Hru92] to prove the following:

2.1 Hrushovski's Lemma *Given a finite graph $\mathfrak{G}$ and a set of partial isomorphisms of $\mathfrak{G}$ we can find a finite graph $\mathfrak{H} \supset \mathfrak{G}$ where all these partial isomorphisms can be extended to automorphisms.*

Herwig [Her95] extended this to arbitrary relational (finite) structures. These results were needed in a completely different context, namely to prove the so called *small index property* for certain theories.

To understand the intuitive contents of the construction method it is best to consider a very special case first. Suppose, we are given a finite relational structure $\mathfrak{A}$ and a partial isomor-

phism p of $\mathfrak{A}$. We shall construct a finite extension $\mathfrak{B}$ of $\mathfrak{A}$ where p can be extended to an automorphism.
Choose an integer $m \geq 1$ such that p^m is the identity on its domain. Such an m always exists, possibly its domain is empty. Let $\mathfrak{B}' = \mathfrak{A}_0 \cup \ldots \cup \mathfrak{A}_{m-1}$ be the disjoint union of m copies of $\mathfrak{A}$. Observe $f : B' \to B'$ mapping each $\mathfrak{A}_i$ to $\mathfrak{A}_{i+1}$ and $\mathfrak{A}_{m-1}$ to $\mathfrak{A}_0$ in the canonical way is an automorphism of $\mathfrak{B}'$. Next, we identify, for any i $(1 \leq i < m)$, the elements of $\mathrm{dom}(p)$ in the copy $\mathfrak{A}_i$ ($\mathfrak{A}_0$) with the elements of $\mathrm{im}(p)$ in the copy $\mathfrak{A}_{i-1}$ ($\mathfrak{A}_{m-1}$, respectively) in a way that a is identified with $p(a)$. The resulting structure is $\mathfrak{B}$.

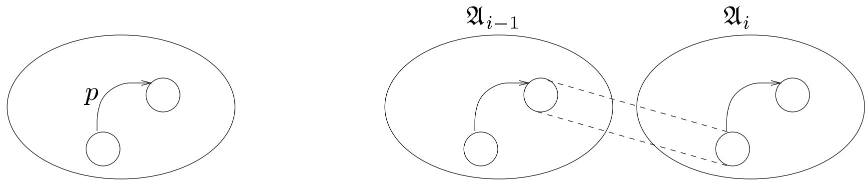

2.2 Figure. *Structure $\mathfrak{A}$ with the partial isomorphism p.* *Two copies of $\mathfrak{A}$ in $\mathfrak{B}$.*

The crucial observation is that f remains an automorphism of $\mathfrak{B}$, and in $\mathfrak{B}$ it extends the partial isomorphism p (taken in any copy $\mathfrak{A}_i$). So all there is left to do is show that the natural homomorphism of $\mathfrak{A}$ onto, say, $\mathfrak{A}_0$ remains an embedding after passing from $\mathfrak{B}'$ to $\mathfrak{B}$. Unfortunately, in general it is only an injective homomorphism. But we do not bother with fixing this, since we will never need it later.[1] It might be helpful for the reader to try and proof the claims we stated in this proof–sketch (f is an automorphism of $\mathfrak{B}$, in particular well–defined, and the natural homomorphism from $\mathfrak{A}$ to $\mathfrak{A}_0$ is injective) before proceeding. They are special cases of more general statements proved below.

We now generalize this idea to the case of several partial ismorphisms. The treatment is somewhat formal and uses a few (simple) notions from algebra. The whole point which makes this construction so useful for us is that we can tell many properties of the resulting structure without knowing how exactly it looks, just by being able to control its ingredients.
Let me add here that the following does not lead to a proof of Hrushovski's lemma, but is only an adaption of some of his ideas for our purposes.

We let σ be a relational signature, $\mathfrak{A}$ a σ–structure, and $p_1, \ldots, p_l$ partial bijections of A.
Let P be the free group generated by l elements $\mathsf{p}_1, \ldots, \mathsf{p}_l$. Each element $\mathsf{p} = \mathsf{p}_{i_1}^{s_1} \ldots \mathsf{p}_{i_r}^{s_r}$ $(r \geq 1, i_j \leq l, s_j \in \{1, -1\})$ of P naturally corresponds to a partial bijection $p = p_{i_1}^{s_1} \ldots p_{i_r}^{s_r}$ of A. (We write $p_{i_1}^{s_1} \ldots p_{i_r}^{s_r}$ for the concatenation of the partial bijections. By $a^{p_{i_1}^{s_1} \ldots p_{i_r}^{s_r}}$ we mean $p_{i_r}^{s_r}(\ldots p_{i_1}^{s_1}(a)\ldots)$. This always includes that a is in the domain of this partial bijection.) We

[1] Actually, in the special case with only one partial isomorphism there is a quite simple way to fix it: Instead of m copies take $2m$ copies and do the same construction, now it works. The following example shows why this is necessary: Let $\mathfrak{A}$ be the graph with universe $\{a, b, c\}$ and edges ab and bc. Let p be the partial isomorphism defined by $p(a) = b, p(b) = c$, and let $m = 3$. Then there is an edge between the images of c and a in $\mathfrak{B}$ under the natural homomorphism mapping $\mathfrak{A}$ to $\mathfrak{A}_0$.

extend this correspondence by associating the unit of P (the empty word) with the (totally defined) identity mapping on A.
Note that we obtain a mapping from P into the set of all partial bijections of A (which is a monoid homomorphism). We are a bit sloppy with this mapping and never refer to it explicitly, but we always let p be the partial bijection corresponding to p.

Let Γ be an arbitrary group and $\rho : \mathsf{P} \longrightarrow \Gamma$ a group homomorphism.
We define an equivalence relation $\sim$ on $A \times \Gamma$ by:

$$(a, \gamma) \sim (b, \delta) \iff \exists \mathsf{p} \in \mathsf{P} : \delta = \rho(\mathsf{p})\gamma \wedge a = b^p$$

In this case we occasionally say $(a, \gamma) \sim (b, \delta)$ *via* p.
It can easily be seen that $\sim$ is indeed an equivalence relation.

We define a σ-structure $\mathfrak{H} = \mathfrak{H}(\mathfrak{A}, p_1, \dots, p_l, \Gamma, \rho)$ with universe $H = A \times \Gamma/_\sim$ by

$$R^{\mathfrak{H}} b_1 \dots b_r \iff \exists \gamma \in \Gamma \ \exists a_1, \dots, a_r \in A : \ R^{\mathfrak{A}} a_1 \dots a_r \wedge (a_i, \gamma)/_\sim = b_i \text{ (for } 1 \leq i \leq r)$$

for each r–ary $R \in \sigma$.

2.3 Lemma. *For each $\gamma \in \Gamma$ the mapping*

$$\pi_\gamma : \mathfrak{A} \to \mathfrak{H} \text{ defined by } \pi_\gamma(a) := (a, \gamma)/_\sim$$

is a homomorphism.

The **proof** is an immediate consequence of the definition of the relations on $\mathfrak{H}$. ■

2.4 Lemma. *For each $\gamma \in \Gamma$ the mapping*

$$f_\gamma : H \to H \text{ defined by } f_\gamma((a, \delta)/_\sim) := (a, \delta\gamma)/_\sim$$

is an automorphism of $\mathfrak{H}$.

Proof. f_γ is well-defined because

$$\begin{aligned}
& (a, \delta)/_\sim = (b, \eta)/_\sim \\
\iff & \exists \mathsf{p} \in \mathsf{P} : \eta = \rho(\mathsf{p})\delta \wedge a = b^p \\
\iff & \exists \mathsf{p} \in \mathsf{P} : \eta\gamma = \rho(\mathsf{p})\delta\gamma \wedge a = b^p \\
\iff & (a, \delta\gamma) \sim (b, \eta\gamma) \\
\iff & f_\gamma((a, \delta)/_\sim) = f_\gamma((b, \eta)/_\sim)
\end{aligned}$$

f_γ is bijective because $(f_\gamma)^{-1} = f_{\gamma^{-1}}$ is also well–defined.

Finally, for r-ary $R \in \sigma$ and $b_1, \dots, b_r \in H$ we have

$$\begin{array}{rl} & R^{\mathfrak{H}} b_1 \dots b_r \\ \iff & \exists \delta \in \Gamma\ \exists a_1, \dots, a_r \in A\ \forall j \leq r:\ (a_j, \delta)/_\sim = b_j\ \wedge\ R^{\mathfrak{A}} a_1 \dots a_r \\ \iff & \exists \eta \in \Gamma\ \exists a_1, \dots, a_r \in A\ \forall j \leq r:\ (a_j, \eta)/_\sim = f_\gamma(b_j)\ \wedge\ R^{\mathfrak{A}} a_1 \dots a_r \\ & \text{(Since } f_\gamma \text{ is bijective.)} \\ \iff & R^{\mathfrak{H}} f_\gamma(b_1) \dots f_\gamma(b_r)\ . \end{array}$$

■

2.5 Remark. Observe that in a sense $f_{\rho(\mathsf{p})}$ extends p since

$$\pi_1(a^p) = (a^p, 1)/_\sim = (a, \rho(\mathsf{p}))/_\sim = f_{\rho(\mathsf{p})}((a,1)/_\sim) = f_{\rho(\mathsf{p})}(\pi_1(a))\ .$$

So it seems that we have nearly proved Hrushovski's Lemma. We just have to show that π_1 is an embedding. Unfortunately in general π_1 is not even injective, e.g. if Γ is the group $\{1\}$. In fact we are not as near as it seems. □

Although the mappings π_γ are not injective in general, they are for all structures we are dealing with because Γ and ρ will always be *suitable* for $p_1, \dots, p_l$ in the following sense:

2.6 Definition. Let $p_1, \dots, p_l$, Γ and ρ be as above. We call Γ, ρ *suitable* for $p_1, \dots, p_l$ if

$$\forall \mathsf{p} \in \mathsf{P}\ : (\rho(\mathsf{p}) = 1 \Longrightarrow p = id\restriction_{\mathrm{dom}(p)})\ .$$

2.7 Lemma. *Suppose Γ and ρ are suitable for $p_1, \dots, p_l$. Then for all $\gamma \in \Gamma$ the mapping π_γ is injective.*

Proof. Suppose $(a, \gamma) \sim (b, \gamma)$ *via* p. Then

$$\rho(\mathsf{p})\gamma = \gamma \quad \Longrightarrow \quad \rho(\mathsf{p}) = 1 \quad \Longrightarrow \quad p = id\restriction_{\mathrm{dom}(p)}$$

But since $(a, \gamma) \sim (b, \gamma)$ *via* p also implies $a = b^p$ we have $a = b$. ■

An important difference between our construction and Hrushovski's and Herwig's is that we are always going to work with Abelian groups Γ (which we write additively). This makes the resulting structure very homogeneous. We immediately get the following:

2.8 Lemma. *Suppose Γ is an Abelian group. Then for all $\mathsf{p} \in \mathsf{P}$, $a \in \mathrm{dom}(p)$, and $\delta \in \Gamma$ we have*

$$(a^p, \delta)/_\sim = f_{\rho(\mathsf{p})}((a,\delta)/_\sim).$$

Proof. $(a^p, \delta)/_\sim = (a, \rho(\mathsf{p}) + \delta)/_\sim = (a, \delta + \rho(\mathsf{p}))/_\sim = f_{\rho(\mathsf{p})}((a,\delta)/_\sim)$. ■

2.2 The Construction

Let us fix some $k \geq 1$ for this section. For each $n \geq 1$ we are going to define a graph $\mathfrak{G}$ such that the transitive closure of the formula

$$\textsc{Edge}_k(\overset{k}{x}, \overset{k}{y}) := \bigwedge_{1 \leq i,j \leq k} Ex_iy_j \wedge \bigwedge_{1 \leq i \neq j \leq k} (Ex_ix_j \wedge Ey_iy_j) \; .$$

is not definable in $(k-1)$–ary simultaneous partial fixed–point logic on these graphs. $\textsc{Edge}_k(\overset{k}{x}, \overset{k}{y})$ simply says that the elements of $\overset{k}{x}\overset{k}{y}$ form a $2k$–clique.

We start by defining a simple graph $\mathfrak{A} = \mathfrak{A}(k,n)$ whereon we apply the construction described in the previous paragraph. It consists of two disjoint $\textsc{Edge}_k$-paths of length n. Its universe is $A = \{1, \ldots, n\} \times \{-1, \ldots, -k, 1, \ldots, k\}$, and the edge–relation is defined by

$$E^{\mathfrak{A}}(I,a)(J,b) \iff (|I-J| = 1 \wedge \mathrm{sgn}(a) = \mathrm{sgn}(b)) \vee (I = J \wedge \mathrm{sgn}(a) = \mathrm{sgn}(b) \wedge a \neq b)^2$$

(see Figure 2.9). Note that $E^{\mathfrak{A}}$ is symmetric and anti–reflexive.
To simplify the notation we let $K := \{-1, \ldots, -k, 1, \ldots, k\}$ and define a mapping $\boldsymbol{c} : K \longrightarrow K$ by $\boldsymbol{c}(a) = -a$. We extend $\boldsymbol{c}$ to A by letting $\boldsymbol{c}(I,a) := (I, -a)$.
We define two more mappings row and col (*column*) on A by

$$\mathrm{row}((I,a)) = I \quad \text{and} \quad \mathrm{col}((I,a)) = |a| \; .$$

The graph $\mathfrak{G}$ we aim at will consist of n disjoint rows and $\textsc{Edge}_k$s between k–tuples in successive rows. There will be a natural embedding of $\mathfrak{A}$ into $\mathfrak{G}$ and (at least) two connected components arising from the two paths in $\mathfrak{A}$. So each element of G naturally belongs to one component, depending on whether it is derived from an element in $\{-1. \ldots, -k\}$ or $\{1. \ldots, k\}$. But for each $(k-1)$–tuple there is an automorphism of $\mathfrak{G}$ mapping it into the "other component" in its rows, but leaving most other rows fixed. (Remember that the components in fact consist of k-tuples and not of single elements !).

To achieve this, for all $I_1, \ldots, I_{k-1} \in \{1, \ldots, n\}$, $a \in \{1, \ldots, k\}$ we define a partial bijection of $\mathfrak{A}$ by

$$(J,b) \mapsto \begin{cases} (J,b) & \text{if } \forall i \leq k-1 : |J - I_i| > 1 \\ (J,-b) & \text{if } \exists i \leq k-1 : J = I_i \text{ and } |b| \neq |a| \\ (\text{undefined} & \text{otherwise}) \end{cases} \; .$$

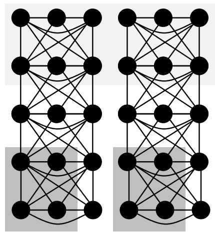

2.9 Figure. *Structure $\mathfrak{A}(3,5)$. The grey areas show the domain of the partial bijection p induced by $I_1 = 1$, $I_2 = 2$, and $a = 3$. p switches the two dark grey areas and is the identity on the light grey area.*

Let $p_1, \ldots, p_l$ be an enumeration of all partial bijections defined in this way and P the free group generated by $\mathsf{p}_1, \ldots, \mathsf{p}_l$. Furthermore, let Γ be the (Abelian) group $(\mathbb{Z}_2^l, +)$ and ρ the extension of the mapping $\mathsf{p}_i \longmapsto (0, \ldots, 0, 1, 0, \ldots, 0)$ (with the 1 at the i-th position) to a homomorphism from P onto Γ.
$\mathfrak{G} = \mathfrak{G}(k,n) = \mathfrak{H}(\mathfrak{A}(k,n), p_1, \ldots, p_l, \Gamma, \rho)$ is the structure we have been heading for.

Note that Γ and ρ are suitable (cf. Definition 2.6) for $p_1, \ldots, p_l$ because for all $i, j \leq l$ we have $p_i = p_i^{-1}$ and $p_i p_j = p_j p_i$. Thus by Lemma 2.7 for every $\gamma \in \Gamma$ the mapping $\pi_\gamma : A \longrightarrow G$ is injective.
Also note that for each $\mathsf{p} \in \mathsf{P}$ the corresponding partial bijection p is a partial isomorphism.

To show that $\mathfrak{G}$ has the desired properties we will have to prove a sequence of lemmata. The reader feeling uncomfortable at the moment will hopefully gain a some familiarity with our structure while doing so. In particular, Lemma 2.10 and Lemma 2.14 might help.
Before we start we introduce some notational conventions:

- We use the variables I, J for the rows.
- We use the variables $a, b, \ldots$ for the elements of K.
- We use the variables $\gamma, \delta, \epsilon, \zeta, \ldots$ for the elements of Γ.
- We write $(I, a, \gamma)/_\sim$ (instead of $\big((I,a),\gamma\big)/_\sim$) for the elements of G or use the variables $\boldsymbol{a}, \boldsymbol{b}, \ldots$.

We extend the mappings row, col, and $\mathfrak{c}$ to G by

$$\begin{aligned}
\mathrm{row}((I,a,\gamma)/_\sim) &:= \mathrm{row}((I,a)) = I \\
\mathrm{col}((I,a,\gamma)/_\sim) &:= \mathrm{col}((I,a)) = |a| \\
\mathfrak{c}((I,a,\gamma)/_\sim) &:= (I,-a,\gamma)/_\sim \,.
\end{aligned}$$

They are well–defined because $p_1, \ldots, p_l$ leave the row and the column of an element fixed. Furthermore, the mapping $\mathfrak{c}$ is an automorphism of $\mathfrak{G}$ because the construction is completely symmetric.

[2] $\mathrm{sgn}(a)$ is defined to be 1 if $a \geq 0$ and -1 otherwise.

To see that $\mathfrak{G}$ is a graph we use the following lemma which implies that $E^{\mathfrak{G}}$ is anti–reflexive. Clearly the relation $E^{\mathfrak{G}}$ is symmetric because $E^{\mathfrak{A}}$ is.

2.10 Lemma. *Let $\boldsymbol{a}, \boldsymbol{b} \in G$ such that $E^{\mathfrak{G}}\boldsymbol{ab}$. Then $\mathrm{col}(\boldsymbol{a}) \neq \mathrm{col}(\boldsymbol{b})$ or $\mathrm{row}(\boldsymbol{a}) \neq \mathrm{row}(\boldsymbol{b})$.*

Proof. Suppose $\mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b}) = I$. $E^{\mathfrak{G}}\boldsymbol{ab}$ implies that there exists a $\zeta \in \Gamma$ and $a', b' \in K$ such that $\boldsymbol{a} = (I, a', \zeta)/_\sim$, $\boldsymbol{b} = (I, b', \zeta)/_\sim$, and $E^{\mathfrak{A}}(I, a')(I, b')$. By the definition of $E^{\mathfrak{A}}$ this means $\mathrm{col}(\boldsymbol{a}) = |a'| \neq |b'| = \mathrm{col}(\boldsymbol{b})$. ■

Recall that $\Gamma = (\mathbb{Z}_2^l, +)$ and let for each $\gamma \in \Gamma$

$$\mathrm{supp}(\gamma) := \{i \leq l \mid \text{The } i\text{th component of } \gamma \text{ is } 1\} .$$

2.11 Lemma. *Let $\gamma, \delta \in \Gamma$, $I \leq n$, $a, b \in K$, and $\mathrm{supp}(\delta - \gamma) = \{i_1, \ldots, i_r\}$. Then*

$$(I, a, \gamma) \sim (I, b, \delta) \quad \Longleftrightarrow \quad (I, a, \gamma) \sim (I, b, \delta) \quad \textit{via} \quad \mathsf{p}_{i_1} \ldots \mathsf{p}_{i_r} .$$

Proof. The direction from the right to the left is trivial; the other direction is an easy consequence of the suitability of ρ and Γ. ■

Most of the work to be done here is to show that $\mathfrak{G}$ has (at least) two $\textsc{Edge}_k$–connected components.

2.12 Definition. We say that a tuple $\overset{k}{\boldsymbol{a}} \in G$ belongs to the *right (left) component* if there exists a permutation π of $\{1, \ldots, k\}$, an $I \leq n$, and a $\gamma \in \Gamma$ such that

$$\forall i \leq k : \ \boldsymbol{a}_{\pi(i)} = (I, (-)i, \gamma)/_\sim .$$

□

It is quite easy to see that

2.13 Lemma. *No k–tuple in $\mathfrak{G}$ belongs to the left and the right component.*

Proof. It clearly suffices to show that for all $\gamma, \delta \in \Gamma, I \leq n$ there exists an $a \in \{1, \ldots, k\}$ such that

$$(I, a, \gamma) \not\sim (I, \text{\ae}(a), \delta) .$$

So suppose there exists some $b \in K$ such that

$$(I, b, \gamma) \sim (I, \text{\ae}(b), \delta) \quad \textit{via} \quad \mathsf{p} .$$

(If no such b exists there is nothing to prove.)

By the previous lemma we can assume that $\mathsf{p} = \mathsf{p}_{i_1} \ldots \mathsf{p}_{i_r}$ where $\{i_1, \ldots, i_r\} = \mathrm{supp}(\delta - \gamma)$. Then there exists an $i \in \{i_1, \ldots, i_r\}$ such that $p_i((I, b)) = (I, \text{æ}(b))$.
By the definition of p_i there exists an $a \in \{1, \ldots, k\}$ such that $(I, a), (I, \text{æ}(a)) \notin \mathrm{dom}(p_i)$. Hence

$$\neg\Big((I, a, \gamma) \sim (I, \text{æ}(a), \delta) \quad \textit{via} \quad \mathsf{p}_{i_1} \ldots \mathsf{p}_{i_r}\Big)$$

which implies $(I, a, \gamma) \not\sim (I, \text{æ}(a), \delta)$, again by Lemma 2.11. ■

The difficulty is to prove that there are no $\textsc{Edge}_k$s between the components. This is done in the following three lemmata.

2.14 Lemma. *Let $I, J \leq n$, $a, b \in K$, $\gamma, \delta \in \Gamma$ such that*

$$E^{\mathfrak{G}}(I, a, \gamma)/_{\sim}(J, b, \delta)/_{\sim} \; .$$

Then there exist $\hat{a}, \hat{b} \in K, \eta \in \Gamma$ such that

(i) $(I, a, \gamma) \sim (I, \hat{a}, \eta)$

(ii) $(J, b, \delta) \sim (J, \hat{b}, \eta)$

(iii) $E^{\mathfrak{A}}(I, \hat{a})(J, \hat{b})$

(iv) $\mathrm{supp}(\eta - \gamma) \,\dot{\cup}\, \mathrm{supp}(\eta - \delta) = \mathrm{supp}(\delta - \gamma)$

Proof. We only have to prove that η can be chosen according to (iv), the rest follows immediately from the definition of $E^{\mathfrak{G}}$. Choose $\eta \in \Gamma, \hat{a}, \hat{b} \in K$ such that (i)–(iii) hold.
Note that

$$(\delta - \gamma) = (\eta - \gamma) + (\eta - \delta) \qquad (*)$$

(in the group $\Gamma = (\mathbb{Z}_2^l, +)$). Thus

$$\mathrm{supp}(\delta - \gamma) \subseteq \mathrm{supp}(\eta - \gamma) \cup \mathrm{supp}(\eta - \delta) \; .$$

For each $i \in \mathrm{supp}(\eta - \gamma) \setminus \mathrm{supp}(\delta - \gamma)$ we have, by $(*)$, $i \in \mathrm{supp}(\eta - \delta) \setminus \mathrm{supp}(\delta - \gamma)$. So we suppose that $\mathrm{supp}(\eta - \gamma) \setminus \mathrm{supp}(\delta - \gamma) = \{i_1, \ldots, i_r\}$ and let

$$\eta' = \rho(\mathsf{p}_{i_1}) + \ldots + \rho(\mathsf{p}_{i_r}) + \eta \; .$$

We still have

$$(\delta - \gamma) = (\eta' - \gamma) + (\eta' - \delta)$$

and in addition

$$\mathrm{supp}(\delta - \gamma) = \mathrm{supp}(\eta' - \gamma) \cup \mathrm{supp}(\eta' - \delta) \; .$$

Suppose now that there exists an

$$i \in \mathrm{supp}(\eta' - \gamma) \cap \mathrm{supp}(\eta' - \delta).$$

Then

$$i \notin \operatorname{supp}((\eta' - \gamma) + (\eta' - \delta)) = \operatorname{supp}(\delta - \gamma)$$

which is a contradiction. Thus (iv) holds for η' instead of η.

To see (i) (for η') note that $\{i_1, \ldots, i_r\} \subseteq \operatorname{supp}(\eta - \gamma)$, thus by (i) (for η) we have $(I, a), (I, -a) \in \operatorname{dom}(p_{i_1} \ldots p_{i_r})$. Select $\hat{\hat{a}}$ such that $(I, \hat{\hat{a}}) = (I, \hat{a})^{p_{i_1} \cdots p_{i_r}}$, then (i) holds for η' (instead of η) and $\hat{\hat{a}}$ (instead of $\hat{a}$). Analogously it can be seen that (ii) holds for η' and $\hat{\hat{b}}$ (where $(J, \hat{\hat{b}}) = (J, \hat{b})^{p_{i_1} \cdots p_{i_r}}$).

To see (iii) we just recall that $p_{i_1} \ldots p_{i_r}$ is a partial isomorphism. ■

Although the next lemma looks quite complicated, it is almost trivial once the statement is understood.

2.15 Lemma. *Let $\overset{k}{\boldsymbol{a}}, \overset{k}{\boldsymbol{b}} \in G$ such that $\mathfrak{G} \models \mathrm{EDGE}_k[\overset{k}{\boldsymbol{a}}, \overset{k}{\boldsymbol{b}}]$ and let $\overset{k}{\boldsymbol{a}}$ belong to the right component.*
Then there exist tuples $\overset{k}{\boldsymbol{a}'}, \overset{k}{\boldsymbol{b}'}$ such that

- $\mathfrak{G} \models \mathrm{EDGE}_k[\overset{k}{\boldsymbol{a}'}, \overset{k}{\boldsymbol{b}'}]$
- *$\overset{k}{\boldsymbol{b}}$ belongs to the right component if and only if $\overset{k}{\boldsymbol{b}'}$ belongs to the right component.*
- *There exist $I \leq n$, $J \in \{I+1, I-1\}$, $b_i \in \{i, -i\}$ (for each $i \leq k$), and $\gamma, \delta_1, \ldots, \delta_k \in \Gamma$ such that*

  $$\boldsymbol{a}'_i = (I, i, \gamma)/_\sim \quad \textit{and} \quad \boldsymbol{b}'_i = (J, b_i, \delta_i)/_\sim$$

  for all $i \leq k$.

  (Note that in particular $\overset{k}{\boldsymbol{a}'}$ also belongs to the right component.)

The analogous statement holds for $\overset{k}{\boldsymbol{a}}$ belonging to the left component.

Proof. Select first $\gamma \in \Gamma$, $I \leq n$, and a permutation π of $\{1, \ldots, k\}$ such that $\boldsymbol{a}'_i := \boldsymbol{a}_{\pi(i)} = (I, i, \gamma)/_\sim$ (this is possible since $\overset{k}{\boldsymbol{a}}$ belongs to the right component).

Note that, since $\mathfrak{G} \models \mathrm{EDGE}_k[\overset{k}{\boldsymbol{a}}, \overset{k}{\boldsymbol{b}}]$, the elements $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_k, \boldsymbol{b}_1, \ldots, \boldsymbol{b}_k$ form a $2k$–clique and use Lemma 2.10 to see that the $\boldsymbol{b}_i$ $(i \leq k)$ either all belong to row $I+1$ or all belong to row $I-1$. Moreover, all $\boldsymbol{b}_i$ belong to distinct columns.
Now we permute $\overset{k}{\boldsymbol{a}}$ and $\overset{k}{\boldsymbol{b}}$ in the desired way. They still form a $2k$–clique and hence an EDGE_k. ■

Now we are prepared to take the crucial step of the proof:

2.16 Lemma. *Let* $\overset{k}{\boldsymbol{a}}, \overset{k}{\boldsymbol{b}} \in G$ *such that* $\mathfrak{G} \models \mathrm{EDGE}_k[\overset{k}{\boldsymbol{a}}, \overset{k}{\boldsymbol{b}}]$.
Then $\overset{k}{\boldsymbol{a}}$ *belongs to the right component (left component) if and only if* $\overset{k}{\boldsymbol{b}}$ *belongs to the right component (left component, respectively).*

Proof. We only prove that if $\overset{k}{\boldsymbol{a}}$ belongs to the right component then $\overset{k}{\boldsymbol{b}}$ also belongs to the right component, the rest follows from the facts that $\boldsymbol{c}$ is an automorphism of $\mathfrak{G}$ which switches the components and that EDGE_k is symmetric.

By the last lemma we can assume without loss of generality that $\overset{k}{\boldsymbol{a}}$ and $\overset{k}{\boldsymbol{b}}$ are of the form $\boldsymbol{a}_i = (I, i, \gamma)/_\sim$ and $\boldsymbol{b}_i = (J, b_i, \delta_i)/_\sim$ (for each $i \leq k$), where $I \leq n$, $J \in \{I+1, I-1\}$, $b_i \in \{i, -i\}$, and $\gamma, \delta_1, \ldots, \delta_k \in \Gamma$.

We only consider the case $J = I+1$, the other possibility $J = I-1$ can be treated analogously.

The basic fact behind the following is that, since there are edges between all pairs $\boldsymbol{a}_i$, $\boldsymbol{b}_j$ the δ_j cannot differ from γ too much. On the first sight, it even seems that if $\delta_j - \gamma = \rho(\mathsf{p})$ then p cannot be undefined for both (I, i) and $(I+1, j)$. Actually, this is not perfectly true, it might happen that $\mathsf{p} = \mathsf{p}'\mathsf{p}''$ where p' is defined on (I, i) and $(p'')^{-1}$ is defined for $(I+1, j)$. This is what makes the proof difficult. However, analysing this and the analogous fact that follows from the existence of edges between the $\boldsymbol{b}_i$, $\boldsymbol{b}_j$ in detail we finally come to the clue that a p as above must fix rows I and $(I+1)$. But then it is harmless because it cannot change the components. Practically it is more convenient to eliminate ps fixing the rows first, we do this in Step 1.

During this proof we deviate from our usual notation and use $x, x_1, \ldots, y, y_1, \ldots$ to denote the indices of the $p_1, \ldots, p_l$.
For each $J \leq n$ we define the set of indices of p_x leaving row J fixed, letting

$$\mathrm{fix}(J) := \{x \leq l \mid \forall a \in K : p_x((J, a)) = (J, a)\} .$$

Note that for all $x \in \mathrm{fix}(J)$, $a \in K$ and $\zeta \in \Gamma$ we have

$$(J, a, \zeta) \sim (J, a, \rho(\mathsf{p}_x) + \zeta) . \tag{$*$}$$

Furthermore, we define the set of indices of p_x that are completely undefined on row J as

$$\mathrm{undef}(J) := \{x \leq l \mid \forall a \in K : (J, a) \notin \mathrm{dom}(p_x)\} .$$

By the definition of the p_x we have

$$\forall J < n : \mathrm{fix}(J) \setminus \mathrm{fix}(J+1) \subseteq \mathrm{undef}(J+1) .$$

In a First Step we show that we can assume without loss of generality that

$$(a) \quad \forall i \leq k : \mathrm{supp}(\delta_i - \gamma) \cap \mathrm{fix}(I+1) = \emptyset$$

and

$$(b) \quad \forall i \leq k : \mathrm{supp}(\delta_i - \gamma) \cap \mathrm{fix}(I) = \emptyset \ .$$

Proof: To see (a) suppose that for some $i \leq k$

$$\mathrm{supp}(\delta_i - \gamma) \cap \mathrm{fix}(I+1) = \{x_1, \ldots, x_r\} \ .$$

Let $\delta_i' := \rho(\mathsf{p}_{x_1}) + \ldots + \rho(\mathsf{p}_{x_r}) + \delta_i$. Then $\mathrm{supp}(\delta_i' - \gamma) \cap \mathrm{fix}(I+1) = \emptyset$ and $(I+1, b_i, \delta_i) \sim (I+1, b_i, \delta_i')$ by $(*)$. Thus we could have taken δ_i' instead of δ_i to guarantee (a).

To prove (b) note first that there cannot exist $i, j \leq k$ such that

$$\exists x \leq l : x \in \mathrm{supp}(\delta_i - \delta_j) \cap \mathrm{undef}(I+1) \ .$$

Otherwise for each $\delta \in \Gamma$ we would have $x \in \mathrm{supp}(\delta - \delta_i)$ or $x \in \mathrm{supp}(\delta - \delta_j)$. The former would imply (since $x \in \mathrm{undef}(I+1)$) that there exists no $\hat{b}_i$ such that $(I+1, b_i, \delta_i) \sim (I+1, \hat{b}_i, \delta)$, the latter that there exists no $\hat{b}_j$ such that $(I+1, b_j, \delta_j) \sim (I+1, \hat{b}_j, \delta)$. This contradicts $E^{\mathfrak{G}} \boldsymbol{b}_i \boldsymbol{b}_j$.

Together with (a) and since $\mathrm{fix}(I) \setminus \mathrm{fix}(I+1) \subseteq \mathrm{undef}(I+1)$ this implies

$$\forall i, j \leq k : \mathrm{supp}(\delta_i - \gamma) \cap \mathrm{fix}(I) = \mathrm{supp}(\delta_j - \gamma) \cap \mathrm{fix}(I).$$

Suppose $\mathrm{supp}(\delta_i - \gamma) \cap \mathrm{fix}(I) = \{y_1, \ldots, y_s\}$ (for some, hence for all $i \leq k$).
Then, as in the proof of (a), we let $\gamma' := \rho(\mathsf{p}_{y_1}) + \ldots + \rho(\mathsf{p}_{y_s}) + \gamma$ and see that for all $i \leq k$ we have $\mathrm{supp}(\delta_i - \gamma') \cap \mathrm{fix}(I) = \emptyset$ and $(I, i, \gamma) \sim (I, i, \gamma')$. Taking γ' instead of γ proves (b).

Second Step: For each $i \leq k$ there exists an $\hat{b_i} \in \{i, -i\}$ such that
$\boldsymbol{b}_i = (I+1, \hat{b_i}, \gamma)/_{\sim}$.

Proof: Let us fix $x \in \mathrm{supp}(\delta_i - \gamma)$. Note that there exists a $j \in \{1, \ldots, k\}$ such that $(I, j), (I, -j) \notin \mathrm{dom}(p_x)$ (since $x \notin \mathrm{fix}(I)$).
Because of $E^{\mathfrak{G}} \boldsymbol{a}_j \boldsymbol{b}_i$, by Lemma 2.14 there exists an $\eta \in \Gamma$ and $\hat{j} \in \{j, -j\}, \hat{b}_i \in \{i, -i\}$ such that

(i) $(I, j, \gamma) \sim (I, \hat{j}, \eta)$

(ii) $(I+1, b_i, \delta_i) \sim (I+1, \hat{b}_i, \eta)$

(iii) $E^{\mathfrak{A}}(I, \hat{j})(I+1, \hat{b}_i)$

(iv) $\mathrm{supp}(\eta - \gamma) \mathbin{\dot{\cup}} \mathrm{supp}(\eta - \delta_i) = \mathrm{supp}(\delta_i - \gamma)$.

Since $(I, j), (I, -j) \notin \mathrm{dom}(p_x)$ we have $x \notin \mathrm{supp}(\eta - \gamma)$ (by (i)) hence $x \in \mathrm{supp}(\eta - \delta_i)$ (because $x \in \mathrm{supp}(\delta_i - \gamma)$). Thus $(I+1, i), (I+1, -i) \in \mathrm{dom}(p_x)$ (by (ii)).

Since $x \in \mathrm{supp}(\delta_i - \gamma)$ was arbitrary this holds for all $x \in \mathrm{supp}(\delta_i - \gamma)$.

Suppose $\mathrm{supp}(\delta_i - \gamma) = \{x_1, \ldots, x_r\}$ and let $p = p_{x_1} \ldots p_{x_r}$. Then $\mathsf{p} \in \rho^{-1}(\delta_i - \gamma)$; and we have just seen that $(I+1, b_i) \in \mathrm{dom}(p)$.
Thus

$$(I+1, b_i, \delta_i) \sim ((I+1, b_i)^p, \gamma)$$

which proves the second step.

We are finished once we have proved:

Third Step: $(I+1, b_i)^p = (I+1, i)$

Let us look back to the beginning of the second step and consider $\eta, \hat{j}, \hat{b}_i$ with properties (i)–(iv) defined there.

Recall that $\mathrm{supp}(\delta_i - \gamma) = \{x_1, \ldots, x_r\}$ and assume (without loss of generality) that $x_1, \ldots, x_r$ are ordered in a way that there exists an $s \leq r$ such that

$$\mathrm{supp}(\eta - \delta_i) = \{x_1, \ldots, x_s\} \quad \text{and} \quad \mathrm{supp}(\eta - \gamma) = \{x_{s+1}, \ldots, x_r\} \ .$$

We have seen in the second step that $p_{x_{s+1}} \ldots p_{x_r}$ is defined for both $(I, \hat{j})$ and $(I+1, \hat{b}_i)$. Since it is a partial isomorphism we have

$$E^{\mathfrak{A}}(I, \hat{j})^{p_{x_{s+1}} \cdots p_{x_r}} (I+1, \hat{b}_i)^{p_{x_{s+1}} \cdots p_{x_r}}$$

hence $(I+1, \hat{b}_i)^{p_{x_{s+1}} \cdots p_{x_r}} = (I+1, i)$ (since $(I, \hat{j})^{p_{x_{s+1}} \cdots p_{x_r}} = (I, j)$).
On the other hand we have

$$(I+1, b_i)^{p_{x_1} \cdots p_{x_s}} = (I+1, \hat{b}_i)$$

thus $(I+1, b_i)^p = (I+1, i)$. ∎

The following theorem summarizes the properties of $\mathfrak{G}$:

2.17 Theorem. *Let $k, n \geq 2$. Then there exists a finite graph $\mathfrak{G} = \mathfrak{G}(k, n)$ such that:*

(1) There exists a mapping $\mathrm{row} : G \longrightarrow \{1, \ldots, n\}$ *such that*

$$\forall \boldsymbol{a}, \boldsymbol{b} \in G : (E^{\mathfrak{G}} \boldsymbol{ab} \Longrightarrow |\mathrm{row}(\boldsymbol{b}) - \mathrm{row}(\boldsymbol{a})| \leq 1) \ .$$

(2) There exists an automorphism $\mathbb{c}$ of $\mathfrak{G}$ which is self–inverse and preserves the rows.

(3) There exist tuples $\overset{k}{\boldsymbol{c}}, \overset{k}{\boldsymbol{d}} \in G$ in the first and last row, respectively, such that $\overset{k}{\boldsymbol{c}}\overset{k}{\boldsymbol{d}}$ belongs to the transitive closure of the $2k$–ary relation defined by $\mathrm{E{\scriptsize DGE}}_k$ *on $\mathfrak{G}$, whereas $\overset{k}{\boldsymbol{c}}\mathbb{c}(\overset{k}{\boldsymbol{d}})$ does not.*

(4) For all $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_{k-1} \in G$ there exists an automorphism f of $\mathfrak{G}$ that is self-inverse, preserves the rows, and maps $\overset{k-1}{\boldsymbol{b}}$ to $\text{\ae}(\overset{k-1}{\boldsymbol{b}})$, but leaves all elements in rows of distance > 1 from $\mathrm{row}(\boldsymbol{b}_1), \ldots, \mathrm{row}(\boldsymbol{b}_{k-1})$ fixed. (More formally, the last statement reads

$$\forall \boldsymbol{b} \in \mathfrak{G} : \big(\forall i \leq k : |\mathrm{row}(\boldsymbol{b}) - \mathrm{row}(\boldsymbol{b}_i)| > 1\big) \quad \Longrightarrow \quad f(\boldsymbol{b}) = \boldsymbol{b} \qquad)$$

Proof. Of course we take the structure $\mathfrak{G}(k,n)$ we have constructed before. Then (1) and (2) are immediate.

Letting $\boldsymbol{c}_i = (1, -i, 0)/_\sim$ and $\boldsymbol{d}_i = (n, -i, 0)/_\sim$, there is obviously an $\textsc{Edge}_k$–path from $\overset{k}{\boldsymbol{c}}$ to $\overset{k}{\boldsymbol{d}}$. On the other hand, by Lemma 2.16 no such path can leave the left component, so there is no path from $\overset{k}{\boldsymbol{c}}$ to $\text{\ae}(\overset{k}{\boldsymbol{d}})$. This proves (3).

To see (4), suppose that $\boldsymbol{b}_i = (I_i, b_i, \delta_i)/_\sim$ (for $i \leq k-1$).
By the definition of $p_1, \ldots, p_l$ there exists an $i \in \{1, \ldots, l\}$ such that $p_i((I,a)) = (I, \text{\ae}(a))$ if $I \in \{I_1, \ldots, I_{k-1}\}$ and $|a| \in \{|b_1|, \ldots, |b_{k-1}|\}$ and $p_i((I,a)) = (I,a)$ if $|I - I_j| > 1$ for all $j \leq k$.

Let $\gamma = \rho(\mathsf{p}_i)$; by Lemma 2.4 f_γ is an automorphism of $\mathfrak{G}$. By Lemma 2.8 we have for all $j \leq k-1$

$$f_\gamma(\boldsymbol{b}_j) = f_\gamma((I_j, b_j, \delta_j)/_\sim) = ((I_j, b_j)^{p_i}, \delta_j)/_\sim = (I_j, \text{\ae}(b_j), \delta_j)/_\sim = \text{\ae}(\boldsymbol{b}_j)$$

and for all I such that $|I - I_j| > 1$ for all $j \leq k$, $a \in K$, $\delta \in \Gamma$

$$f_\gamma((I, a, \delta)/_\sim) = ((I,a)^{p_i}, \delta)/_\sim = (I, a, \delta)/_\sim \; .$$

Clearly, f_γ is self–inverse and preserves the rows. ■

3 Fixed–Point Logics

As has been mentioned in the introduction, we are going to work with a generalization of the common fixed–point logics which allows simultaneous inductions. It is known that this does not increase the expressive power in general, but the situation looks different when we restrict the arity of the formulae.

The class s–FP of *simultaneous fixed–point formulae* is given by means of the calculus consisting of the first–order clauses and the clause

$$\text{(s–FP)} \quad \frac{\varphi_1, \ldots, \varphi_m}{[\mathrm{FP}_{\bar{x}_1, X_1, \ldots, \bar{x}_m, X_m}\ \varphi_1, \ldots, \varphi_m]\,\bar{u}}$$

where $m \geq 1$ and, for each $i \leq m$, X_i is a relation variable whose arity matches the length of $\bar{x}_i$, and $\bar{u}$ is a tuple of terms of the same length as $\bar{x}_1$.

The free (first and second order) variables of a simultaneous fixed–point formula are defined inductively in the standard way adding the clause

$$\mathrm{free}([\mathrm{FP}_{\bar{x}_1,X_1,\ldots,\bar{x}_m,X_m}\ \varphi_1,\ldots,\varphi_m]\,\bar{u}) = \mathrm{free}(\bar{u}) \cup \bigcup_{i=1}^{m}(\mathrm{free}(\varphi_i)\setminus\{\bar{x}_i,X_i\}).$$

To define the semantics, for each simultaneous fixed–point formula

$$[\mathrm{FP}_{\bar{x}_1,X_1,\ldots,\bar{x}_m,X_m}\ \varphi_1,\ldots,\varphi_m]\,\bar{u}$$

and interpretation $\mathfrak{I} = (\mathfrak{A},\alpha)$[3] we define sequences $(X^{\mathfrak{I}}_{ij})_{j\geq 0}$ $(1 \leq i \leq m)$ of relations on A by

$$\begin{aligned}
X^{\mathfrak{I}}_{i0} &:= \emptyset\\
X^{\mathfrak{I}}_{i(j+1)} &:= \{\bar{a}\in A \mid \mathfrak{I}\models\varphi_i[\bar{a},X^{\mathfrak{I}}_{1j},\ldots,X^{\mathfrak{I}}_{mj}]\}\\
X^{\mathfrak{I}}_{i\infty} &:= \begin{cases} X^{\mathfrak{I}}_{ik} & \text{where } k=\min\{j \mid \bigwedge_{i'=1}^{m} X^{\mathfrak{I}}_{i'j} = X^{\mathfrak{I}}_{i'(j+1)}\}\\ & \text{(if such a } k \text{ exists)}\\ \emptyset & \text{otherwise}\end{cases}
\end{aligned}$$

We let

$$\mathfrak{I}\models[\mathrm{FP}_{\bar{x}_1,X_1,\ldots,\bar{x}_m,X_m}\ \varphi_1,\ldots,\varphi_m]\,\bar{u} \quad\Longleftrightarrow\quad \alpha(\bar{u})\in X^{\mathfrak{I}}_{1\infty}$$

and define the semantics of *simultaneous partial fixed–point logic* **s–PFP** $= (\text{s–FP},\models)$ inductively.
FP is the subclass of s–FP built by the first–order clauses and the restriction of the (s–FP)–clause to the case $m = 1$. The corresponding sublogic *partial fixed–point logic* **PFP** $= (\mathrm{FP},\models)$ is known to have already the same expressive as **s–PFP**.

A simultaneous fixed–point operator $[\mathrm{FP}_{\bar{x}_1,X_1,\ldots,\bar{x}_m,X_m}\ldots]$ is *k–ary* if k is the maximum of the arities of $X_1,\ldots,X_m$. A (simultaneous) fixed–point formula is *k–ary* if it contains at most k–ary (simultaneous) fixed–point operators. **s–PFP**^{k} denotes the k–ary fragment of **s–PFP**.

There are several other fixed–point logics being studied. Let us briefly mention the following:

- *Inductive fixed–point logic* **IFP** $= (\mathrm{FP},\models_I)$ has the same syntax as **PFP** but different semantics. The easiest way to define $\models_I$ here is by

$$\mathfrak{I}\models_I[\mathrm{FP}_{\bar{x},X}\ \varphi]\,\bar{u} \quad\Longleftrightarrow\quad \mathfrak{I}\models[\mathrm{FP}_{\bar{x},X}\ X\bar{x}\vee\varphi]\,\bar{u}.$$

- *Least Fixed–Point Logic* **LFP** is the sublogic of **PFP** whose formulae only contain applications of the fixed–point operator to formulae positive in the relation variable of the operator. A well–known result of Gurevich and Shelah [GS86] says that (on finite structures) we have **LFP** = **IFP**.

[3] Recall that an interpretation is a pair $(\mathfrak{A},\alpha)$, where $\mathfrak{A}$ is a structure over an appropriate signature and α assigns a value in $\mathfrak{A}$ to each variable. We extend α to all terms.

- We call a fixed–point formula existential if it contains no universal quantifiers and negation symbols only occur in front of atomic subformulae. *Existential least fixed–point logic* $\exists$**LFP** is the sublogic of **LFP** whose formulae are existential.

  We can further restrict $\exists$**LFP** to existential fixed–point formulae where no negations symbols at all occur and obtain *positive existential least fixed–point logic* **pos**$\exists$**LFP**.

As for **PFP** we define the k–ary fragments of these logics. Obviously we have

$$\mathbf{pos}\exists\mathbf{LFP}^k \subseteq \exists\mathbf{LFP}^k \subseteq \mathbf{LFP}^k \subseteq \mathbf{IFP}^k \subseteq \mathbf{PFP}^k \subseteq \mathbf{s\text{–}PFP}^k.$$

It is also clear how to define the simultaneous versions of all these logics.

Remember the formula $\textsc{Edge}_k$ defined at the beginning of Subsection 2.2 and note that it is a conjunction of atomic formulae.
The transitive closure of $\textsc{Edge}_k$ is defined by the $\mathbf{pos}\exists\mathbf{LFP}^k$–formula

$$\textsc{Path}_k(\overset{k}{u},\overset{k}{v}) = \left[\mathrm{FP}_{\overset{k}{x},X}\textsc{Edge}_k(\overset{k}{u},\overset{k}{x}) \vee \exists\overset{k}{y}(X\overset{k}{y} \wedge \textsc{Edge}_k(\overset{k}{y},\overset{k}{x}))\right]\overset{k}{v}.$$

Thus as a corollary of our main lemma we obtain $\mathbf{pos}\exists\mathbf{LFP}^k \not\subseteq \mathbf{s\text{–}PFP}^{k-1}$ for each $k \geq 1$ and hence:

3.1 Corollary. *The arity hierarchies are strict for the fixed–point logics* **pos**$\exists$**LFP**, $\exists$**LFP**, **LFP**, **IFP**, **PFP**, *and for their simultaneous variants* **s–pos**$\exists$**LFP**, ... , **s–PFP**. *This holds even on the class of graphs.*

3.1 An Ehrenfeucht-Fraïssé Game

The following game is a variant of the well–known pebble game which corresponds to the infintary logic $\mathbf{L}^{\omega}_{\infty\omega}$. Just as the pebble game has turned out to be useful in proving non-expressibility results for fixed–point logics (using the fact that they are contained in $\mathbf{L}^{\omega}_{\infty\omega}$), our game will be used for proving non-expressibility in k–ary fixed–point logics.

3.2 Definition. Let $r, k \geq 0$. The *k–ary r–pebble game* is played by two players, the spoiler and the duplicator, on a pair $\mathfrak{A}, \mathfrak{B}$ of structures of the same signature with $2r$ pebbles $P_1, Q_1, \ldots, P_r, Q_r$.
We say that a pair (P_i, Q_i) of pebbles is *on the board* in a situation of the game if it is placed on the structures in that situation. The other pebbles are called *free*.
Each situation of the game and each pebble on the board have a depth ≥ 0. The game starts in the situation with depth 0 and all pebbles free.
In each situation of the game the spoiler selects one of the following moves:

$\exists$**-move** : The spoiler places a free pebble P_i on an element $a_i \in A$. The duplicator places the corresponding pebble Q_i on a $b_i \in B$. The depth of P_i and Q_i is defined to be the current depth of the game.

∀-move : The spoiler places a free pebble Q_i on $b_i \in B$. The duplicator places P_i on $a_i \in A$. The depth of P_i and Q_i is defined to be the current depth of the game.

I-move : The depth of the game is increased by 1.

R-move : This move is only allowed if the current depth of the game is ≥ 1. The spoiler reduces the depth of the game to an arbitrary $d \geq 1$ less than or equal to the current depth. Then he selects $l \leq k$ pairs (P_i, Q_i) of depth $\geq d$ to be left on the board. Their depth is reduced to d. All other pebbles of depth $\geq d$ are removed from the structures.

In each situation of the game the pairs of pebbled elements (a_i, b_i) and the pairs $(c^{\mathfrak{A}}, c^{\mathfrak{B}})$ of constants are called *pairs of corresponding elements.*
The duplicator wins the game if in each situation the pairs of corresponding elements form a partial isomorphism between $\mathfrak{A}$ and $\mathfrak{B}$. □

Observe that for all $k \geq r-1$ the k–ary r–pebble game has the same strength as the usual r–pebble game (the spoiler just has to start with an I–move).

To relate the game to fixed–point logics we need a notion of quantifier–rank for these logics. We extend the usual inductive definition by adding the clause

$$\mathrm{qr}([\mathrm{FP}_{\overset{l_1}{x_1}, X_1, \ldots, \overset{l_m}{x_m}, X_m} \ \varphi_1, \ldots, \varphi_m]\overset{l_1}{u}) = \max\{\mathrm{qr}(\varphi_i) + l_i \mid i \leq m\} \ .$$

3.3 Theorem. *Let $r, k \geq 0$ and $\mathfrak{A}$ and $\mathfrak{B}$ be two structures such that the duplicator has a winning strategy for the k–ary r–pebble game on $\mathfrak{A}, \mathfrak{B}$.*
Then the same **s–PFP**^{k}*–sentences of quantifier–rank $\leq r$ hold in $\mathfrak{A}$ and $\mathfrak{B}$.*

Proof: Let χ be an **s–PFP**^{k}–sentence of quantifier rank $\leq r$ such that $\mathfrak{A} \models \chi$. We shall prove that $\mathfrak{B} \models \chi$.

Although the following proof is much longer than it is supposed to be, it is straightforward and not very deep. The technical difficulty we have when trying to prove the theorem in the usual Ehrenfeucht-Fraïssé manner is that the game overestimates the logic, i.e. that we only have one direction of the desired equivalence. So we cannot prove the result by a simple induction on the formulae, looking back just to the immediate subformulae in the induction step. Instead, we have to proceed as follows: Starting from χ we have to go all the way down to the atomic formulae in structure $\mathfrak{A}$. Then we can pass to structure $\mathfrak{B}$ using the fact that winning positions for the duplicator are partial isomorphisms. Finally, we have to go back to χ in structure $\mathfrak{B}$.
Before we do so, in the first step of the proof we show how to describe simultaneous fixed–point operators by infinite disjunctions. Recalling that our game is a restriction of the usual pebble game for $\mathbf{L}^r_{\infty\omega}$ it is clear that this will be needed sooner or later. Then, in the second step, we take our way from χ down to the atomic formulae in $\mathfrak{A}$, and in the third step we pass to $\mathfrak{B}$ and return to χ.

There is a canonical way to construct for each **s–PFP**–formula χ an equivalent formula where negation symbols only occur in front of atomic subformulae or in front of FP–operators. We denote it by $\overline{\chi}$.

Without loss of generality we assume that $\chi = \overline{\chi}$. Furthermore, we may assume that each variable in χ is quantified at most once (i.e. for each x there exists at most one quantifier $\exists x$ or $\forall x$ or fixed-point operator $[\mathrm{FP}_{\bar{x}_1, X_1 \ldots, \bar{x}_m, X_m} \ldots]$ where x occurs among the variables in $\bar{x}_1, \ldots, \bar{x}_m$ in χ).

STEP 1: DEFINING THE STAGES

We consider a subformula $\psi = [\mathrm{FP}_{\overset{l_1}{x}_1, X_1, \ldots, \overset{l_m}{x}_m, X_m} \varphi_1, \ldots, \varphi_m]\overset{l_1}{u}$ of χ.

For simplicity we assume that for each subformula of ψ of the form $X_i \overset{l_i}{t}$ we have $\overset{l_i}{t} = \overset{l_i}{x}_i$.
(There are two ways to deal with the fact that the original formula ψ might have a subformula $X_i \overset{l_i}{t}$ where $\overset{l_i}{t} \neq \overset{l_i}{x}_i$: The first is to rename variables in the formulae ψ_{ij} (defined below) suitably. This is the the proper way to do it, but it would be annoying to carry the renaming of the variables through the whole proof.
The second way is to replace $X_i \overset{l_i}{t}$ by $\exists \overset{l_i}{z} \big(\overset{l_i}{t} = \overset{l_i}{z} \wedge \exists \overset{l_i}{x}_i (\overset{l_i}{z} = \overset{l_i}{x}_i \wedge X \overset{l_i}{x}_i) \big)$ where $z_1, \ldots, z_{l_i}$ are variables that do not occur in $\mathrm{free}(\overset{l_i}{t}) \cup \{x_{i1}, \ldots, x_{il_i}\}$. Note that this may increase the quantifier–rank of the formula so we have to reformulate the theorem a bit.)

Now, we let $\psi_{i0} := x_{i1} \neq x_{i1}$ and for each $j \geq 0$ we let $\psi_{i(j+1)}$ be the formula obtained from φ_i by replacing, for each $i' \leq m$, each subformula of the form $X_{i'} \overset{l_{i'}}{x}_{i'}$ by $\psi_{i'j}$.

Then clearly ψ_{ij} defines the set $X_{ij}^{\mathfrak{I}}$ for each interpretation $\mathfrak{I}$ (note that the restriction that none of the free variables of ψ is quantified in a φ_i is needed here) and we have

$$\mathfrak{I} \models \psi \quad \Longleftrightarrow \quad \exists j \geq 1 : \mathfrak{I} \models \underbrace{\exists \overset{l_1}{x}_1 (\overset{l_1}{x}_1 = \overset{l_1}{u} \wedge \psi_{1j}) \wedge \bigwedge_{i=1}^{m} \forall \overset{l_i}{x}_i \big(\psi_{ij} \leftrightarrow \psi_{i(j+1)} \big)}_{=: \psi_j^*(\overset{l_1}{u})} \ .$$

(Note that this can also be written as an infinite disjunction.)
We call subformulae ψ_{ij} and $\neg\psi_{ij}$ of the $\psi_{i'j'}$ $(i, i' \leq m,\ 0 \leq j < j')$ which are obtained by replacing subformulae $X \overset{l_i}{x}_i$ or $\neg X \overset{l_i}{x}_i$ of $\varphi_{i'}$ *R(eplacement)–formulae of* ψ.

STEP 2: PROOF–SEQUENCES FOR χ

DEFINITION: A *proof–situation* s consists of a formula ψ and a mapping ι from the free variables of ψ to $\{1, \ldots, r\}$. In addition, it may contain a move of the k–ary r–pebble game.
A *proof–sequence* S is a sequence of proof–situations such that

- The sequence of moves in the k–ary r–pebble game that belongs to S is according to the rules of the game, and the duplicator's moves are according to her winning strategy.

Hence each proof–situation s in S corresponds to a situation in the k–ary r–pebble game (obtained by the moves of the previous proof–situations (including s) in S).

- If ι is the mapping that belongs to s and i is in the range of ι then the pebbles P_i and Q_i are on the board in this situation.

Thus the mapping ι induces two (partial) assignments α and β (for the free variables of the formula ψ) defined by $\alpha(x) = a$ and $\beta(x) = b$ if $\iota(x) = i$ and P_i and Q_i are placed on $a \in A$ and $b \in B$ respectively.

- We always have $(\mathfrak{A}, \alpha) \models \psi$.

If S and T are two sequences then $S^\frown T$ denotes their concatenation.

We are going to define a set $\mathcal{S}$ of proof–sequences. The formulae that occur in these proof–sequences are essentially subformulae of χ or formulae $\psi_j^*(\overset{l}{u})$.
Depending on the form of such formulae ψ, we inductively define the set $\mathcal{S}$. The induction will be "from the outside to the inside", i.e. we start with χ and continue with its subformulae and their subformulae and so on.

Let s_0 be the proof–situation with formula χ, no move in the k–ary r–pebble game and the empty mapping ι (χ is a sentence !).
We start our induction with $\mathcal{S} = \{s_0\}$ (here we do not distinguish between the proof–situation s_0 and the proof–sequence consisting only of s_0). Note that s_0 is indeed a proof–sequence since $\mathfrak{A} \models \chi$.

Now, for each proof–sequence $S \in \mathcal{S}$ we consider the formula ψ of its last proof–situation. Let ι be the mapping belonging to this situation and α the assignment for the free variables of ψ in $\mathfrak{A}$ it induces. Since S is a proof–sequence we have $(\mathfrak{A}, \alpha) \models \psi$.
If ψ is not atomic or negated atomic and not an R–formula we proceed as follows, depending on the form of ψ (see below for the case that ψ is an R–formula):

$\psi = \varphi_1 \vee \varphi_2$: Since $(\mathfrak{A}, \alpha) \models \psi$ we have $(\mathfrak{A}, \alpha) \models \varphi_1$ or $(\mathfrak{A}, \alpha) \models \varphi_2$; without loss of generality we assume the first. Let s_1 be the proof–situation consisting of no move, the formula φ_1, and the restriction of ι to free(φ_1) as its mapping. $(\mathfrak{A}, \alpha) \models \varphi_1$ guarantees that $S^\frown s_1$ is a proof–sequence.

We replace S in $\mathcal{S}$ by $S^\frown s_1$.

$\psi = \varphi_1 \wedge \varphi_2$: Let s_i $(i = 1, 2)$ be the proof–situation consisting of no move, the formula φ_i, and the restriction of ι to free(φ_i). Then $S^\frown s_i$ is a proof–sequence.

We remove S from $\mathcal{S}$ and put in $S^\frown s_1$ and $S^\frown s_2$.

$\psi = \exists x \varphi(x)$: Select an $a \in A$ such that $(\mathfrak{A}, \alpha) \models \varphi[a]$.

Let s_a be the proof–situation whose move is the following $\exists$–move: The spoiler places a free pebble, say P_i, on a. The duplicator places Q_i on $\mathfrak{B}$ according to her winning strategy. (Our treatment of R–formulae described below guarantees that there are always enough free pebbles.)

Let φ be the formula of s_a and $(\iota \cup \{(x, i)\})\restriction_{\text{free}(\varphi)}$ its mapping.

We replace S in $\mathcal{S}$ by $S^\frown s_a$.

$\psi = \forall x \varphi(x)$: For each $b \in B$ let s_b be the proof–situation whose move is the following $\forall$–move: The spoiler places a free pebble, say Q_i, on b. The duplicator places P_i on $\mathfrak{A}$ according to her winning strategy.

Let φ be the formula of s_b and $(\iota \cup \{(x,i)\})\restriction_{\text{free}(\varphi)}$ its mapping.

We replace S in $\mathcal{S}$ by all sequences $S^\frown s_b$ $(b \in B)$.

Observe that in the cases $\psi = \exists x \varphi(x)$ and $\psi = \forall x \psi(x)$ the pebbles P_i, Q_i are connected to the variable x. This connection remains as long as P_i, Q_i are not removed from the board, which may only happen in an R–move occurring with an R–formula (see below). Note also that each variable may only be connected to one pair of pebbles since we have assumed that each variable is only quantified once in the formula χ.

$\psi = [\mathbf{FP}_{\overset{\iota_1}{x_1},X_1,\ldots,\overset{\iota_m}{x_m},X_m}\ \varphi_1,\ldots,\varphi_m]\overset{l_1}{u}$: Select $j \geq 1$ such that $(\mathfrak{A},\alpha) \models \psi_j^*(\overset{l_1}{u})$.

Let s_j be the proof–situation consisting of an I–move, the formula $\overline{\psi_j^*(\overset{l_1}{u})}$ and the mapping ι restricted to $\text{free}(\psi_j^*(\overset{l_1}{u}))(\subseteq \text{free}(\psi))$.

We replace S in $\mathcal{S}$ by $S^\frown s_j$.

$\psi = \neg[\mathbf{FP}_{\overset{\iota_1}{x_1},X_1,\ldots,\overset{\iota_m}{x_m},X_m}\ \varphi_1,\ldots,\varphi_m]\overset{l_1}{u}$: Note that for each $j \geq 1$ we have

$$(\mathfrak{A},\alpha) \models \neg\psi_j^*(\overset{l_1}{u}).$$

Let s_j be the proof situation consisting of an I–move, the formula $\overline{\neg\psi_j^*(\overset{l_1}{u})}$ and the mapping ι restricted to $\text{free}(\psi_j^*(\overset{l_1}{u}))$.

We replace S in $\mathcal{S}$ by the sequences $S^\frown s_j$ $(j \geq 1)$.

However, since the quantifier–rank of formulae of the form $\psi_j^*(\overset{l}{u})$ is not necessarily $\leq r$ it may happen that we run out of free pebbles. To prevent this, whenever ψ is an R–formula (cf. Step 1) we proceed as follows:

ψ R–formula: In this case we have to extend S twice, the first time to include an R–move into our sequence that sets some pebbles free, and the second time in the usual way (described above) to take care of ψ. (Being an R–formula is only an additional feature and does not say anything about the form of ψ.)

FIRST EXTENSION: Suppose ψ is an R–formula $\overline{\xi_{ij}}$ or $\overline{\neg\xi_{ij}}$[4] for a formula

$$\xi = (\neg)[\mathrm{FP}_{\overset{\iota_1}{x_1},X_1,\ldots,\overset{\iota_m}{x_m},X_m}\ \varphi_1,\ldots,\varphi_m]\overset{l_1}{u}\ .$$

[4]More precisely, ξ_{ij} or $\neg\xi_{ij}$ is an R–formula (instead of $\overline{\xi_{ij}}$ or $\overline{\neg\xi_{ij}}$).

Then there exists a proof–situation s in S with formula ξ. The move of the situation following s must have been an I–move which increased the depth of the game to, say, d.

Let s' be the situation consisting of the formula ψ, the mapping ι (of the current situation), and the following R–move:

The spoiler reduces the depth to d and removes all pebbles of depth $\geq d$ except those pairs (P_n, Q_n) with $n \in \{\iota(x_{ih}) \mid h \leq l_i\}$ from the board.

Note that, since $\mathrm{free}(\xi_{ij}) \subseteq \mathrm{free}(\xi) \cup \{x_{i1}, \ldots, x_{il_i}\}$, all pebbles (P_n, Q_n) with $n \in \{\iota(x) \mid x \in \mathrm{free}(\psi)\}$ are still on the board as they are either of depth $< d$ (they have been played, and thus connected with their variable in the sense explained above, before the fixed–point operator of ξ occured, hence before the depth was increased to d) or of the form (P_n, Q_n) with $n \in \{\iota(x_{ih}) \mid h \leq l_i\}$.

Thus $S' = S^\frown s'$ is a proof–sequence (because S is one). We replace S by S'.

Second Extension: We now proceed with S' as if ψ (the formula belonging to s') was not an R–formula. (I.e. we apply one of the "standard" operations defined above to S', depending on the form of ψ.)

We continue with this procedure until each proof–sequence in $\mathcal{S}$ ends with an atomic formula. Note that, although $\mathcal{S}$ may be infinite (because of the infinite disjunctions used to deal with fixed–point operators), each sequence $S \in \mathcal{S}$ is of finite length. (Formally, we can prove this by transfinite induction.)

Note also that our treatment of the R–formulae guarantees that we never run out of free pebbles. This can also be verified by induction (on the (finite) length of a proof–sequence).

Step 3: Proving $\mathfrak{B} \models \chi$ with χ's Proof–Sequences

Inductively we prove that for each proof–situation s occurring in a sequence $S \in \mathcal{S}$:

Claim: *If ψ is the formula of s and β the assignment for the free variables of ψ in $\mathfrak{B}$ induced by the mapping ι of s we have $(\mathfrak{B}, \beta) \models \psi$.*

This implies $\mathfrak{B} \models \chi$.

Proof: This time the induction is "from the inside to the outside":

ψ atomic or negated atomic: Let α be the assignment for the free variables of ψ in $\mathfrak{A}$ induced by the mapping ι of s.

By Step 2 we have $(\mathfrak{A}, \alpha) \models \psi$ (by the definition of proof sequences). Furthermore, since the duplicator has always played according to her winning strategy in the k–ary r–pebble game the mapping between the corresponding elements is a partial isomorphism between $\mathfrak{A}$ and $\mathfrak{B}$.

By the definition of α and β this means that the same atomic or negated atomic formulae with variables in $\mathrm{free}(\psi)$ hold in $(\mathfrak{A}, \alpha)$ and $(\mathfrak{B}, \beta)$. In particular, we have $(\mathfrak{B}, \beta) \models \psi$.

The cases $\psi = \varphi_1 \vee \varphi_2$, $\psi = \varphi_1 \wedge \varphi_2$, $\psi = \exists x \varphi$, and $\psi = \forall x \varphi$ are all easy. To illustrate this we consider only the last:

$\boldsymbol{\psi = \forall x \varphi}$ **:** By the definition of $\mathcal{S}$ for each $b \in B$ there exists a sequence $S \in \mathcal{S}$ where situation s is followed by a situation s_b with formula φ and induced assignment $\beta\frac{b}{x}$ extending β by $\beta\frac{b}{x}(x) = b$.

By the induction hypothesis we know that $(\mathfrak{B}, \beta\frac{b}{x}) \models \varphi$. Hence $(\mathfrak{B}, \beta) \models \psi$.

The fixed–point cases are also easy, e.g.:

$\boldsymbol{\psi = [\mathbf{FP}_{\overset{l_1}{x_1}, X_1, \ldots, \overset{l_m}{x_m}, X_m} \ \varphi_1, \ldots, \varphi_m]\overset{l_1}{u}}$ **:** By the definition of $\mathcal{S}$ there exists a $j \geq 1$ and a sequence $S \in \mathcal{S}$ where situation s is followed by situation s_j with formula $\overline{\psi_j^*(\overset{l_1}{u})}$ and induced assignment $\beta' \subseteq \beta$.

By the induction hypothesis we know that $(\mathfrak{B}, \beta') \models \overline{\psi_j^*(\overset{l_1}{u})}$ and thus $(\mathfrak{B}, \beta) \models \psi$.

The case $\psi = \neg[\mathrm{FP}_{\overset{l_1}{x_1}, X_1, \ldots, \overset{l_m}{x_m}, X_m} \ \varphi_1, \ldots, \varphi_m]\overset{l_1}{u}$ can be treated similarly.

Note that if ψ is an R–formula it occurs in two successive proof–situations, but that does not affect $(\mathfrak{B}, \beta) \models \psi$. ■

3.2 Proof of the Main Lemma

Again we fix some $k \geq 1$. Recall that the main lemma states that the transitive closure of the formula $\textsc{Edge}_k(\overset{k}{x}, \overset{k}{y})$ is not $\mathbf{s}$–$\mathbf{PFP}^{k-1}$–definable.

We let $r \geq 0$ (intended to be the quantifier rank of an $\mathbf{s}$–$\mathbf{PFP}^{k-1}$–formula) and $n = 2^{r+1} + 1$. Furthermore, we choose a graph $\mathfrak{G} = \mathfrak{G}(k, n)$ according to Theorem 2.17. Augmenting our signature by constant symbols $\underline{c}_1, \ldots, \underline{c}_k$ and $\underline{d}_1, \ldots, \underline{d}_k$, we expand it to the structures $\mathfrak{G}_1 := (\mathfrak{G}, \overset{k}{\boldsymbol{c}}, \overset{k}{\boldsymbol{d}})$ and $\mathfrak{G}_2 := (\mathfrak{G}, \overset{k}{\boldsymbol{c}}, \boldsymbol{c}(\overset{k}{\boldsymbol{d}}))$ where $\overset{k}{\boldsymbol{c}}$ and $\overset{k}{\boldsymbol{d}}$ are tuples chosen according to 2.17(3).[5]
Clearly we are finished once we have proved that the same $\mathbf{PFP}^{k-1}$–sentences of quantifier rank $\leq r$ hold in $\mathfrak{G}_1$ and $\mathfrak{G}_2$. By Theorem 3.3 it suffices to prove:

3.4 Proposition. *The duplicator has a winning strategy for the $(k-1)$–ary r–pebble game on $\mathfrak{G}_1$ and $\mathfrak{G}_2$.*

Proof. Recall that the pairs of corresponding elements in a situation of the game are the pairs of constants and pairs of pebbled elements.

We prove inductively that the duplicator can preserve the following throughout the whole game:
Suppose $r' \leq r$ pairs of pebbles are on the board.
Then there exist $M, N \in \{1, \ldots, n\}$ and an automorphism h of $\mathfrak{G}_1$ such that:

[5]Our notation is explained as follows: $\mathfrak{G}_1$ is the structure whose $\{E\}$–reduct is $\mathfrak{G}$ and where $\underline{c}_i^{\mathfrak{G}_1} = c_i$ and $\underline{d}_i^{\mathfrak{G}_1} = {}_i$ for each $i \leq k$. $\mathfrak{G}_2(k, n)$ denotes the same structure, except that $d_i^{\mathfrak{G}_2} = \boldsymbol{c}(d_i)$.

(i) $N - M \geq 2^{r+1-r'}$

(ii) For pairs $(\boldsymbol{a}, \boldsymbol{b})$ of corresponding elements we have:

$$\mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b}) \leq M \quad \text{or} \quad \mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b}) \geq N \; .$$

(iii) h is self–inverse and preserves the rows.

(iv) For each pair $(\boldsymbol{a}, \boldsymbol{b})$ of corresponding elements the following holds:

- If $\mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b}) \leq M$ then $h(\boldsymbol{a}) = \boldsymbol{b}$.
- If $\mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b}) \geq N$ then $\mathbb{c}(h(\boldsymbol{a})) = \boldsymbol{b}$.

Note that this implies that the pairs of corresponding elements form a partial isomorphism from $\mathfrak{G}_1$ to $\mathfrak{G}_2$ since $\mathbb{c}$ is an automorphism of $\mathfrak{G}$ and we only have edges between successive rows of $\mathfrak{G}$.

Clearly (i)–(iv) hold in the beginning with $M = 1, N = n, h = id$ since then the only corresponding elements are the pairs of constants $(\boldsymbol{c}_1, \boldsymbol{c}_1), \ldots, (\boldsymbol{c}_k, \boldsymbol{c}_k)$ in row 1 and $(\boldsymbol{d}_1, \mathbb{c}(\boldsymbol{d}_1)), \ldots, (\boldsymbol{d}_k, \mathbb{c}(\boldsymbol{d}_k))$ in row $n = 2^{r+1} + 1$.

So let there be $r' \leq r$ pairs of pebbles on the board and let M, N, h be defined such that (i)–(iv) hold. We are going to define the duplicator's answer to the next move and M', N', h' to preserve (i)–(iv) in the next situation. We have to consider the following moves:

$\exists$-move: The spoiler places a pebble on $\boldsymbol{a} \in \mathfrak{G}_1$. Let $I = \mathrm{row}(\boldsymbol{a})$.

> **Case 1:** $I - M \leq N - I$
>
> Let $M' := \max\{M, I\}$ and $N' := N$. The duplicator places the corresponding pebble on $\boldsymbol{b} := h(\boldsymbol{a})$.
>
> **Case 2:** $I - M > N - I$
>
> Let $M' := M$ and $N' := \min\{I, N\}$. The duplicator places the corresponding pebble on $\boldsymbol{b} := \mathbb{c}(h(\boldsymbol{a}))$.
>
> Note that in both cases
>
> $$N' - M' \geq \left\lfloor \frac{N-M}{2} \right\rfloor \geq 2^{r+1-r'-1}$$
>
> so (i) holds. Letting $h' = h$ guarantees (iii) and (iv). Since h preserves the rows we have (ii).

$\forall$-move: Can be treated analogously to the $\exists$-move.

I-move: Nothing has to be done here. Just note the following:

> Say, the depth is increased from $d - 1$ to d. Then all pebbles which have depth $< d$ at the moment cannot be moved until the depth is reduced below d.

R–move: The spoiler reduces the depth to $d \geq 1$ and leaves only the pairs $(\boldsymbol{a}_1, \boldsymbol{b}_1), \ldots, (\boldsymbol{a}_l, \boldsymbol{b}_l)$ $(l \leq k-1)$ of depth $\geq d$ pebbled.

All other pebbles that are still on the board (those of depth $< d$) have already been in their position when the last I–move increasing the depth of the game from $d-1$ to d was made. We look back to that situation; say M'', N'', h'' satisfied (i)–(iv) then. Moreover, let r'' be the number of pebbles that have been on the board then, thus $(r''+l)$ are on the board now.

By the pigeonhole principle there exist M', N' such that $M'' \leq M' < N' \leq N''$,

$$N' - M' \geq \frac{N'' - M''}{l+1} \geq \frac{2^{r+1-r''}}{l+1} \geq 2^{r+1-(r''+l)},$$

and none of the pairs of corresponding elements that are still on the board are in a row between M' and N'. So (i) and (ii) hold for M', N'.

Note that by the induction hypothesis there are also no pairs of corresponding elements in a row between M and N.

But M' and N' have nothing to do with M and N so it may happen that some of the $\boldsymbol{a}_1, \boldsymbol{b}_1, \ldots, \boldsymbol{a}_l, \boldsymbol{b}_l$ are on the "wrong side". In fact the following two cases may cause some trouble:

Case 1: Some of the pairs $(\boldsymbol{a}_i, \boldsymbol{b}_i)$ $(i \leq l)$ are in a row between N' and M (see Figure 3.5).

Let $\{\boldsymbol{a}_{i_1}, \ldots, \boldsymbol{a}_{i_m}\}$ be the maximal subset of $\{\boldsymbol{a}_1, \ldots, \boldsymbol{a}_l\}$ with

$$N' \leq \mathrm{row}(\boldsymbol{a}_{i_1}), \ldots, \mathrm{row}(\boldsymbol{a}_{i_m}) \leq M\ .$$

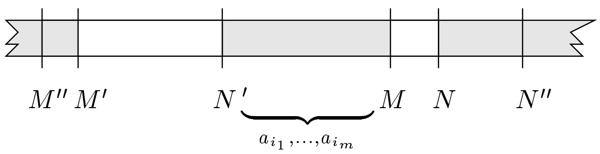

3.5 Figure.

Then for each pair of corresponding elements

$$(\boldsymbol{a}, \boldsymbol{b}) \notin \{(\boldsymbol{a}_{i_1}, \boldsymbol{b}_{i_1}), \ldots, (\boldsymbol{a}_{i_m}, \boldsymbol{b}_{i_m})\}$$

we have

$$\forall j \leq m: \ |\mathrm{row}(\boldsymbol{a}) - \mathrm{row}(\boldsymbol{a}_{i_j})| > 1$$

(since $N' - M', M - N \geq 2^1 = 2$).

Since $m \leq l \leq (k-1)$, by Theorem 2.17(4) there exists an automorphism f of $\mathfrak{G}$ such that

- f is self–inverse and preserves the rows.
- $\forall j \leq m : \; f(\boldsymbol{b}_{i_j}) = \mathbb{c}(\boldsymbol{b}_{i_j})$
- If $\boldsymbol{a} \in G$ such that $\forall j \leq m : |\text{row}(\boldsymbol{a}) - \text{row}(\boldsymbol{b}_{i_j})| > 1$ then $f(\boldsymbol{a}) = \boldsymbol{a}$.

We let $h' := f \circ h$. Since $M'' \geq 1$ and $N'' \leq n$ the automorphism f keeps the first and the last row fixed, in particular the constants $\overset{k}{\boldsymbol{c}}$ and $\overset{k}{\boldsymbol{d}}$, thus h' is in fact an automorphism of $\mathfrak{G}_1$. Moreover, it is clearly self–inverse and preserves the rows hence (iii) holds.

We get (iv) since for each pair of corresponding elements $(\boldsymbol{a}, \boldsymbol{b})$ we have:

- $\text{row}(\boldsymbol{a}) \leq M' \Longrightarrow h'(\boldsymbol{a}) = f(h(\boldsymbol{a})) = h(\boldsymbol{a}) = \boldsymbol{b}$
- $N' \leq \text{row}(\boldsymbol{a}) \leq M \Longrightarrow \exists j \leq m : (\boldsymbol{a}, \boldsymbol{b}) = (\boldsymbol{a}_{i_j}, \boldsymbol{b}_{i_j})$
  $\Longrightarrow \mathbb{c}(h'(\boldsymbol{a})) = \mathbb{c}(f(h(\boldsymbol{a}))) = \mathbb{c}(f(\boldsymbol{b})) = \mathbb{c}(\mathbb{c}(\boldsymbol{b})) = \boldsymbol{b}$
- $\text{row}(\boldsymbol{a}) \geq N (\geq N') \Longrightarrow \mathbb{c}(h'(\boldsymbol{a})) = \mathbb{c}(f(h(\boldsymbol{a}))) = \mathbb{c}(h(\boldsymbol{a})) = \boldsymbol{b}$.

Case 2: Some of the $\boldsymbol{a}_i$ $(i \leq l)$ are in a row between N and M'.

This case can be treated analogously.

If none of these cases occurs we have no problems anyway; we can just let $h' = h$.

This finishes our proof of Proposition 3.4. ■

4 Datalogics

We only give a very brief introduction into the database language Datalog (which is strongly related to the fixed–point logics); all definitions are given in a way that fits immediately into our formal frame. Datalog was first in introduced in [CH85]; for a more detailed presentation of the definitions and results given here we refer the reader to [Kan90, EF95].

4.1 Definition. A *general Datalog program* Π is a finite set of *clauses* of the form

$$\gamma \leftarrow \gamma_1, \ldots, \gamma_l$$

where $l \geq 0$ and $\gamma_1, \ldots, \gamma_l$ are atomic or negated atomic formulae. Moreover, γ must be of the form $X\bar{x}$ where X is a relation variable. It is called the *head* of the clause; $\gamma_1, \ldots, \gamma_l$ form the *body* of the clause.
Relation variables that occur in the heads of clauses in Π are called *intentional*; all other relation symbols and relation variables that occur in Π are called *extensional*. The *signature* of a general Datalog program is the set of all extensional relation symbols and constant symbols that occur in the program.

One intentional relation variable of Π is distinguished to be the *goal variable.* $\square$

To define the semantics, let Π be a general Datalog program, $X_1, \ldots, X_m$ the intentional relation variables of Π, and X_1 the goal variable.
For brevity we assume that for each $j \leq m$ there is a tuple $\bar{x}_j$ of distinct individual variables such that whenever X_j occurs in the head of a clause of Π this clause is of the form $X_j\bar{x}_j \leftarrow \ \ldots$.
We let

$$\varphi_j(\bar{x}_j) := \bigvee \{\exists \bar{y}(\gamma_1 \wedge \ldots \wedge \gamma_l) \mid X_j\bar{x}_j \leftarrow \gamma_1, \ldots, \gamma_l \in \Pi, \{\bar{y}\} = \mathrm{free}(\gamma_1 \wedge \ldots \wedge \gamma_l) \setminus \{\bar{x}_j\}\}.$$

We define Π to be equivalent to the simultaneous fixed–point formula

$$\varphi_\Pi := [\mathrm{FP}_{\bar{x}_1, X_1, \ldots, \bar{x}_m, X_m}\ \varphi_1, \ldots, \varphi_m]\bar{x}\ ,$$

i.e. for each interpretation $\mathfrak{I}$ we let $\mathfrak{I} \models \Pi \ :\Longleftrightarrow\ \mathfrak{I} \models \varphi_\Pi$.[6]
To be able to define Boolean queries even by Datalog programs not containing any constant symbols we allow 0–ary relation variables in general Datalog programs.
We let **P–Datalog** be the logic whose formulae are the general Datalog programs and whith the semantics defined as shown. Thus obviously we have **P–Datalog**$\subseteq$ **PFP**. In fact equality holds, a theorem which is essentially due to Abiteboul and Vianu [AV91]. In the present form it can be found in [EF95].

There are other *datalogics* (i.e. logics whose formulae are general Datalog programs), actually **P–Datalog** is a generalization of the original **Datalog** which was introduced only quite recently ([AV91, EF95]).

- **Datalog** is the sublogic of **P–Datalog** whose formulae are *(pure) Datalog programs*, i.e. general Datalog program whithout any negation symbols. It can be seen that a query is definable by a pure Datalog program if and only if it is definable by a **pos$\exists$LFP**–formula.
- By allowing negation symbols only in front of extensional relation symbols we obtain **Datalog$_\neg$**, which corresponds to **$\exists$LFP**.
- There is also a datalogic **I–Datalog** corresponding to **IFP** (and hence **LFP**). As for **P–Datalog**, its formulae are all general Datalog programs, whereas the semantics is defined by using $\models_I$ instead of $\models$. Again, a theorem of Abiteboul and Vianu [AV91] shows that **I–Datalog=IFP**.

The *arity* of a general Datalog program is defined to be the maximal arity of its intentional relation variables. Denoting the k–ary fragment of a datalogic by the superscript ^{k} we obviously have

$$\mathbf{Datalog}^k \subseteq \mathbf{Datalog}_\neg{}^k \subseteq \mathbf{I\text{–}Datalog}^k \subseteq \mathbf{P\text{–}Datalog}^k \subseteq \mathbf{s\text{–}PFP}^k$$

for each $k \geq 0$.

[6] This is just the usual fixed–point semantics for Datalog.

Since we cannot define a k–ary query by a $(k-1)$–ary Datalog program (because the goal variable is at most $(k-1)$–ary) it is reasonable to say that the arity hierarchy of a datalogic is strict if, for each $k \geq 1$, there is a Boolean query which is definable in the k–ary fragment, but not in the $(k-1)$–ary fragment.[7]

Thus, to be able to apply the main lemma to the datalogics we have to modify the query we consider. We augment our signature by constants $\underline{c}_1, \ldots, \underline{c}_k$, $\underline{d}_1, \ldots, \underline{d}_k$ and define ρ to be the Boolean query

$$\textit{There exists an } \mathrm{EDGE}_k\textit{–path from } \overset{k}{\underline{c}} \textit{ to } \overset{k}{\underline{d}}.$$

It is defined by the following k–ary (pure) Datalog program (with 0–ary goal variable Z):

$$\begin{array}{rcl}
Y\overset{k}{y} & \leftarrow & E\underline{c}_1 y_1, E\underline{c}_1 y_2, \ldots, E\underline{c}_k y_k, E\underline{c}_1\underline{c}_2, E\underline{c}_1\underline{c}_3, \ldots, E\underline{c}_{k-1}\underline{c}_k, \\
 & & Ey_1y_2, Ey_1y_3, \ldots, Ey_{k-1}y_k \\
Y\overset{k}{y} & \leftarrow & Y\overset{k}{x}, Ex_1y_1, Ex_1y_2, \ldots, Ex_ky_k, Ey_1y_2, Ey_1y_3, \ldots, Ey_{k-1}y_k \\
Z & \leftarrow & Y\overset{k}{\underline{d}}
\end{array}$$

Proving the main lemma we have in fact shown that for each $r \geq 1$ there are $\{E, \underline{c}_1, \ldots, \underline{c}_k, \underline{d}_1, \ldots, \underline{d}_k\}$–structures $\mathfrak{G}_1$ and $\mathfrak{G}_2$ which cannot distinguished by an **s–PFP**$^{k-1}$–formula of quantifier–rank $\leq r$ such that ρ holds in $\mathfrak{G}_1$, but not in $\mathfrak{G}_2$. This implies **Datalog**$^k \not\subseteq$ **s–PFP**$^{k-1}$.

However, we would prefer to have a pure graph query (without any constants) to obtain the result. We will take care of this in the next paragraph, where we show that we can modify the structures $\mathfrak{G}_1$ and $\mathfrak{G}_2$ in a way that the constants are definable, even by an existential positive first–order formula and hence by a pure Datalog program. This gives us the desired graph query, and we have:

4.2 Corollary. *The arity hierarchies of* **Datalog**, **Datalog**$_\neg$, **I–Datalog**, *and* **P–Datalog** *are strict (even on the class of graphs).*

[7] To circumvent this problem, instead of the arity Afrati and Cosmadakis [AC89] consider the *right hand side–width (rhs–width)* of a Datalog program which is defined to be the maximal arity of the intentional variables that occur in the body of a clause. The advantage is that it is possible to define queries of arbitrary arity by programs of bounded rhs–width.

Note that every Datalog program of rhs–width $\leq k$ is equivalent to a program where all intensional variables except the goal–variable have an arity $\leq k$. Thus a Datalog program that defines a Boolean query is of arity $\leq k$ if and only if it is of rhs–width $\leq k$. Since we are going to work with Boolean queries our results also hold for "rhs–width" instead of "arity".

(We prefer our notion of arity because it immediately implies **P–Datalog**$^k \subseteq$ **s–PFP**^{k}.)

4.1 Eliminating the Constants

We want to show how we can extend our $\{E, \underline{c}_1, \ldots, \underline{c}_k, \underline{d}_1, \ldots, \underline{d}_k\}$–structures $\mathfrak{G}_1$ and $\mathfrak{G}_2$ in order to make the constant–tuples $\overset{k}{\underline{c}}$ and $\overset{k}{\underline{d}}$ definable. This is rather simple in most cases, only for pure Datalog there occur some difficulties.

We will first sketch how to make the tuples definable by existential first–order formulae which suffices for all results except those incorporating pure Datalog.
Then we turn to the more involved question how to make them definable by existential first–order formulae without any negation symbols (hence by pure Datalog programs).

Consider for example the structure $\mathfrak{G}_1 = (\mathfrak{G}, \overset{k}{\boldsymbol{c}}, \overset{k}{\boldsymbol{d}})$. We are going to define a new structure $\mathfrak{G}_1^*$ that still has all the properties we need, but instead of the constants definable elements.
Remember that by Lemma 2.10 in $\mathfrak{G} = \mathfrak{G}(k, n)$ there are no edges between two elements of the same row and column. Since there are only edges between elements of the same row or neighboured rows this implies that $\mathfrak{G}$ contains no cliques of size $\geq 2k + 1$.
We augment $\mathfrak{G}$ by cliques of size $(2k+3), (2k+4), \ldots, (4k+2)$ (clique–sizes $(2k+1)$ and $(2k+2)$ are reserved for the next section) and add edges between $\boldsymbol{c}_1$ and a vertex of the $(2k+3)$–clique, $\boldsymbol{c}_2$ and a vertex of the $(2k+4)$–clique, ... , $\boldsymbol{c}_k$ and a vertex of the $(3k+2)$–clique, $\boldsymbol{d}_1$ and a vertex of the $(3k+3)$–clique, ... , $\boldsymbol{d}_k$ and a vertex of the $(4k+2)$–clique.

This suffices to make $\overset{k}{\boldsymbol{c}}$ and $\overset{k}{\boldsymbol{d}}$ definable: $\boldsymbol{d}_k$ is the unique vertex that is connected with a $(4k+2)$–clique in which it is not contained, $\boldsymbol{d}_{k-1}$ is the unique vertex distinct from $\boldsymbol{d}_k$ that is connected with a $(4k+1)$–clique in which it is not contained, and so on.
This can easily be formulated in existential first–order logic.

However, if we have no negation symbols available it is not so easy. For example, without negation symbols we cannot say that $\boldsymbol{d}_k$ is not contained in the $(4k+2)$–clique it is connected with or that $\boldsymbol{d}_{k-1}$ is distinct from $\boldsymbol{d}_k$.

For a graph $\mathfrak{A}$ with universe $A = \{a_1, \ldots, a_n\}$ we let

$$\Phi_{\mathfrak{A}} := \bigwedge \{Ex_i x_j \mid E^{\mathfrak{A}} a_i a_j\}.$$

Observe that for each graph $\mathfrak{B}$ we have $\mathfrak{B} \models \exists x_1 \ldots \exists x_n \Phi_{\mathfrak{A}}$ if and only if there exists a homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$.

4.3 Lemma. *There exist graphs $\mathfrak{C}_1, \ldots, \mathfrak{C}_{2k}$ such that for all $i \neq j \leq 2k$ we have:*

- *There exists no homomorphism from $\mathfrak{C}_i$ to $\mathfrak{C}_j$.*
- *$\mathfrak{C}_i$ is connected.*
- *Each vertex of $\mathfrak{C}_i$ is contained in a $(2k+3)$–clique.*

Once we have seen this the rest is easy: Let $\tilde{\mathfrak{G}}_1^*$ be the disjoint union of $\mathfrak{G}_1$, $\mathfrak{C}_1$, ..., $\mathfrak{C}_{2k}$ with additional edges between $\boldsymbol{c}_1$ and a vertex of $\mathfrak{C}_1$, $\boldsymbol{c}_2$ and a vertex of $\mathfrak{C}_2$, ... , $\boldsymbol{c}_k$ and a vertex of $\mathfrak{C}_k$, $\boldsymbol{d}_1$ and a vertex of $\mathfrak{C}_{k+1}$, ... , $\boldsymbol{d}_k$ and a vertex of $\mathfrak{C}_{2k}$.

For each $i \leq 2k$ let

$$\Phi_i(x) := \exists x_1 \dots \exists x_{n_i}(\Phi_{\mathfrak{C}_i} \wedge \bigvee_{i=1}^{n_i} Exx_i)$$

(where n_i is the size of $\mathfrak{C}_i$). Φ_i is a positive existential first–order formula that says

x is connected with a homomorphic image of $\mathfrak{C}_i$.

Then obviously we have $\tilde{\mathfrak{G}}_1^* \models \Phi_i[\boldsymbol{c}_i]$ (for each $i \leq k$). Moreover, there is no vertex $\boldsymbol{a} \in G \setminus \{\boldsymbol{c}_i\}$ [8] such that $\tilde{\mathfrak{G}}_1^* \models \Phi_i[\boldsymbol{a}]$ for the following reasons:

- There is no homomorphic image of $\mathfrak{C}_i$ in $\tilde{\mathfrak{G}}_1^*$ that contains vertices of $\mathfrak{G}_1$ because each vertex of $\mathfrak{C}_i$ is part of a $(2k+3)$–clique.
- There is no homomorphic image of $\mathfrak{C}_i$ in $\tilde{\mathfrak{G}}_1^*$ that contains vertices of distinct graphs $\mathfrak{C}_j$, $\mathfrak{C}_{j'}$ since $\mathfrak{C}_i$ is connected.
- There is no homomorphic image of $\mathfrak{C}_i$ in $\tilde{\mathfrak{G}}_1^*$ that is contained in $\mathfrak{C}_j$ for a $j \neq i$ since there is no homomorphism from $\mathfrak{C}_i$ to $\mathfrak{C}_j$.

Hence all homomorphic images of $\mathfrak{C}_i$ in $\tilde{\mathfrak{G}}_1^*$ are subgraphs of $\mathfrak{C}_i$ and $\boldsymbol{c}_i$ is the only vertex in G that is connected with $\mathfrak{C}_i$.
Similarly it can be seen that $\boldsymbol{d}_i$ is the unique vertex of G such that $\tilde{\mathfrak{G}}_1^* \models \Phi_{k+i}[\boldsymbol{d}_i]$ (for each $i \leq k$).

Hence $\overset{k}{\boldsymbol{c}}$ is the unique tuple in $\tilde{\mathfrak{G}}_1^*$ that satisfies

$$\mathrm{C}(\overset{k}{x}) := \bigwedge_{i=1}^{k} \Phi_i(x_i) \wedge \bigwedge_{i \neq j \leq k} Ex_i x_j$$

and $\overset{k}{\boldsymbol{d}}$ is the unique tuple in $\tilde{\mathfrak{G}}_1^*$ that satisfies

$$\mathrm{D}(\overset{k}{x}) := \bigwedge_{i=1}^{k} \Phi_{k+i}(x_i) \wedge \bigwedge_{i \neq j \leq k} Ex_i x_j \ .$$

Similarly we can define a structure $\tilde{\mathfrak{G}}_2^*$ where $\overset{k}{\boldsymbol{c}}$ and $\boldsymbol{e}(\overset{k}{\boldsymbol{d}})$ are the unique tuples satisfying $\mathrm{C}(\overset{k}{x})$ and $\mathrm{D}(\overset{k}{x})$ respectively.

To prove the lemma we make use of a theorem of Erdös [Erd59].
The *girth* of a graph $\mathfrak{C}$ is the size of the smallest cycle. The *chromatic number* $\mathrm{chr}(\mathfrak{C})$ is the minimal number of colours that is needed to colour the vertices of $\mathfrak{C}$ in a way that adjacent vertices have distinct colours.

[8] Recall that G denotes the universe of structure $\mathfrak{G}$ hence of the structure $\mathfrak{G}_1$ we started with.

4.4 Theorem. **[Erd59]** *For all $l, m \geq 1$ there exists a graph of girth $> l$ and chromatic number $> m$.*

A simple proof of this theorem can be found in [AS92] (on page 35). ■

Observe that if there exists a homomorphism from a graph $\mathfrak{C}$ to a graph $\mathfrak{D}$ we have $\mathrm{chr}(\mathfrak{D}) \geq \mathrm{chr}(\mathfrak{C})$ because each colouring of $\mathfrak{D}$ gives rise to a colouring of $\mathfrak{C}$ with the same colours
In particular this implies that if a graph $\mathfrak{C}$ has a chromatic number $> m$ there exists no homomorphism from $\mathfrak{C}$ to a graph of size $\leq m$.

4.5 Lemma. *For each $r \geq 1$ there exist graphs $\tilde{\mathfrak{C}}_1, \ldots, \tilde{\mathfrak{C}}_r$ such that for all $i \neq j \leq r$ we have:*

- *There exists no homomorphism from $\tilde{\mathfrak{C}}_i$ to $\tilde{\mathfrak{C}}_j$.*
- *$\tilde{\mathfrak{C}}_i$ is connected.*
- $\mathrm{girth}(\tilde{\mathfrak{C}}_i) > 3$.

Proof. We define the graphs inductively. Select a graph $\tilde{\mathfrak{C}}_1$ with $\mathrm{girth}(\tilde{\mathfrak{C}}_1) > 3$ and $\mathrm{chr}(\tilde{\mathfrak{C}}_1) > 2$ by Theorem 4.4. Without loss of generality we can assume that $\tilde{\mathfrak{C}}_1$ is connected (otherwise there exists a connected subgraph with the same properties). Let n_1 be the size of $\mathfrak{C}_1$.
Suppose now that $\tilde{\mathfrak{C}}_i$ is already defined and that n_i is the size of $\tilde{\mathfrak{C}}_i$. Let $\tilde{\mathfrak{C}}_{i+1}$ be a connected graph with $\mathrm{girth}(\tilde{\mathfrak{C}}_{i+1}) > n_i$ and $\mathrm{chr}(\tilde{\mathfrak{C}}_{i+1}) > n_i$. Note that this implies that the size n_{i+1} of $\tilde{\mathfrak{C}}_{i+1}$ is $> n_i$.

Having defined $\tilde{\mathfrak{C}}_1, \ldots, \tilde{\mathfrak{C}}_r$ like this it is clear from the observation above that there exists no homomorphism from $\tilde{\mathfrak{C}}_i$ to $\tilde{\mathfrak{C}}_j$ if $i > j$ (since $\mathrm{chr}(\tilde{\mathfrak{C}}_i) > n_j$).
Suppose for contradiction that there exists a homomorphism $h : \tilde{\mathfrak{C}}_j \to \tilde{\mathfrak{C}}_i$. Then the size of $h(\tilde{\mathfrak{C}}_j)$ is $\leq n_j$. Since $\mathrm{girth}(\tilde{\mathfrak{C}}_i) > n_{i-1} \geq n_j$ this means that $h(\tilde{\mathfrak{C}}_j)$ is cycle free. It is easy to see that a cycle free graph has a chromatic number ≤ 2.
Thus $\mathrm{chr}(\tilde{\mathfrak{C}}_j) \leq \mathrm{chr}(h(\tilde{\mathfrak{C}}_j)) \leq 2$, a contradiction. ■

Now it is easy to prove Lemma 4.3. Select $\tilde{\mathfrak{C}}_1, \ldots, \tilde{\mathfrak{C}}_{2k}$ by Lemma 4.5 and define for each $i \leq 2k$ a graph $\mathfrak{C}_i$ by

$$C_i := \tilde{C}_i \times \{1, \ldots, 2k+3\}$$

and

$$E^{\mathfrak{C}_i}(a, x)(b, y) \iff (a = b \wedge x \neq y) \vee E^{\tilde{\mathfrak{C}}_i} ab\,.$$

Then obviously each vertex of $\mathfrak{C}_i$ is contained in a $(2k+3)$–clique and $\mathfrak{C}_i$ is connected.

Suppose that there exists a homomorphism $H : \mathfrak{C}_i \to \mathfrak{C}_j$ for some $i \neq j$. We are going to derive a homomorphism $h : \tilde{\mathfrak{C}}_i \to \tilde{\mathfrak{C}}_j$ which yields a contradiction.
Let $a, b \in \tilde{\mathfrak{C}}_i$ such that $E^{\tilde{\mathfrak{C}}_i} ab$. Then

$$\{(a, z) \mid z \leq 2k+3\} \cup \{(b, z) \mid z \leq 2k+3\}$$

form a $(4k+6)$–clique in $\mathfrak{C}_i$ hence their images under H form a $(4k+6)$–clique in $\mathfrak{C}_j$. Since $\tilde{\mathfrak{C}}_j$ is triangle–free there exist $a', b' \in \tilde{\mathfrak{C}}_j$ such that $E^{\tilde{\mathfrak{C}}_j} a'b'$ and

$$\begin{aligned} & H\Big(\{(a,z) \mid z \leq 2k+3\} \cup \{(b,z) \mid z \leq 2k+3\}\Big) \\ & \quad = \quad \{(a',z) \mid z \leq 2k+3\} \cup \{(b',z) \mid z \leq 2k+3\}. \end{aligned}$$

We can assume without loss of generality that there exists a $z' \leq 2k+3$ such that $H((a,1)) = (a',z')$ (otherwise we just swap a' and b'). Furthermore, we can assume that for some $y' \leq 2k+3$ we have $H((b,1)) = (b',y')$. (Since $H((a,1)) = (a',z')$, there must exist some $y, y' \leq 2k+3$ such that $H((b,y)) = (b',y')$. Thus we can swap (b,y) and $(b,1)$ to achieve $H((b,1)) = (b',y')$.) We start with the definition of the homomorphism $h : \tilde{\mathfrak{C}}_i \to \tilde{\mathfrak{C}}_j$ by letting $h(a) = a'$ and $h(b) = b'$. If there exists a $c \in \tilde{\mathfrak{C}}_i$ such that $E^{\tilde{\mathfrak{C}}_i} bc$ then there exists a $c' \in \tilde{\mathfrak{C}}_j$ such that $E^{\tilde{\mathfrak{C}}_j} b'c'$ and

$$\begin{aligned} & H\Big(\{(b,z) \mid z \leq 2k+3\} \cup \{(c,z) \mid z \leq 2k+3\}\Big) \\ & \quad = \quad \{(b',z) \mid z \leq 2k+3\} \cup \{(c',z) \mid z \leq 2k+3\}. \end{aligned}$$

(note the b' remains the same here). Again we can assume without loss of generality that $H((c,1)) = (c',x')$ for some $x' \leq 2k+3$ since we already know that $H((b,1)) = (b',y')$; and we let $h(c) = c'$. We continue in this manner to define h on the connected component of a, and since $\tilde{\mathfrak{C}}_i$ is connected this yields the complete definition of h.
Hence Lemma 4.3 is proved. ∎

Having seen this gives us the reassuring feeling that we can continue using the constants.

5 Implicit Definability

5.1 Definition. Let σ be a signature and **L** a logic.
(1) Let $R_1, \ldots, R_m$ be relation symbols not contained in σ. A sentence

$$\varphi \in \mathrm{L}[\sigma \cup \{R_1, \ldots, R_m\}]$$

defines $R_1, \ldots, R_m$ *implicitly* if for each σ–structure there exists exactly one expansion to a $(\sigma \cup \{R_1, \ldots, R_m\})$–structure which is a model of φ.[9]

(2) A σ–query ρ of arity ≥ 1 is *implicitly definable in* **L** if there exist relation symbols $\{R_1, \ldots, R_m\}$ not contained in σ (where R_1 has the same arity as ρ) and a sentence $\varphi \in L[\sigma \cup \{R_1, \ldots, R_m\}]$ such that

- φ defines $R_1, \ldots, R_m$ implicitly.
- For each $(\sigma \cup \{R_1, \ldots, R_m\})$–structure $\mathfrak{A} \models \varphi$ we have $R_1^{\mathfrak{A}} = \rho(\mathfrak{A})$.

[9] Sometimes the notion given here is called "strongly" implicit definition to distinguish it from the other natural way to define (weakly) implicit definitions by replacing "exactly one expansion" by "at most one expansion".

A Boolean query ρ is *implicitly definable in* $\mathbf{L}$ if there exists an implicitly definable query ρ' (of arity ≥ 1) such that for all structures $\mathfrak{A}$ we have $\rho(\mathfrak{A}) = \text{TRUE}$ if and only if $\rho'(\mathfrak{A}) \neq \emptyset$.

$\mathbf{IMP(L)}$ denotes the class of all queries that are implicitly definable in $\mathbf{L}$.

(3) A query ρ is *k–ary* implicitly definable in $\mathbf{L}$ if the relation symbols $R_1, \ldots, R_m$ used for its implicit definition are all of arity $\leq k$.
$\mathbf{IMP(L)}^k$ denotes the class of all k–ary implicitly definable queries. □

The investigation of implicit definability on finite structures was initiated by Kolaitis [Kol90]. He conjectured that for each $k \geq 1$ there exists a k–ary query in

$$\mathbf{IMP}^{2k}(\mathbf{FO}) \setminus \mathbf{IMP}^{2k-1}(\mathbf{FO})$$

(and proved it for $k = 1$).

In the next paragraph we are going to define, for each $k \geq 1$, a Boolean graph–query in

$$\mathbf{IMP}^{k}(\mathbf{FO}) \setminus \mathbf{IMP}^{k-1}(\mathbf{s\text{–}PFP}^{k-1}).$$

Obviously this implies Kolaitis' conjecture.

5.2 Remark. Kolaitis' conjecture is also an easy consequence of a theorem of Ajtai [Ajt83] which implies that for a k–ary relation P the query

P contains an even number of elements

is not definable in $(k-1)$–ary Σ_1^1, even on the class of ordered structures (and hence not $(k-1)$–ary implicitly definable).
But it is easy to see that the query is k–ary implicitly definable in first–order logic on ordered structures.

However, this proof is a bit unsatisfying because it requires a new signature containing the k–ary relation symbol P for each k. □

5.1 The Hierarchy Theorem

5.3 Theorem. *For each $k \geq 2$ there is a Boolean graph–query in*

$$\mathbf{IMP}^{k}(\mathbf{FO}) \setminus \mathbf{IMP}^{k-1}(\mathbf{s\text{–}PFP}^{k-1}).$$

Proof. As usually we extend our signature by the constant–tuples $\overset{k}{\underline{c}}$ and $\overset{k}{\underline{d}}$ and consider the Boolean query

There exists an EDGE_k*–path from* $\overset{k}{\underline{c}}$ *to* $\overset{k}{\underline{d}}$.

(In Subsection 4.1 we have seen how we can avoid the use of constants.)

The main problem we have to deal with here is that there is no obvious way to define the query k–ary implicitly in first–order logic. To cope with it we extend our structures to make the order of the rows definable.

Let $k, n \geq 2$ and recall that the structure $\mathfrak{G} = \mathfrak{G}(k,n)$ does not contain any cliques of size $\geq 2k+1$. We are going to define a structure $\mathfrak{G}^* = \mathfrak{G}^*(k,n)$ which is an extension of $\mathfrak{G}$ and still has properties (1)–(4) described in Theorem 2.17. But moreover, the order of the rows is definable in $\mathfrak{G}^*$.
Therefore we augment each row of $\mathfrak{G}$ by a $(2k+1)$–clique and by a $(2k+2)$–clique. For each row I we select one vertex a_I that belongs to the $(2k+1)$–clique of this row and one vertex b_I that belongs to the $(2k+2)$–clique. We include an edge between each vertex of row I and a_I and an edge between each vertex of row I and b_I.
Furthermore, for all $I < J \leq n$ we include an edge between a_I and b_J. The resulting structure is $\mathfrak{G}^*$.
We extend the mapping row to the new vertices in the obvious way, i.e. the vertices of the cliques associated with row I belong to row I. We extend $\boldsymbol{c}$ to $\mathfrak{G}^*$ by letting it be the identity on the new vertices.
It is not difficult to define first–order formulae $\textsc{old}_k(x)$, $\prec_k (x,y)$, $\equiv_k (x,y)$ and $\textsc{Succ}_k(x,y)$ such that for all $\boldsymbol{a}, \boldsymbol{b} \in G^*$:

$$\begin{array}{lcl} \mathfrak{G}^* \models \textsc{old}_k[\boldsymbol{a}] & \iff & \boldsymbol{a} \in G \\ \mathfrak{G}^* \models \boldsymbol{a} \prec_k \boldsymbol{b} & \iff & \mathrm{row}(\boldsymbol{a}) < \mathrm{row}(\boldsymbol{b}) \\ \mathfrak{G}^* \models \boldsymbol{a} \equiv_k \boldsymbol{b} & \iff & \mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b}) \\ \mathfrak{G}^* \models \textsc{Succ}_k[\boldsymbol{a}, \boldsymbol{b}] & \iff & \mathrm{row}(\boldsymbol{a}) + 1 = \mathrm{row}(\boldsymbol{b}). \end{array}$$

Finally, we define the formula

$$\text{O–}\textsc{Edge}_k(\overset{k}{x}, \overset{k}{y}) := \textsc{Edge}_k(\overset{k}{x}, \overset{k}{y}) \wedge \bigwedge_{i=1}^{k} (\textsc{old}_k(x_i) \wedge \textsc{old}_k(y_i)) \wedge \bigwedge_{i,j \leq k} \textsc{Succ}_k(x_i, y_j)$$

that we consider instead of $\textsc{Edge}_k$ from now on.

Let Φ_k be the first–order formula saying

- $\equiv_k$ is an equivalence relation.
- $\prec_k$ is a linear order of the $\equiv_k$–classes.
- $\textsc{Succ}_k$ is the successor–relation that belongs to $\prec_k$.

Calling elements $\boldsymbol{a} \in G^*$ with $\mathfrak{G}^* \models \textsc{Old}_k[\boldsymbol{a}]$ *old elements* and the rest *new elements*, our structure $\mathfrak{G}^*$ has the following properties:

(0) $\mathfrak{G}^* \models \Phi_k$

(1) There exists a mapping row : $G^* \longrightarrow \{1,\dots,n\}$ such that

$$\mathrm{row}(\boldsymbol{a}) < \mathrm{row}(\boldsymbol{b}) \quad \Longleftrightarrow \quad \mathfrak{G}^* \models \boldsymbol{a} \prec_k \boldsymbol{b},$$

and for all old elements $\boldsymbol{a}, \boldsymbol{b}$, $E^{\mathfrak{G}^*}\boldsymbol{ab}$ implies $|\mathrm{row}(\boldsymbol{b}) - \mathrm{row}(\boldsymbol{a})| \leq 1$.

(2) There exists an automorphism $\text{\ae}$ of $\mathfrak{G}^*$ which is self inverse and preserves the rows. Furthermore, $\text{\ae}$ is the identity on the new elements.

(3) There exist tuples $\overset{k}{\boldsymbol{c}}, \overset{k}{\boldsymbol{d}}$ of old elements in the first and last row, respectively, such that $\overset{k}{\boldsymbol{c}}\overset{k}{\boldsymbol{d}}$ belongs to the transitive closure of O–$\mathrm{E{\scriptsize DGE}}_k$, whereas $\overset{k}{\boldsymbol{c}}\text{\ae}(\overset{k}{\boldsymbol{d}})$ does not.

(4) For all $\boldsymbol{b}_1,\dots,\boldsymbol{b}_{k-1} \in G^*$ there exists an automorphism f of $\mathfrak{G}^*$ which is self-inverse, preserves the rows, and maps $\overset{k-1}{\boldsymbol{b}}$ to $\text{\ae}(\overset{k-1}{\boldsymbol{b}})$, but leaves all elements in rows of distance > 1 from $\mathrm{row}(\boldsymbol{b}_1),\dots,\mathrm{row}(\boldsymbol{b}_{k-1})$ fixed.

We let $\mathfrak{G}_1^*$ and $\mathfrak{G}_2^*$ be the $\{E, \underline{c}_1,\dots,\underline{d}_k\}$–structures $(\mathfrak{G}^*, \overset{k}{\boldsymbol{c}}, \overset{k}{\boldsymbol{d}})$ and $(\mathfrak{G}^*, \overset{k}{\boldsymbol{c}}, \text{\ae}(\overset{k}{\boldsymbol{d}}))$ respectively, where $\overset{k}{\boldsymbol{c}}$ and $\overset{k}{\boldsymbol{d}}$ are chosen according to property (3) of $\mathfrak{G}^*$.
Let ρ be the k–ary query that defines the O–$\mathrm{E{\scriptsize DGE}}_k$–connected component of $\overset{k}{\underline{c}}$. It is definable by the k–ary **LFP**–formula

$$[\mathrm{FP}_{\overset{k}{x},X}\, \overset{k}{x} = \overset{k}{\underline{c}} \vee \exists \overset{k}{y} \in X \text{ O–}\mathrm{E{\scriptsize DGE}}_k(\overset{k}{y}, \overset{k}{x})]\overset{k}{x}.$$ [10]

Thus by a result of Kolaitis [Kol90] (which says that $\mathbf{LFP}^k \subseteq \mathbf{IMP}^{2k}(\mathbf{FO})$ for each $k \geq 1$) it is definable in $\mathbf{IMP}^{2k}(\mathbf{FO})$.
However, on the class of models of Φ_k it is k–ary implicitly definable by the following first–order formula:

$$\begin{array}{rcl} \varphi(R) & = & R\overset{k}{\underline{c}} \\ & \wedge & \forall \overset{k}{x}, \overset{k}{y}((R\overset{k}{x} \wedge \text{O–}\mathrm{E{\scriptsize DGE}}_k(\overset{k}{x}, \overset{k}{y})) \rightarrow R\overset{k}{y}) \\ & \wedge & \forall \overset{k}{x}(R\overset{k}{x} \rightarrow (\overset{k}{x} = \overset{k}{\underline{c}} \vee \exists \overset{k}{y}\,(R\overset{k}{y} \wedge \text{O–}\mathrm{E{\scriptsize DGE}}_k(\overset{k}{y}, \overset{k}{x})))) \ . \end{array}$$

Because let $\mathfrak{A} \models \Phi_k$ and $C \subseteq A^k$ the (uniquely determined) O–$\mathrm{E{\scriptsize DGE}}_k$–connected component of $\overset{k}{\underline{c}}$. Then clearly $\mathfrak{A} \models \varphi[C]$.
Let now $B \subseteq A^k$ such that $\mathfrak{A} \models \varphi[B]$. By the first and the second conjunct we have $C \subseteq B$. Moreover, by the third conjunct, for each $\overset{k}{a} \in B$ there exists a sequence $\overset{k}{a}_1 = \overset{k}{a}, \overset{k}{a}_2, \dots$ such that for each $i \geq 1$

$$\mathfrak{G}^* \models \text{O–}\mathrm{E{\scriptsize DGE}}_k[\overset{k}{a}_{i+1}, \overset{k}{a}_i] \quad \text{or} \quad \overset{k}{a}_i = \overset{k}{\underline{c}}.$$

[10] Since we are always a bit vague with the notion "connected component" the reader may take this formula as a definition of the query.

We are finished if the latter case occurs. Since the former case implies $\mathfrak{G}^* \models \overset{k}{a}_{i+1} \prec_k \overset{k}{a}_i$ (by the definition of the formula O–$\textsc{Edge}_k$), and since $\prec_k$ is well–founded, the sequence cannot be infinitely long thus eventually the latter case must occur.

We consider the query ρ^* that says:

If Φ_k holds then there exists an O–$\textsc{Edge}_k$*–path from* $\overset{k}{\underline{c}}$ *to* $\overset{k}{\underline{d}}$.

It can be implicitly defined by the formula

$$\xi(P,R) := \Big(\neg\Phi_k \wedge \forall x Px \wedge \forall \overset{k}{x} R \overset{k}{x}\Big) \\ \vee \big(\Phi_k \wedge \varphi(R) \wedge ((R\overset{k}{\underline{d}} \wedge \forall x Px) \vee (\neg R\overset{k}{\underline{d}} \wedge \forall x \neg Px))\big).$$

So we shall prove that ρ^* is not $(k-1)$–ary implicitly definable in $\mathbf{s}$–$\mathbf{PFP}^{k-1}$. Basically, the proof that it is not (explicitly) definable in $\mathbf{s}$–$\mathbf{PFP}^{k-1}$ just goes through. We explore the fact that implicitly defined relations must be closed under automorphisms, and that for any $(k-1)$–tuple there is an automorphism of our structure mapping it to the other component. Hence $(k-1)$–ary implicitly defined relations cannot distinguish between the components.

To carry out this argument in detail, suppose for contradiction that ρ^* is in $\mathbf{IMP}^{k-1}(\mathbf{s}$–$\mathbf{PFP}^{k-1})$. Then there exists an $\mathbf{s}$–$\mathbf{PFP}^{k-1}$ sentence $\psi(R_1,\dots,R_m)$ of some signature $\{E,\overset{k}{\underline{c}},\overset{k}{\underline{d}},R_1,\dots,R_m\}$ (where $R_1,\dots,R_m$ are relation symbols of arity $\leq (k-1)$) defining $R_1,\dots,R_m$ implicitly such that for all $\{E,\overset{k}{\underline{c}},\overset{k}{\underline{d}},R_1,\dots,R_m\}$–structures $\mathfrak{A}$ we have

$$\rho^*(\mathfrak{A}) = \textsc{True} \iff R_1^{\mathfrak{A}} \neq \emptyset.$$

Without loss of generality we assume that the arity of each R_i $(i \leq m)$ is $(k-1)$.

Let r be the quantifier–rank of ψ and $n = 2^{r+1}+1$. Consider the structures $\mathfrak{G}^* = \mathfrak{G}^*(k,n)$, $\mathfrak{G}_1^* = \mathfrak{G}_1^*(k,n)$, and $\mathfrak{G}_2^* = \mathfrak{G}_2^*(k,n)$ defined above. For each $i \leq m$ let $A_i := R_i^{\mathfrak{G}_1^*}$.
We define a bijection $f : (G^*)^{k-1} \longrightarrow (G^*)^{k-1}$ as follows:
Let $\overset{k-1}{\boldsymbol{a}} \in (G^*)^{k-1}$. Then there exist uniquely determined sets $A, A' \subseteq G^*$ such that

- $A \cup A' = \{\boldsymbol{a}_1,\dots,\boldsymbol{a}_{k-1},\boldsymbol{c}_1,\dots,\boldsymbol{c}_k,\boldsymbol{d}_1,\dots,\boldsymbol{d}_k\}$
- $\forall \boldsymbol{a} \in A, \boldsymbol{a}' \in A' : \mathrm{row}(\boldsymbol{a}) < \mathrm{row}(\boldsymbol{a}') - 1$
- A is chosen minimal with respect to the first two conditions.

We let $\boldsymbol{b}_i = \boldsymbol{a}_i$ if $\boldsymbol{a}_i \in A$ and $\boldsymbol{b}_i = \boldsymbol{c}(\boldsymbol{a}_i)$ if $\boldsymbol{a}_i \in A'$ (for each $i \leq k-1$). Now we define $f(\overset{k-1}{\boldsymbol{a}}) = \overset{k-1}{\boldsymbol{b}}$. It is not hard to see that $f^{-1} = f$ thus f is indeed a bijection.

Now $B_1,\dots,B_m \subseteq (G^*)^{k-1}$ are defined by $B_i = f(A_i)$.

We are going to prove (making use of the $(k-1)$–ary r–pebble game again) that the same $\mathbf{s}$–$\mathbf{PFP}^{k-1}$–sentences of quantifier–rank r hold in $(\mathfrak{G}_1^*, A_1, \ldots, A_m)$ and $(\mathfrak{G}_2^*, B_1, \ldots, B_m)$. In particular $\mathfrak{G}_2^* \models \psi(B_1, \ldots, B_m)$, thus for each $i \leq m$ we have $B_i = R_i^{\mathfrak{G}_2^*}$. But then, since $A_1 \neq \emptyset$ we have $R_1^{\mathfrak{G}_2^*} = B_1 \neq \emptyset$ thus $\rho^*(\mathfrak{G}_2^*) =$TRUE which is a contradiction.

We say that two tuples $\overset{k-1}{\boldsymbol{a}}, \overset{k-1}{\boldsymbol{b}} \in (G^*)^{k-1}$ are *similar* (and write $\overset{k-1}{\boldsymbol{a}} \approx \overset{k-1}{\boldsymbol{b}}$) if there exist sets $A, A', A'' \subseteq G^*$ such that:

- $A \cup A' \cup A'' = \{\boldsymbol{a}_1, \ldots, \boldsymbol{a}_{k-1}, \boldsymbol{c}_1, \ldots, \boldsymbol{c}_k, \boldsymbol{d}_1, \ldots, \boldsymbol{d}_k\}$
- $\forall \boldsymbol{a} \in A, \boldsymbol{a}' \in A' : \mathrm{row}(\boldsymbol{a}) < \mathrm{row}(\boldsymbol{a}') - 1$
- $\forall \boldsymbol{a}' \in A', \boldsymbol{a}'' \in A'' : \mathrm{row}(\boldsymbol{a}') < \mathrm{row}(\boldsymbol{a}'') - 1$
- $\forall i \leq k-1 : \ \boldsymbol{b}_i = \begin{cases} \boldsymbol{a}_i & \text{if } \boldsymbol{a}_i \in A \cup A'' \\ \boldsymbol{c}(\boldsymbol{a}_i) & \text{if } \boldsymbol{a}_i \in A' \end{cases}$

5.4 Lemma. *Suppose $C \subset (G^*)^{k-1}$ is closed under automorphisms of $\mathfrak{G}_1$. Then C is closed under similarity, i.e. $\overset{k-1}{\boldsymbol{a}} \in C$ and $\overset{k-1}{\boldsymbol{a}} \approx \overset{k-1}{\boldsymbol{b}}$ implies $\overset{k-1}{\boldsymbol{b}} \in C$.*

Proof. Suppose $\overset{k-1}{\boldsymbol{a}} \approx \overset{k-1}{\boldsymbol{b}}$ via $A, A', A'' \subseteq G^*$. The claim follows from property (4) of structure $\mathfrak{G}^*$ applied to a tuple formed of the at most $(k-1)$–elements of A'. ∎

5.5 Corollary. *Let $A, A' \subset G^*$ such that $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_k \in A$, $\boldsymbol{d}_1, \ldots, \boldsymbol{d}_k \in A'$, and for all $\boldsymbol{a} \in A, \boldsymbol{a}' \in A' : \mathrm{row}(\boldsymbol{a}) < \mathrm{row}(\boldsymbol{a}') - 1$. For a tuple $\overset{k-1}{\boldsymbol{a}} \in (A \cup A')^{k-1}$ we define $\overset{k-1}{\boldsymbol{b}}$ by*

$$\boldsymbol{b}_i := \begin{cases} \boldsymbol{a}_i & \text{if } \boldsymbol{a}_i \in A \\ \boldsymbol{c}(\boldsymbol{a}_i) & \text{if } \boldsymbol{a}_i \in A' \end{cases} .$$

Then for each $i \leq m$ we have

$$\overset{k-1}{\boldsymbol{a}} \in A_i \quad \Longleftrightarrow \quad \overset{k-1}{\boldsymbol{b}} \in B_i .$$

Proof. Note first that the (implicitly defined) sets A_i $(i \leq m)$ are closed under automorphisms of $\mathfrak{G}_1^*$. Note next that $\overset{k-1}{\boldsymbol{a}} \approx f^{-1}(\overset{k-1}{\boldsymbol{b}})$. Hence for all $i \leq m$ we have

$$\overset{k-1}{\boldsymbol{a}} \in A_i \quad \Longleftrightarrow \quad f^{-1}(\overset{k-1}{\boldsymbol{b}}) \in A_i \quad \Longleftrightarrow \quad \overset{k-1}{\boldsymbol{b}} \in B_i .$$

∎

The proof of the theorem is now easily finished:

In Subsection 3.2 we proved that the duplicator wins the $(k-1)$–ary r–pebble game on $\mathfrak{G}_1$ and $\mathfrak{G}_2$ hence on $\mathfrak{G}_1^*$ and $\mathfrak{G}_2^*$ because we only needed properties (1)–(4) of $\mathfrak{G}$, and $\mathfrak{G}^*$ has got the same properties.
In fact we proved that she can preserve the following throughout the whole game:
Suppose $r' \leq r$ pairs of pebbles are on the board. Then there exist $M, N \in \{1, \ldots, n\}$ and an automorphism h of $\mathfrak{G}_1^*$ such that:

(i) $N - M \geq 2^{r+1-r'}$

(ii) For pairs $(\boldsymbol{a}, \boldsymbol{b})$ of corresponding elements we have:

$$\mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b}) \leq M \quad \text{or} \quad \mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b}) \geq N\ .$$

(iii) h is self–inverse and preserves the rows.

(iv) For each pair $(\boldsymbol{a}, \boldsymbol{b})$ of corresponding elements the following holds:

- If $\mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b}) \leq M$ then $h(\boldsymbol{a}) = \boldsymbol{b}$.
- If $\mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b}) \geq N$ then $\boldsymbol{c}(h(\boldsymbol{a})) = \boldsymbol{b}$.

This implies that the pairs of corresponding elements form a partial isomorphism not only between $\mathfrak{G}_1^*$ and $\mathfrak{G}_2^*$, but between $(\mathfrak{G}_1^*, A_1, \ldots, A_m)$ and $(\mathfrak{G}_2^*, B_1, \ldots, B_m)$.
Because let $A = \{\boldsymbol{a} \in G^* \mid \mathrm{row}(\boldsymbol{a}) \leq M\}$ and $A' = \{\boldsymbol{a} \in G^* \mid \mathrm{row}(\boldsymbol{a}) \geq N\}$ and define a partial mapping p on G^* by

$$p(\boldsymbol{a}) = \begin{cases} h(\boldsymbol{a}) & \text{if } \boldsymbol{a} \in A \\ \boldsymbol{c}(h(\boldsymbol{a})) & \text{if } \boldsymbol{a} \in A' \\ \text{undefined} & \text{otherwise} \end{cases}$$

Since $N - M > 2^{r-r'} \geq 1$ and $h, \boldsymbol{c}$ are automorphisms of $\mathfrak{G}^*$, p is a partial isomorphism between $\mathfrak{G}_1^*$ and $\mathfrak{G}_2^*$. Corollary 5.5 applied to the sets $h(A)$ and $h(A')$ together with the fact that the A_i are closed under automorphisms (in particular under the automorphism h of $\mathfrak{G}_1^*$) does the rest.

Thus the duplicator wins the the $(k-1)$–ary r–pebble game on $(\mathfrak{G}_1^*, A_1, \ldots, A_m)$ and $(\mathfrak{G}_2^*, B_1, \ldots, B_m)$ and Theorem 5.3 is proved. ∎

6 Transitive Closure Logics

All through this paper we have worked with a query which is the transitive closure of a first–order formula. *Transitive closure logic* **TC** is the extension of first–order logic with an operator defining such queries.
Formally, the class TC of *transitive closure formulae* is given by the usual first–order clauses and the new clause:

$$(\mathrm{TC}) \quad \frac{\varphi}{[\mathrm{TC}_{\overset{k}{x},\overset{k}{y}}\ \varphi]\overset{k}{u},\overset{k}{v}} \quad \text{where } k \geq 1 \text{ and } \overset{k}{u} \text{ and } \overset{k}{v} \text{ are } k\text{–tuples of terms.}$$

$[\mathrm{TC}_{\overset{k}{x},\overset{k}{y}} \ldots]$ is called a k–ary *transitive closure operator.*

The semantics of the new operator is defined by

$$\mathfrak{I} \models [\mathrm{TC}_{\overset{k}{x},\overset{k}{y}}\, \varphi]\overset{k}{u},\overset{k}{v} \quad \iff \quad \alpha(\overset{k}{u}\overset{k}{v}) \in \mathrm{TC}(\{\overset{k}{a}\overset{k}{b} \mid \mathfrak{I} \models \varphi[\overset{k}{a},\overset{k}{b}]\})$$

(for any **TC**–formula $[\mathrm{TC}_{\overset{k}{x},\overset{k}{y}}\, \varphi]\overset{k}{u},\overset{k}{v}$ and interpretation $\mathfrak{I} = (\mathfrak{A}, \alpha)$).

The *k–ary fragment* $\mathbf{TC}^k$ of **TC** is the sublogic whose formulae contain at most k–ary TC–operators. Note that our main lemma implies that the arity hierarchy of **TC** is strict. However, we must disappoint the reader who thinks we can stop with this because there is an important sublogic of **TC**, *deterministic transitive closure logic* **DTC**, for which the problem still has to be solved.
In the next subsection we introduce **DTC** and its k–ary fragments and show why we cannot just extend our main lemma to the statement that $\mathbf{DTC}^k \not\subseteq \mathbf{s}\text{–}\mathbf{PFP}^{k-1}$ for all $k \geq 1$: Actually **DTC** is already contained in $\mathbf{IFP}^2$. However, the arity hierarchy of **DTC** is strict, in Subsections 6.2–6.4 we show that we even have $\mathbf{DTC}^k \not\subseteq \mathbf{TC}^{k-1}$ for all $k \geq 1$; as usual this already holds on the class of graphs.

6.1 Deterministic Transitive Closure Logic

One reason for the interest in transitive closure logic is the fact (due to Immerman [Imm87]) that on ordered structures a query is **NLOGSPACE**–computable if and only if it is definable in **TC**. Immerman also showed that there is a natural sublogic of **TC**, *deterministic transitive closure logic* **DTC**, that captures **LOGSPACE** on ordered structures.

Instead of the transitive closure of a relation $R \in A^{2k}$, a DTC–operator defines its *deterministic transitive closure*

$$\mathrm{DTC}(R) = \Big\{\overset{k}{a}\overset{k}{b} \mid \exists n \geq 2, \overset{k}{a}_1 = \overset{k}{a}, \overset{k}{a}_2, \ldots, \overset{k}{a}_n = \overset{k}{b} \forall i < n : \ \big(R\overset{k}{a}_i\overset{k}{a}_{i+1} \wedge \forall \overset{k}{c}(R\overset{k}{a}_i\overset{k}{c} \rightarrow \overset{k}{c} = \overset{k}{a}_{i+1})\big)$$

A convenient way to introduce **DTC** is the following:
For each formula $\varphi(\bar{x}, \bar{y})$ we consider its *deterministic version*

$$\varphi_D(\bar{x}, \bar{y}) := \varphi(\bar{x}, \bar{y}) \wedge \forall \bar{z}(\varphi(\bar{x}, \bar{z}) \rightarrow \bar{z} = \bar{y}).$$

DTC is the sublogic of **TC** whose formulae are those **TC**–formulae where the (TC)–clause is only applied to formulae of the form φ_D. We usually write $[\mathrm{DTC}_{\bar{x},\bar{y}}\varphi]\bar{u}, \bar{v}$ instead of $[\mathrm{TC}_{\bar{x},\bar{y}}\varphi_D]\bar{u}, \bar{v}$.
Note that for any **DTC**–formula $[\mathrm{DTC}_{\overset{k}{x},\overset{k}{y}}\, \varphi]\overset{k}{u},\overset{k}{v}$ and interpretation $\mathfrak{I} = (\mathfrak{A}, \alpha)$ we have

$$\mathfrak{I} \models [\mathrm{DTC}_{\overset{k}{x},\overset{k}{y}}\, \varphi]\overset{k}{u},\overset{k}{v} \quad \iff \quad \alpha(\overset{k}{u}\overset{k}{v}) \in \mathrm{DTC}(\{\overset{k}{a}\overset{k}{b} \mid \mathfrak{I} \models \varphi[\overset{k}{a},\overset{k}{b}]\}).$$

Again we are interested in the *k–ary fragment* $\mathbf{DTC}^k$ which is defined in the obvious way. The reader may ask whether we can strengthen our main lemma in a way such that it implies $\mathbf{DTC}^k \not\subseteq \mathbf{s}\text{–}\mathbf{PFP}^{k-1}$. Unfortunately, we cannot:

6.1 Theorem. $\mathbf{DTC} \subseteq \mathbf{s\text{–}PFP}^1$

Proof. The idea is rather simple:
We can code a path $\overset{k}{a}_1, \ldots, \overset{k}{a}_n$ by sets X_{ij} $(1 \leq i \leq k, 1 \leq j \leq n)$ where $X_{ij} = \{a_{ji}\}$. If the path belongs to a deterministic transitive closure formula these sets can be defined to be the stages of a simultaneous fixed–point process. We need the path to be deterministic because otherwise we would get in trouble mixing up different successors of one tuple.
Since this proof is quite similar to the proof of Theorem 6.3 we skip the details here. ■

6.2 Corollary. $\mathbf{DTC} \subseteq \mathbf{PFP}^2$

Proof. Using standard techniques to eliminate simultaneous inductions it is not hard to see that $\mathbf{s\text{–}PFP}^k \subseteq \mathbf{PFP}^{k+1}$ for each $k \geq 0$. ■

The situation is not so easy for **IFP**, but we can still prove that **DTC** is contained in $\mathbf{IFP}^2$.
To see this we take two steps:
The first is to show that **DTC** is contained in $\mathbf{LFP}^2$ (hence $\mathbf{IFP}^2$) on ordered structures.[11]
Analysing this proof we will see that we only need a definable order of sufficient length to obtain the result. (Which length is "sufficient" depends on the length of the deterministic path we want to simulate.)
Then, in the second step, we are going to show that indeed an order of suffficient length is definable in $\mathbf{IFP}^2$.

6.3 Theorem. $\mathbf{DTC} \subseteq \mathbf{LFP}^2$ *on ordered structures.*

Proof. The problem we have here, in contrast to the **PFP**–case, is that we cannot remove tuples from the relation X defined in the fixed–point process of an **LFP**–formula. If we just let them in we have the problem of mixing up tuples that belong to different stages. To avoid this we label the (unique) tupel that is reached after s steps with the sth element of our order (for each s less than or equal to the size of the structure we are in).
Hence in structures of size n we have n labels (=elements), which allows us to encode deterministic paths of length up to n in the described way. However, a k–ary deterministic path might have length n^k so we have to iterate our construction. This leads to nested fixed–point operators.

For convenience we assume that the universe of an ordered structure is always an initial segment of $\mathbb{N}$. We can code the tuple $\overset{k}{x}$ which is the sth tuple of a deterministic path by the set $\{(s-1, 0, x_1), \ldots, (s-1, k-1, x_k)\}$. To make this binary we code the first two places into one by a suitable pairing function:

Let us fix $k \geq 1$. Without loss of generality we only consider structures of size $\geq k$.

[11] I.e. structures whose signature contains a fixed binary relation symbol $<$ which is always interpreted as a linear order of the universe.

We define the function $< \cdot,\cdot >: \mathbb{N} \times \{1,\dots,k\} \longrightarrow \mathbb{N}$ by $< s,l >= sk+l-1$. Since it is easy to define addition and multiplication in $\mathbf{LFP}^2$ on ordered structures (to define $x+y=z$ or $x\cdot y=z$, for each x we define a fixed–point process computing the values for y and z and using x as a parameter) we can find an $\mathbf{LFP}^2$–formulae $\psi_{<\cdot,\cdot>}(x,y,z)$ defining the function on ordered structures.

Let $\varphi(\overset{k}{x},\overset{k}{y})$ be an $\mathbf{LFP}^2$–formula. We shall define an $\mathbf{LFP}^2$–formula $\chi'(\overset{k}{t},\overset{k}{u})$ that is equivalent to

$$\chi(\overset{k}{t},\overset{k}{u}) := [\mathrm{DTC}_{\overset{k}{x},\overset{k}{y}}\varphi(\overset{k}{x},\overset{k}{y})]\overset{k}{t},\overset{k}{u}\ .$$

Without loss of generality we can assume that $\chi(\overset{k}{t},\overset{k}{u}) \models \overset{k}{t} \neq \overset{k}{u}$. If this is not the case, we can pass to the formula

$$\varphi_D(\overset{k}{t},\overset{k}{u}) \vee \exists \overset{k'}{t}\big(\varphi_D(\overset{k}{t},\overset{k'}{t}) \wedge [\mathrm{DTC}_{\overset{k}{x},\overset{k}{y}}\overset{k'}{t} \neq \overset{k}{u} \wedge \varphi(\overset{k}{x},\overset{k}{y})]\overset{k'}{t},\overset{k}{u}\big).$$

It can be easily checked that this formula is equivalent to χ; and its DTC–operator satisfies the desired condition.
Before we start with the definition of χ' we pass to the following formula:

$$\tilde{\chi}(\overset{k}{t},\overset{k}{u}) = \Big[\mathrm{DTC}_{\overset{k}{x},\overset{k}{y}}\underbrace{\overset{k}{t} \neq \overset{k}{u} \wedge \big((\overset{k}{x} \neq \overset{k}{u} \wedge \varphi(\overset{k}{x},\overset{k}{y})) \vee (\overset{k}{x} = \overset{k}{u} \wedge \overset{k}{y} = \overset{k}{u})\big)}_{\tilde{\varphi}(\overset{k}{x},\overset{k}{y})}\Big]\overset{k}{t},\overset{k}{u}.$$

Clearly $\tilde{\chi}$ is equivalent to χ (but only because we assumed $\chi(\overset{k}{t},\overset{k}{u}) \models \overset{k}{t} \neq \overset{k}{u}$). Observe furthermore that whenever $\chi(\overset{k}{t},\overset{k}{u})$ holds in a structure via a deterministic φ–path of length m then for any $m' \geq m$ there exists a deterministic $\tilde{\varphi}$–path of length m' from $\overset{k}{t}$ to $\overset{k}{u}$.
Thus working with $\tilde{\chi}$ gives us the possibility to use paths which are actually too long.

We now define an $\mathbf{LFP}^2$–formula $\tilde{\varphi}^{\frac{n}{\lfloor k \rfloor}}(\overset{k}{x},\overset{k}{y})$ saying

> *"There exists a deterministic $\tilde{\varphi}$–path of length $(\frac{n}{\lfloor k \rfloor})$ from $\overset{k}{x}$ to $\overset{k}{y}$."*

by:

$$\bigwedge_{i=1}^{k}\Big[\mathrm{FP}_{\overset{2}{z},Z}\quad \exists\overset{k}{w}\big(\tilde{\varphi}_D(\overset{k}{x},\overset{k}{w}) \wedge \bigvee_{j=1}^{k}(z_1 =< 1,j > \wedge z_2 = w_j)\big)$$

We start the path with the $\tilde{\varphi}_D$–successor of $\overset{k}{x}$.

$$\vee \exists\overset{k}{v}\exists\overset{k}{w}\exists t\Big(\bigwedge_{j=1}^{k} Z < t,j > v_j \wedge \tilde{\varphi}_D(\overset{k}{v},\overset{k}{w})$$
$$\wedge \bigvee_{j=1}^{k}(z_1 =< t+1,j > \wedge z_2 = w_j)\Big)\Big] < \tfrac{n}{\lfloor k \rfloor},i > y_i.$$

Then we walk along the $\tilde{\varphi}_D$–path.

(It is clear how this formula can be turned into a well–defined $\mathbf{LFP}^2$–formula using the formula $\psi_{<\cdot,\cdot>}(x,y,z)$, and explicit definitions of the numbers $1,\ldots,k$ and $\frac{n}{\lfloor k\rfloor}$.)
Note that the formula $\tilde{\varphi}^{\frac{n}{\lfloor k\rfloor}}(\overset{k}{x},\overset{k}{y})$ is deterministic, i.e. $\tilde{\varphi}^{\frac{n}{\lfloor k\rfloor}}(\overset{k}{x},\overset{k}{y})$ is equivalent to $\tilde{\varphi}_D^{\frac{n}{\lfloor k\rfloor}}(\overset{k}{x},\overset{k}{y})$. Thus we can apply the same construction to $\tilde{\varphi}^{\frac{n}{\lfloor k\rfloor}}(\overset{k}{x},\overset{k}{y})$ instead of $\tilde{\varphi}$. Iterating this, we let $\tilde{\varphi}^0 := \tilde{\varphi}$ and $\tilde{\varphi}^{i+1} := (\tilde{\varphi}^i)^{\frac{n}{\lfloor k\rfloor}}$ for each $i \geq 1$. It is easy to see that $\tilde{\varphi}^i$ is an LFP^2–formula that says:

"There exists a deterministic $\tilde{\varphi}$–path of length $(\frac{n}{\lfloor k\rfloor})^i$ from $\overset{k}{x}$ to $\overset{k}{y}$."

This completes the proof because we can always select an i such that $(\frac{n}{\lfloor k\rfloor})^i \geq n^k$ for every $n \geq 2$, and we have chosen $\tilde{\varphi}$ in a way that too long paths do not matter. ■

We say that a formula $\psi(x,y,\overset{k}{z})$ *defines an order with parameters* $\overset{k}{z}$ if for each structure $\mathfrak{A}$ and $\overset{k}{a} \in A$ the binary relation

$$\psi(\_\_,\_\_,\overset{k}{a})^{\mathfrak{A}} = \{ab \mid \mathfrak{A} \models \psi[a,b,\overset{k}{a}]\}$$

is a linear order of its support[12].

Similarly to Theorem 6.3 we can prove:

6.4 Lemma. *Consider a formula $\chi = [\mathrm{DTC}_{\overset{k}{x},\overset{k}{y}}\varphi]\overset{k}{u},\overset{k}{v}$ where $\varphi \in \mathbf{IFP}^2$ and suppose that there exists an $\mathbf{IFP}^2$–formula $\psi(x,y,\overset{k}{z})$ defining an order with parameters $\overset{k}{z}$ such that the deterministic φ–path starting at $\overset{k}{z}$ is always contained in the support of this order. More precisely, for each structure $\mathfrak{A}$ and $\overset{k}{a},\overset{k}{b} \in A$ we have*

$$\mathfrak{A} \models [\mathrm{DTC}_{\overset{k}{x},\overset{k}{y}}\varphi(\overset{k}{x},\overset{k}{y})]\overset{k}{a},\overset{k}{b} \quad \Longrightarrow \quad \overset{k}{b} \text{ is contained in the support of the order } \psi(\_\_,\_\_,\overset{k}{a})^{\mathfrak{A}}.$$

Then there exists an $\mathbf{IFP}^2$–formula that is equivalent to χ.

So we shall prove:

6.5 Lemma. *Let $\varphi(\overset{k}{x},\overset{k}{y}) \in \mathbf{IFP}^2$. Then there exists an $\mathbf{IFP}^2$–formula $\psi(x,y,\overset{k}{z})$ defining an order with parameters $\overset{k}{z}$ such that the deterministic φ–path starting at $\overset{k}{z}$ is always contained in the support of this order.*

Proof. The idea is to define inductively an order $<_{\varphi,\overset{k}{z}}$ as follows:

[12] The support of a relation R on a set A is the set of all elements of A that occur in a tuple $\bar{a}$ with $R\bar{a}$.

- We start with the order $z_1 <_{\varphi,\overset{k}{z}} \dots <_{\varphi,\overset{k}{z}} z_k$ (assuming for a moment that they are all distinct).
- The order defined up to now induces a (lexicographic) order of the k–tuples in its support S. We consider the first tuple $\overset{k}{x} \in S^k$ with respect to this order that has a φ_D–successor $\overset{k}{y} \notin S^k$. We extend our order by the first y_i not contained in S.

Let

$$\mathrm{INI}_{\overset{k}{z}}(x,y) = \bigvee_{i<j\leq k} \big(x = z_i \wedge \neg \bigvee_{i'=1}^{i-1} z_{i'} = z_i \wedge y = z_j \wedge \neg \bigvee_{j'=1}^{j-1} z_{j'} = z_j\big).$$

(This formula takes care of the first step of our definition, taking into account that the z_i may not all be distinct.)
Let $\prec$ be a binary relation variable and

$$\begin{aligned} \mathrm{SUPP}_{\prec}(x) &:= \exists w(x \prec w \vee w \prec x) \vee x = z_1 \\ \mathrm{LEX}_{\prec}(\overset{k}{x},\overset{k}{y}) &:= \bigvee_{i=1}^{k} (x_i \prec y_i \wedge \bigwedge_{j=i+1}^{k} x_j = y_j) \end{aligned}$$

($\mathrm{SUPP}_{\prec}$ defines the support of $\prec$ (or the singleton $\{z_1\}$, which is needed for the case that the support is empty; our order will be defined in a way that z_1 is always its minimal element). Furthermore, if $\prec$ is an order then $\mathrm{LEX}_{\prec}$ defines the associated k–ary lexicographical order.)

Let

$\xi(x,y,\prec) = \quad x \prec y \vee \mathrm{INI}_{\overset{k}{z}}(x,y)$

We already have $x \prec y$, or $x \prec y$ are in the initial order, or:

$\vee \exists \overset{k}{x}, \overset{k}{y} \Big(\quad \bigwedge_{i=1}^{k} \mathrm{SUPP}_{\prec}(x_i) \wedge \bigvee_{i=1}^{k} \neg \mathrm{SUPP}_{\prec}(y_i) \wedge \varphi_D(\overset{k}{x},\overset{k}{y})$

There exist $\overset{k}{x}, \overset{k}{y}$ such that $\overset{k}{x}$ is already contained in the support of $\prec$ whereas $\overset{k}{y}$ is not, and $\overset{k}{y}$ is a φ_D–successor of $\overset{k}{x}$.

$\wedge \forall \overset{k}{u}, \overset{k}{v} (\mathrm{LEX}_{\prec}(\overset{k}{u},\overset{k}{x}) \wedge \varphi_D(\overset{k}{u},\overset{k}{v}) \rightarrow \bigwedge_{i=1}^{k} \mathrm{SUPP}_{\prec}(v_i))$

Actually, $\overset{k}{x}$ is the lexicographically smallest tuple for which such a $\overset{k}{y}$ exists.

$\wedge \mathrm{SUPP}_{\prec}(x) \wedge \neg \mathrm{SUPP}_{\prec}(y) \wedge \bigvee_{i=1}^{k} \big(y = y_i \wedge \bigwedge_{j=1}^{i-1} \mathrm{SUPP}_{\prec}(y_j)\big)\Big)$

Furthermore, x is already in the support of $\prec$, whereas y is the first y_i which is not. (This means that the new element is larger than any element already in the order.)

Then $\psi(x,y) := [\mathrm{FP}_{x,y,\prec}\xi(x,y,\prec)]xy$ is an $\mathbf{IFP}^2$–formula defining the order $<_{\varphi,\overset{k}{z}}$ described above.

Since the support of our order contains $\overset{k}{z}$ and is closed under φ_D–successors, it contains the deterministic φ–path starting at $\overset{k}{z}$. ■

An immediate consequence of the last two lemmata is:

6.6 Theorem. $\mathbf{DTC} \subseteq \mathbf{IFP}^2$

Using standard techniques to transform **IFP** into **LFP** formulae this also shows $\mathbf{DTC} \subseteq \mathbf{LFP}^5$ (a result which can probably be strengthened). So we might suspect that a similar result holds with **TC** instead of **IFP** or even that the arity hierarchy of **DTC** collapses. However, this is not the case:

6.7 Theorem. *For each $k \geq 1$ there exists a graph–query which is definable in $\mathbf{DTC}^k$, but not in $\mathbf{TC}^{k-1}$.*

The rest of this section is devoted to a proof of the theorem. The strategy of this proof is similar to the strategy of the proof of the main lemma: We are going to define a class of graphs which has a certain property of homogeneity (in Subsection 6.2) and show with the help of an Ehrenfeucht-Fraïssé game (introduced in Subsection 6.3) that the deterministic transitive closure of a first–order formula $\text{D–E}\textsc{dge}_k(\overset{k}{x}, \overset{k}{y})$ is not definable in $\mathbf{TC}^{k-1}$ on this class (this will be shown in Subsection 6.4).
We should already remark at this point that there is an easier proof that the arity hierachy of **DTC** is strict, i.e. that for each $k \geq 1$ we have $\mathbf{DTC}^{k-1} \subsetneq \mathbf{DTC}^k$. The reader who is merely interested in this result should read Remark 6.15 before going through the details of the next subsection.

6.2 The Structures

The graphs we need here are very similar to those we used before. But to make a k-ary DTC-Operater applicable, each k-tuple should have at most one $\text{D–E}\textsc{dge}_k$-successor. The price for this is that we cannot find automorphisms of the structures that move a $(k-1)$–tuple of elements to another component and leave most other rows fixed. But it will suffice here to have automorphisms that move some $(k-1)$–tuples and leave some $(k-1)$–tuples in different rows fixed.
To deal with nested TC–operators we have to iterate the construction.

6.8 Theorem. *Let $k, n \geq 2, t \geq 1$. Then there exists a graph $\mathfrak{H}^t = \mathfrak{H}^t(k,n)$ and a first–order formula* $\text{D–E}\textsc{dge}_k(\overset{k}{x}, \overset{k}{y})$ *such that*

(1) There exists a mapping row : $H^t \longrightarrow \{1, \ldots, n\}$ *such that*

$$\forall \boldsymbol{a}, \boldsymbol{b} \in H^t : (E^{\mathfrak{H}^t} \boldsymbol{ab} \Longrightarrow 0 \leq \text{row}(\boldsymbol{b}) - \text{row}(\boldsymbol{a}) \leq 1) \ .$$

(2) There exists an automorphism $\mathfrak{E}$ *of* $\mathfrak{H}^t$ *which is self–inverse and preserves the rows.*

(3) There exist tuples $\overset{k}{\boldsymbol{c}}, \overset{k}{\boldsymbol{d}} \in H^t$ *in the first and last row respectively such that*

$$\mathfrak{H}^t \models [\mathrm{DTC}_{\overset{k}{x},\overset{k}{y}} \mathrm{D\text{–}E{\scriptstyle DGE}}_k(\overset{k}{x},\overset{k}{y})]\overset{k}{\boldsymbol{c}},\overset{k}{\boldsymbol{d}}$$

$$\textit{and} \quad \mathfrak{H}^t \not\models [\mathrm{DTC}_{\overset{k}{x},\overset{k}{y}} \mathrm{D\text{–}E{\scriptstyle DGE}}_k(\overset{k}{x},\overset{k}{y})]\overset{k}{\boldsymbol{c}},\mathfrak{E}(\overset{k}{\boldsymbol{d}}) \; .$$

(4) There exist sets $\mathcal{F}^s$ $(0 \leq s \leq t)$ *of automorphisms of* $\mathfrak{H}^t$ *(which are in fact subgroups of the automorphism–group) such that* $\mathcal{F}^0 = \{id, \mathfrak{E}\}$ *and for each* $s \leq t$ *the following holds:*
Let $A, B \subset H^t$ *such that* $\forall \boldsymbol{a} \in A, \boldsymbol{b} \in B : |\mathrm{row}(\boldsymbol{a}) - \mathrm{row}(\boldsymbol{b})| > 1$ *and in each row of* $\mathfrak{H}^t$ *there are at most* $(k-1)$ *elements of* $A \cup B$*. Then for all* $f_*, f^* \in \mathcal{F}^{s-1}$ *there exists an* $f \in \mathcal{F}^s$ *which is self–inverse and preserves the rows such that*

$$\forall \boldsymbol{a} \in A : f(\boldsymbol{a}) = f_*(\boldsymbol{a}) \quad \textit{and} \quad \forall \boldsymbol{b} \in B : f(\boldsymbol{b}) = f^*(\boldsymbol{b}) \; .$$

We fix $k, n \geq 2$ for the construction.
Let $P_1, \ldots, P_k$ be unary relation symbols. It is convenient to construct $\{E, P_1, \ldots, P_k\}$–structures with the desired properties first; they can easily be turned into graphs later. The relation P_i is intended to contain the elements of the ith column. We let

$$\mathrm{D\text{–}E{\scriptstyle DGE}}_k(\overset{k}{x},\overset{k}{y}) := \bigwedge_{1 \leq i,j \leq k} Ex_iy_j \wedge \bigwedge_{1 \leq i \leq k} (P_ix_i \wedge P_iy_i) \; .$$

As the reader might guess we start our construction with a structure $\mathfrak{H}^0$ consisting of two disjoint D–Edge$_k$–paths. Formally it is defined as follows:
Let Γ^0 be the group $(\mathbb{Z}_2, +)$ and

$$\begin{aligned} H^0 &:= \{1, \ldots, n\} \times \{1, \ldots, k\} \times \Gamma^0 \\ E^{\mathfrak{H}^0} &:= \big\{ \big((I, a, \gamma), (I+1, b, \gamma)\big) \;\big|\; 1 \leq I < n,\ 1 \leq a, b \leq k,\ \gamma \in \Gamma^0 \big\} \end{aligned}$$

and for each $i \leq k$

$$P_i^{\mathfrak{H}^0} \;:=\; \{(I, i, \gamma) \mid 1 \leq I \leq n, \gamma \in \Gamma^0\}.$$

The automorphism $\mathfrak{E}$ exchanging the components of $\mathfrak{H}^0$ is defined by

$$\mathfrak{E}((I, a, \gamma)) = (I, a, \gamma + 1) \; .$$

We let $\mathcal{F}^0 = \{id, \mathfrak{E}\}$ and, furthermore, $\mathrm{row}((I, a, \gamma)) := I$ and $\mathrm{col}((I, a, \gamma)) := a$.

We are going to define structures $\mathfrak{H}^t$ for $t \geq 1$ inductively. To define $\mathfrak{H}^t$ we do not only need the previously defined structures $\mathfrak{H}^s$ for $s < t$, but also the groups Γ^s that were used to define

$\mathfrak{H}^s$ and groups $\mathcal{F}^s$ of automorphisms of $\mathfrak{H}^s$. Moreover, the mappings row and col need to be extended to $\mathfrak{H}^s$.
So suppose $t \geq 1$ and everything is defined for each $s < t$.
For all $f \in \mathcal{F}^{t-1}$, $R_A, R_B \subseteq \{1, \ldots, n\}$ such that $\forall I \in R_A, J \in R_B : |I - J| > 1$, and $a_I \in \{1, \ldots, k\}$ (for each $I \in R_A \cup R_B$) we define a partial bijection of $\mathfrak{H}^{t-1}$ by

$$\boldsymbol{a} \mapsto \begin{cases} \boldsymbol{a} & \text{if } I := \mathrm{row}(\boldsymbol{a}) \in R_A \text{ and } \mathrm{col}(\boldsymbol{a}) \neq a_I \\ f(\boldsymbol{a}) & \text{if } J := \mathrm{row}(\boldsymbol{a}) \in R_B \text{ and } \mathrm{col}(\boldsymbol{a}) \neq a_J \\ (\text{undefined} & \text{otherwise}) \end{cases} .$$

We let $p_1^t, \ldots, p_{l_t}^t$ be an enumeration of all these partial bijections and P^t the free group generated by $\mathsf{p}_1^t, \ldots, \mathsf{p}_{l_t}^t$.
Note that for each $\mathsf{p} \in \mathsf{P}^t$ the corresponding partial bijection p is a partial isomorphism. Similarly as in the construction of $\mathfrak{G}$ we let $\Gamma^t := (\mathbb{Z}_2^{l_t}, +)$ and define the mapping ρ^t to be the extension of $\mathsf{p}_i^t \longmapsto (0, \ldots, 0, 1, 0, \ldots, 0)$ (with the 1 at the ith position) to a homomorphism from P^t to Γ^t. We let $\mathfrak{H}^t := \mathfrak{H}(\mathfrak{H}^{t-1}, p_1^t, \ldots, p_{l_t}^t, \Gamma^t, \rho^t)$.

To simplify the notation we let $\bar{\Gamma} = \bigoplus_{i=0}^t \Gamma^i$ and write $\bar{\gamma}, \bar{\delta}, \ldots$ or $(\gamma_0, \ldots, \gamma_t)$, $(\delta_0, \ldots, \delta_t), \ldots$ for elements of $\bar{\Gamma}$. Elements of $\mathfrak{H}^t$ are written in the form $(I, a, \bar{\gamma})/_\sim$ or $(I, a, \gamma_0, \ldots, \gamma_t)/_\sim$ instead of

$$((\ldots((I, a, \gamma_0), \gamma_1)/_\sim, \ldots, \gamma_{t-1})/_\sim, \gamma_t)/_\sim.$$

For each $\bar{\gamma} \in \bar{\Gamma}$ we define a mapping $f_{\bar{\gamma}} : \mathfrak{H}^t \longrightarrow \mathfrak{H}^t$ by

$$f_{\bar{\gamma}}((I, a, \bar{\delta})/_\sim) = (I, a, \bar{\delta} + \bar{\gamma})/_\sim ,$$

and we let $\mathcal{F}^t := \{f_{\bar{\gamma}} \mid \bar{\gamma} \in \bar{\Gamma}\}$. We will have to prove (cf. Lemma 6.9) that the $f_{\bar{\gamma}}$ are well–defined. Once we have done this it is obvious that each $f_{\bar{\gamma}}$ is self–inverse, that for $\bar{\gamma}, \bar{\delta} \in \bar{\Gamma}$ we have $f_{\bar{\gamma}} \circ f_{\bar{\delta}} = f_{\bar{\delta}} \circ f_{\bar{\gamma}}$, and that $\mathcal{F}^t$ is closed under concatenation.

However, having seen that the $f_{(\gamma_0, \ldots, \gamma_{t-1})} \in \mathcal{F}^{t-1}$ have these properties it is obvious that Γ^t and ρ^t are suitable for $p_1^t, \ldots, p_{l_t}^t$ (since for all $i, j \leq l_t$ we have $(p_i^t)^{-1} = p_i^t$ and $p_i^t p_j^t = p_j^t p_i^t$ by the analogous properties of the $f_{(\gamma_0, \ldots, \gamma_{t-1})} \in \mathcal{F}^{t-1}$ and the choice of the domains of the p_i^t). Furthermore, the $f_{(\gamma_0, \ldots, \gamma_{t-1})} \in \mathcal{F}^{t-1}$ respect rows and columns hence the mappings row and col defined as usual by $\mathrm{row}((I, a, \bar{\gamma})/_\sim) := I$ and $\mathrm{col}((I, a, \bar{\gamma})/_\sim) := a$ are well–defined.

Before we prove the next lemma we shall have a closer look at the partial bijections p for $\mathsf{p} \in \mathsf{P}^t$. By its definition, on each row each p_i^t $(1 \leq i \leq l_t)$ is the restriction of an $f \in \mathcal{F}^{t-1}$. Thus for each fixed row I, on the intersection of this row with its domain p is the concatenation of mappings in $\mathcal{F}^{t-1}$. Since $\mathcal{F}^{t-1}$ is closed under concatenation it is itself in $\mathcal{F}^{t-1}$. Hence by the definition of $\mathcal{F}^{t-1}$ we have:
For all $I \leq n, \mathsf{p} \in \mathsf{P}^t$ there exist $\zeta_0 \in \Gamma^0, \ldots, \zeta_{t-1} \in \Gamma^{t-1}$ such that

$$\forall \boldsymbol{a} \in \mathrm{dom}(p) : \Big(\mathrm{row}(\boldsymbol{a}) = I \Longrightarrow \boldsymbol{a}^p = f_{(\zeta_0, \ldots, \zeta_{t-1})}(\boldsymbol{a})\Big) . \qquad (*)$$

6.9 Lemma. *For each $\bar{\gamma} \in \bar{\Gamma}$ the mapping $f_{\bar{\gamma}}$ is an automorphism of $\mathfrak{H}^t$.*

Proof. Although not completely trivial, the proof is a rather straightforward induction on t. The main difficulty is to prove that $f_{\bar{\gamma}}$ is well–defined; to see this we proceed as follows: Suppose that $(I, a, \delta_0, \ldots, \delta_t)/_\sim = (I, a, \eta_0, \ldots, \eta_t)/_\sim$. Then for some $\mathsf{p} \in (\rho^t)^{-1}(\eta_t - \delta_t)$ we have $(I, a, \delta_0, \ldots, \delta_{t-1})/_\sim = \big((I, a, \eta_0, \ldots, \eta_{t-1})/_\sim\big)^p$. By $(*)$ we can choose $\zeta_0 \in \Gamma^0, \ldots, \zeta_{t-1} \in \Gamma^{t-1}$ such that

$$\forall \boldsymbol{a} \in \mathrm{dom}(p) : \Big(\mathrm{row}(\boldsymbol{a}) = I \Longrightarrow \boldsymbol{a}^p = f_{(\zeta_0, \ldots, \zeta_{t-1})}(\boldsymbol{a})\Big) .$$

So we have

$$(I, a, \delta_0, \ldots, \delta_{t-1})/_\sim = (I, a, \eta_0 + \zeta_0, \ldots, \eta_{t-1} + \zeta_{t-1})/_\sim$$

which implies, by the induction hypothesis and since $\bar{\Gamma}$ is Abelian,

$$(I, a, \delta_0 + \gamma_0, \ldots, \delta_{t-1} + \gamma_{t-1})/_\sim = \big((I, a, \eta_0 + \gamma_0, \ldots, \eta_{t-1} + \gamma_{t-1})/_\sim\big)^p .$$

(Note that $(I, a, \eta_0 + \gamma_0, \ldots, \eta_{t-1} + \gamma_{t-1})/_\sim \in \mathrm{dom}(p)$ because the domain of p only depends on the first two components, the row and the column, and we already now that $(I, a, \eta_0, \ldots, \eta_{t-1})/_\sim \in \mathrm{dom}(p)$.)

Remembering that $\mathsf{p} \in (\rho^t)^{-1}(\eta_t - \delta_t)$ we now have

$$(I, a, \delta_0 + \gamma_0, \ldots, \delta_{t-1} + \gamma_{t-1}, \delta_t)/_\sim = (I, a, \eta_0 + \gamma_0, \ldots, \eta_{t-1} + \gamma_{t-1}, \eta_t)/_\sim,$$

and an application of Lemma 2.4 yields the desired result that

$$f_{(\gamma_0, \ldots, \gamma_t)}((I, a, \delta_0, \ldots, \delta_t)/_\sim) = f_{(\gamma_0, \ldots, \gamma_t)}((I, a, \eta_0, \ldots, \eta_t)/_\sim).$$

■

For $s < t$ we can identify $\mathcal{F}^s$ with the subset $\{f_{\bar{\gamma}} \mid \bar{\gamma} \in \bigoplus_{i=0}^s \Gamma^i\}$ of $\mathcal{F}^t$. But we must be aware that $f_{(\gamma_0, \ldots, \gamma_s)}$ is an automorphism of $\mathfrak{H}^s$ and that $f_{(\gamma_0, \ldots, \gamma_s, 0, \ldots, 0)}$ extends it to an automorphism of $\mathfrak{H}^t$. In particular $f_{(1,0,\ldots,0)}$ extends $\mathfrak{c}$ to an automorphism of $\mathfrak{H}^t$.

Now we can already prove that the structure $\mathfrak{H}^t$ has property (4) stated in Theorem 6.8:

6.10 Lemma. *Let $A, B \subset H^t$ such that $\forall \boldsymbol{a} \in A, \boldsymbol{b} \in B : |\mathrm{row}(\boldsymbol{a}) - \mathrm{row}(\boldsymbol{b})| > 1$ and in each row of $\mathfrak{H}^t$ there are at most $(k-1)$ elements of $A \cup B$. Then for all $s \le t$, $f_*, f^* \in \mathcal{F}^{s-1}$ there exists an $f \in \mathcal{F}^s$ which is self–inverse and preserves the rows such that*

$$\forall \boldsymbol{a} \in A : f(\boldsymbol{a}) = f_*(\boldsymbol{a}) \quad \textit{and} \quad \forall \boldsymbol{b} \in B : f(\boldsymbol{b}) = f^*(\boldsymbol{b}) .$$

Proof. Say, $f_*^{-1} \circ f^*$ is the automorphism $f_{(\gamma_0, \ldots, \gamma_{s-1}, 0, \ldots, 0)}$ of $\mathfrak{H}^t$.
We can assume without loss of generality that A and B are of the following form (otherwise we extend them):

There exist sets $R_A, R_B \subseteq \{1, \ldots, n\}$ such that $\forall I \in R_A, J \in R_B : |I - J| > 1$, and for each $I \in R_A \cup R_B$ there exists an $a_I \in \{1, \ldots, k\}$ such that

$$A = \{\boldsymbol{a} \in H^t \mid \mathrm{row}(\boldsymbol{a}) \in R_A, \mathrm{col}(\boldsymbol{a}) \neq a_{\mathrm{row}(\boldsymbol{a})}\}$$

and

$$B = \{\boldsymbol{b} \in H^t \mid \mathrm{row}(\boldsymbol{b}) \in R_B, \mathrm{col}(\boldsymbol{b}) \neq a_{\mathrm{row}(\boldsymbol{b})}\}.$$

For each $r \leq t$ we let

$$A_r = \{\boldsymbol{a} \in H^r \mid \mathrm{row}(\boldsymbol{a}) \in R_A, \mathrm{col}(\boldsymbol{a}) \neq a_{\mathrm{row}(\boldsymbol{a})}\}$$

and

$$B_r = \{\boldsymbol{b} \in H^r \mid \mathrm{row}(\boldsymbol{b}) \in R_B, \mathrm{col}(\boldsymbol{b}) \neq a_{\mathrm{row}(\boldsymbol{b})}\} .$$

Thus, in particular, $A = A_t$ and $B = B_t$.
There exists an $i \leq l_s$ such that the partial bijection p_i^s of $\mathfrak{H}^{s-1}$ is the identity on A_{s-1} and maps each element of B_{s-1} to its image under $f_{(\gamma_0, \ldots, \gamma_{s-1})}$.
Let $\delta_s := \rho^s(\mathsf{p}_i^s)$. By Lemma 2.4 $f_{(0,\ldots,0,\delta_s)}$ is an automorphism of $\mathfrak{H}^s$ which, by Lemma 2.8, is the identity on A_s and maps each element of B_s to its image under $f_{(\gamma_0,\ldots,\gamma_{s-1},0)}$.
Then for each $\boldsymbol{a} \in A_t$ we have $f_{(0,\ldots,0,\delta_s,0,\ldots,0)}(\boldsymbol{a}) = \boldsymbol{a}$ and for each $\boldsymbol{b} \in B_t$

$$f_{(0,\ldots,0,\delta_s,0,\ldots,0)}(\boldsymbol{b}) = f_{(\gamma_0,\ldots,\gamma_{s-1},0,\ldots,0)}(\boldsymbol{b}) = f_*^{-1} \circ f^*(\boldsymbol{b}) .$$

We complete the proof by letting $f := f_* \circ f_{(0,\ldots,0,\delta_s,0,\ldots,0)}$. ■

In the following we proceed quite similar to Subsection 2.2. Instead of proving that our structure has at least 2 EDGE_k–connected components we shall now prove that our paths are deterministic, i.e. that each tuple has at most one D–EDGE_k–successor (at least each tuple of the form $(I, 1, \bar{\gamma})/_\sim \ldots (I, k, \bar{\gamma})/_\sim$).
For each $\gamma_t \in \Gamma^t$ we define $\mathrm{supp}(\gamma_t)$ to be set set of places in the 0–1–vector γ_t that are 1. Completely analogous to Lemma 2.11 we obtain:

6.11 Lemma. *Let* $\bar{\gamma} = (\gamma_1, \ldots, \gamma_t), \bar{\delta} = (\delta_1, \ldots \delta_t) \in \bar{\Gamma}, I \leq n$ *and suppose that* $\mathrm{supp}(\delta_t - \gamma_t) = \{i_1, \ldots, i_r\}$. *Then*

$$(I, a, \bar{\gamma}) \sim (I, b, \bar{\delta}) \quad \Longleftrightarrow \quad (I, a, \bar{\gamma}) \sim (I, b, \bar{\delta}) \quad \textit{via} \quad \mathsf{p}_{i_1}^t \ldots \mathsf{p}_{i_r}^t .$$

The next lemma is a strong version of Lemma 2.13:

6.12 Lemma. *Let $I \leq n, \bar{\gamma} \neq \bar{\delta} \in \bar{\Gamma}$. Then there exists an $a \in \{1, \ldots, k\}$ such that*

$$(I, a, \bar{\gamma}) \not\sim (I, a, \bar{\delta}) \ .$$

Proof. The proof is by induction on t, and again the case $t = 0$ is trivial.
Suppose the lemma is true for $(t-1)$ and let $I \leq n$ and $(\gamma_0, \ldots, \gamma_t) \neq (\delta_0, \ldots, \delta_t) \in \bar{\Gamma}$. We can assume that there exists a $b \in \{1, \ldots, k\}$ such that

$$(I, b, \gamma_0, \ldots, \gamma_t) \sim (I, b, \delta_0, \ldots, \delta_t) \ .$$

(If no such b exists there is nothing to prove.)
Then

$$(I, b, \gamma_0, \ldots, \gamma_t) \sim (I, b, \delta_0, \ldots, \delta_t) \quad via \quad \mathsf{p}^t_{i_1} \ldots \mathsf{p}^t_{i_r}$$

where $\{i_1, \ldots, i_r\} = \mathrm{supp}(\delta_t - \gamma_t\}$.
If $r > 0$ we can choose $a \in \{1, \ldots, k\}$ such that $(I, a, \delta_0, \ldots, \delta_{t-1}) \notin \mathrm{dom}(p^t_{i_1})$.

If $r = 0$, i.e. $\gamma_t = \delta_t$, then $(\gamma_0, \ldots, \gamma_{t-1}) \neq (\delta_0, \ldots, \delta_{t-1})$. By the induction hypothesis there exists an $a \in \{1, \ldots, k\}$ such that

$$(I, a, \gamma_0, \ldots, \gamma_{t-1}) \not\sim (I, a, \delta_0, \ldots, \delta_{t-1}) \ .$$

Since π_{γ_t} is injective, by Lemma 2.7 this implies

$$(I, a, \gamma_0, \ldots, \gamma_t) \not\sim (I, a, \delta_0, \ldots, \delta_t) \ .$$

■

To prove that our paths are really deterministic, which is in a sense the analogon of Lemma 2.16 for this case, in particular we need the following preliminary Lemma 6.13 corresponding to Lemma 2.14.

6.13 Lemma. *Let $a, b \in H^{t-1}$ and $\gamma_t, \delta_t \in \Gamma^t$ such that $E^{\mathfrak{H}^t}(a, \gamma_t)/_\sim(b, \delta_t)/_\sim$. Then there exist $\hat{a}, \hat{b} \in H^{t-1}$ and $\eta_t \in \Gamma_t$ such that*

(i) $(a, \gamma_t) \sim (\hat{a}, \eta_t)$

(ii) $(b, \delta_t) \sim (\hat{b}, \eta_t)$

(iii) $E^{\mathfrak{H}^{t-1}} \hat{a}\hat{b}$

(iv) $\mathrm{supp}(\eta_t - \gamma_t) \,\dot{\cup}\, \mathrm{supp}(\eta_t - \delta_t) = \mathrm{supp}(\delta_t - \gamma_t)$

The **proof** can be given analogously to the proof of Lemma 2.14. ■

6.14 Lemma. *Let* $\overset{k}{\boldsymbol{a}}, \overset{k}{\boldsymbol{b}} \in H^t$ *such that* $\mathfrak{H}^t \models \text{D–}\textsc{Edge}_k[\overset{k}{\boldsymbol{a}}, \overset{k}{\boldsymbol{b}}]$ *and*

$$\exists \bar{\gamma} \in \bar{\Gamma}, I \leq n \; \forall i \leq k : \boldsymbol{a}_i = (I, i, \bar{\gamma})/_\sim \, .$$

Then $\forall i \leq k : \boldsymbol{b}_i = (I+1, i, \bar{\gamma})/_\sim$.

Recalling the proof of Lemma 2.16 we might expect that the proof of this lemma is not too easy. On top of everything we have to incorporate the additional parameter t here.
However, it becomes a bit simpler because the direction of the paths is determined by the direction of E and the columns are fixed by the P_i. So an analogue of Lemma 2.15 is not necessary here. Moreover, the p_i^t never keep whole rows fixed thus we can leave out the first step of the proof of Lemma 2.16. For this reason we also do not need edges between x_i and x_j $(i, j \leq k)$ or between y_i and y_j $(i, j \leq k)$ in the formula $\text{D–}\textsc{Edge}_k(\overset{k}{x}, \overset{k}{y})$.

The **proof** is by induction on t:
It can be immediately seen from the definition of $\mathfrak{H}^0$ that the statement of the lemma holds for $t = 0$. So suppose that $t \geq 1$ and $\bar{\gamma} = (\gamma_0, \ldots, \gamma_t)$.
$\mathfrak{H}^t \models \text{D–}\textsc{Edge}_k[\overset{k}{\boldsymbol{a}}, \overset{k}{\boldsymbol{b}}]$ means that $\forall i, j \leq k : E^{\mathfrak{H}^t} \boldsymbol{a}_i \boldsymbol{b}_j$ and $\forall i \leq k : P_i^{\mathfrak{H}^t} \boldsymbol{a}_i \wedge P_i^{\mathfrak{H}^t} \boldsymbol{b}_i$. Note that for all $i \leq k$ the former implies $\text{row}(\boldsymbol{b}_i) = I + 1$ and the latter implies $\text{col}(\boldsymbol{a}_i) = \text{col}(\boldsymbol{b}_i) = i$.
We shall prove:

Claim: $\forall h \leq k \; \forall j < t \; \exists \delta_{hj} \in \Gamma^j : \; \boldsymbol{b}_h = (I+1, h, \delta_{h0}, \ldots, \delta_{h(t-1)}, \gamma_t)/_\sim$
where

$$\begin{aligned}\mathfrak{H}^{t-1} \models \text{D–}\textsc{Edge}_k\big[&(I, 1, \gamma_0, \ldots, \gamma_{t-1})/_\sim, \ldots, (I, k, \gamma_0, \ldots, \gamma_{t-1})/_\sim, \\ &(I+1, 1, \delta_{10}, \ldots, \delta_{1(t-1)})/_\sim, \ldots, (I+1, k, \delta_{k0}, \ldots, \delta_{k(t-1)})/_\sim\big]\end{aligned}$$

Once we have proved this claim we can apply the induction hypothesis and are finished. More precisely, the induction hypothesis implies that for each $h \leq k$

$$(I+1, h, \delta_{h0}, \ldots, \delta_{h(t-1)})/_\sim = (I+1, h, \gamma_0, \ldots, \gamma_{t-1})/_\sim.$$

Unraveling our notational convention, this shows

$$\begin{aligned}\boldsymbol{b}_h &= (I+1, h, \delta_{h0}, \ldots, \delta_{h(t-1)}, \gamma_t)/_\sim \\ &= \big((I+1, h, \delta_{h0}, \ldots, \delta_{h(t-1)})/_\sim, \gamma_t\big)/_\sim \\ &= (I+1, h, \gamma_0, \ldots, \gamma_{t-1}, \gamma_t)/_\sim.\end{aligned}$$

Proof (of the claim). We fix a $h \leq k$ and suppose that

$$\boldsymbol{b}_h = (I+1, h, \delta_0, \ldots, \delta_t)/_\sim.$$

By the definition of the partial bijections $p_1^t, \ldots, p_{l_t}^t$ for each $x \in \mathit{supp}(\delta_t - \gamma_t)$ there exists an $i_x \in \{1, \ldots, k\}$ such that $a \notin \text{dom}(p_x^t)$ for each $a \in H^{t-1}$ with $\text{row}(a) = I$ and $\text{col}(a) = i_x$. Let us fix such an x and i_x.

By $E^{\mathfrak{H}^t} \boldsymbol{a}_{i_x} \boldsymbol{b}_h$ and the previous Lemma 6.13 there exist $\eta_t \in \Gamma^t$ and $\gamma'_i, \delta'_i \in \Gamma^i$ (for all $i \leq t-1$) such that

$$\left.\begin{array}{ll} & (I, i_x, \gamma_0, \ldots, \gamma_t) \sim (I, i_x, \gamma'_0, \ldots, \gamma'_{t-1}, \eta_t) \\ \wedge & (I+1, h, \delta_0, \ldots, \delta_t) \sim (I+1, h, \delta'_0, \ldots, \delta'_{t-1}, \eta_t) \\ \wedge & E^{\mathfrak{H}^{t-1}}(I, i_x, \gamma'_0, \ldots, \gamma'_{t-1})/_\sim(I+1, h, \delta'_0, \ldots, \delta'_{t-1})/_\sim \\ \wedge & \operatorname{supp}(\eta_t - \gamma_t) \,\dot\cup\, \operatorname{supp}(\eta_t - \delta_t) = \operatorname{supp}(\delta_t - \gamma_t) \qquad . \end{array}\right\} (*)$$

Since for all $a \in H^{t-1}$ with $\operatorname{row}(a) = I$ and $\operatorname{col}(a) = i_x$ we have $a \notin \operatorname{dom}(p^t_x)$, this implies

$$\begin{array}{ll} & x \notin \operatorname{supp}(\eta_t - \gamma_t) \\ \Longrightarrow & x \in \operatorname{supp}(\eta_t - \delta_t) \qquad (\text{since } x \in \operatorname{supp}(\delta_t - \gamma_t)) \\ \Longrightarrow & \forall b \in H^{t-1} \text{ with } \operatorname{row}(b) = I+1 \text{ and } \operatorname{col}(b) = h: \ b \in \operatorname{dom}(p^t_x). \qquad (**) \end{array}$$

As x was arbitrary we have proved this for all $x \in \operatorname{supp}(\delta_t - \gamma_t)$.

Let $\operatorname{supp}(\delta_t - \gamma_t) = \{x_1, \ldots, x_r\}$ be ordered such that $\operatorname{supp}(\delta_t - \eta_t) = \{x_1, \ldots, x_s\}$ and $\operatorname{supp}(\eta_t - \gamma_t) = \{x_{s+1}, \ldots, x_r\}$ for an $s \leq r$ (where η_t is taken from $(*)$).
Then by $(*)$ we have

$$\begin{array}{rcl} \big((I+1, h, \delta_0, \ldots, \delta_{t-1})/_\sim\big)^{p^t_{x_1} \cdots p^t_{x_s}} &=& (I+1, h, \delta'_0, \ldots, \delta'_{t-1})/_\sim \ , \\ \big((I, i_x, \gamma'_0, \ldots, \gamma'_{t-1})/_\sim\big)^{p^t_{x_{s+1}} \cdots p^t_{x_r}} &=& (I, i_x, \gamma_0, \ldots, \gamma_{t-1})/_\sim \ . \end{array}$$

(Since $\forall i, j \leq l_t : p^t_i = (p^t_i)^{-1} \wedge p^t_i p^t_j = p^t_j p^t_i$.) Moreover, by $(**)$ we can choose $\delta_{h0}, \ldots, \delta_{h(t-1)}$ such that

$$\big((I+1, h, \delta'_0, \ldots, \delta'_{t-1})/_\sim\big)^{p^t_{x_{s+1}} \cdots p^t_{x_r}} = (I+1, h, \delta_{h0}, \ldots, \delta_{h(t-1)})/_\sim \ .$$

Hence $\boldsymbol{b}_h = (I+1, h, \delta_{h0}, \ldots, \delta_{h(t-1)}, \gamma_t)/_\sim$.
Since $p^t_{x_{s+1}} \cdots p^t_{x_r}$ is a partial isomorphism

$$E^{\mathfrak{H}^{t-1}}(I, i_x, \gamma'_0, \ldots, \gamma'_{t-1})/_\sim(I+1, h, \delta'_0, \ldots, \delta'_{t-1})/_\sim$$

implies

$$E^{\mathfrak{H}^{t-1}}(I, i_x, \gamma_0, \ldots, \gamma_{t-1})/_\sim(I+1, h, \delta_{h0}, \ldots, \delta_{h(t-1)})/_\sim \ .$$

So the claim is proved and we are finished. ∎

Proof (of Theorem 6.8): The previous lemmata show that $\mathfrak{H}^t$ has all the desired properties:
(1) is obvious. (2) follows from Lemma 6.9 with $\boldsymbol{\mathfrak{c}} = f_{(1,0,\ldots,0)}$.
To prove (3) let $\boldsymbol{c}_i = (1, i, 0, \ldots, 0)/_\sim$ and $\boldsymbol{d}_i = (n, i, 0, \ldots, 0)/_\sim$ (for each $i \leq k$). Then there exists a D–EDGE$_k$–path from $\overset{k}{\boldsymbol{c}}$ to $\overset{k}{\boldsymbol{d}}$, and by Lemma 6.14 it is deterministic. Hence there exists no path from $\overset{k}{\boldsymbol{c}}$ to $\boldsymbol{\mathfrak{c}}(\overset{k}{\boldsymbol{d}})$ because by Lemma 6.12 we have $\overset{k}{\boldsymbol{d}} \neq \boldsymbol{\mathfrak{c}}(\overset{k}{\boldsymbol{d}})$.
We have already proved (4) in Lemma 6.10.

Unfortunately, $\mathfrak{H}^t$ is a $\{E, P_1, \ldots, P_k\}$–structure and not a graph. But recalling the techniques used in in 4.1 and 5.1 it should be clear that we can extend $\mathfrak{H}^t$ in a way that the direction of $E^{\mathfrak{H}^t}$ becomes first–order definable thus we can close $E^{\mathfrak{H}^t}$ under symmetry. Then we extend the resulting structure again in order to make the columns definable and remove $P_1, \ldots, P_k$. We obtain a graph which still has properties (1)–(4). ■

6.15 Remark. If we only wanted to prove that $\mathbf{DTC}^k \not\subseteq \mathbf{DTC}^{k-1}$ we could have already worked with the structure $\mathfrak{H}^1$. The reason for this is that $(k-1)$–tuples cannot define any $(k-1)$–tuples in other rows of $\mathfrak{H}^1$. So there are no deterministic paths of $(k-1)$–tuples crossing longer distances and $\mathbf{DTC}^{k-1}$–formulae are completely ineffective.
But note that proving $\mathfrak{H}^1$ has the properties stated in Theorem 6.8 is not much easier than proving the whole theorem. Basically, it only simplifies the notation. What makes the alternative proof really easier is that it avoids the complicated arguments needed in Subsection 6.4.
□

6.3 An Ehrenfeucht-Fraïssé Game for Transitive Closure Logic

We want to prove that the deterministic transitive closure of D–EDGE$_k$ is not TC^{k-1}–expressible on the class $\{\mathfrak{H}^t(k,n) \mid n \geq 2, t \geq 1\}$ of structures defined in the last paragraph. Again the proof uses an Ehrenfeucht-Fraïssé game, which was introduced by Grädel [Grä91].

6.16 Definition. The k–ary t–nested r–move TC–game is played by two players, the spoiler and the duplicator, on a pair $\mathfrak{A}, \mathfrak{B}$ of structures of the same signature with r pairs of pebbles $P_1, Q_1, \ldots, P_r, Q_r$ following the moves defined below.

We say that two elements $a \in A$ and $b \in B$ correspond in a situation of the game if they are the interpretations of the same constant in $\mathfrak{A}$ and $\mathfrak{B}$ respectively or if a pair P_i, Q_i of pebbles is placed on them.

In each situation of the game the spoiler selects one of the following moves, respecting that throughout the whole game at most t TC or $\neg$TC–moves are allowed.

∃-move: The spoiler places a yet unused pebble P_i on (an element of structure) $\mathfrak{A}$, the duplicator answers by placing the corresponding pebble Q_i on $\mathfrak{B}$.

∀-move: The spoiler places Q_i on $\mathfrak{B}$, the duplicator answers by placing P_i on $\mathfrak{A}$.

TC-move: The spoiler selects an $l \leq k$ such that there are still $2l$ yet unused pairs of pebbles. Then he selects a sequence $\overset{l}{a}_1, \ldots, \overset{l}{a}_m$ of l-tuples in $\mathfrak{A}$ where $\overset{l}{a}_1$ and $\overset{l}{a}_m$ consist of elements already pebbled and constants.

The duplicator selects a sequence $\overset{l}{b}_1, \ldots, \overset{l}{b}_n$ of l-tuples in $\mathfrak{B}$ where b_{1i} and b_{ni} $(1 \leq i \leq l)$ correspond to a_{1i} and a_{mi} respectively.

Now the spoiler selects an $i < n$ and places $2l$ yet unused pebbles on $b_{i1}, \ldots, b_{il}, b_{(i+1)1}, \ldots, b_{(i+1)l}$ respectively.

The duplicator selects a $j < m$ and places the corresponding pebbles on $a_{j1}, \ldots, a_{jl}, a_{(j+1)1}, \ldots, a_{(j+1)l}$ respectively.

¬TC-move: Analogous to the TC–move with reversed boards.

The duplicator wins if after all pebbles are placed on the structures the mapping between the corresponding elements is a partial isomorphism between $\mathfrak{A}$ and $\mathfrak{B}$. $\square$

To relate the game to transitive closure logic we need the following notions of rank for **TC**–formulae: We let $qr([TC_{\overset{l}{x},\overset{l}{y}}\, \varphi(\overset{l}{x}, \overset{l}{y})]\overset{l}{u}\overset{l}{v}) = qr(\varphi) + 2l$ and define the *quantifier-rank* inductively in the natural way. The *TC–rank* $tr(\varphi)$ of a **TC**–formula φ is the maximal number of nested TC–operators in φ.

6.17 Theorem [Grä91]. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be structures of the same signature. Then for all $k, t, r \geq 0$ the following two statements are equivalent:*

(1) For each $\mathbf{TC}^k$–sentence φ with $qr(\varphi) \leq r$ and $tr(\varphi) \leq t$ we have

$$\mathfrak{A} \models \varphi \quad \Longleftrightarrow \quad \mathfrak{B} \models \varphi.$$

(2) The duplicator has a winning strategy for the k–ary t–nested r–move TC–game on $\mathfrak{A}, \mathfrak{B}$.

6.4 The Game

To calculate the number of rows of our structures we define a function

$$D_k : \mathbb{N}^2 \longrightarrow \mathbb{N} \quad \text{by} \quad D_k(r,t) := (3^k k!)^t 2^{r+1}.$$

We fix $r, t \geq 0$ which are intended to be the quantifier rank and the TC–rank of a formula and let $n := D_k(r,t) + 1$ and $\mathfrak{H} := \mathfrak{H}^{t(k-1)}(k, n)$ (chosen according to Theorem 6.8). We select $\overset{k}{\boldsymbol{c}}, \overset{k}{\boldsymbol{d}} \in H$ satisfying property (3) of the structure and let $\mathfrak{H}_1 := (\mathfrak{H}, \overset{k}{\boldsymbol{c}}, \overset{k}{\boldsymbol{d}})$ and $\mathfrak{H}_2 := (\mathfrak{H}, \overset{k}{\boldsymbol{c}}, \boldsymbol{\alpha}(\overset{k}{\boldsymbol{d}}))$; hence $\mathfrak{H}_1$ and $\mathfrak{H}_2$ can be distinguished by a $\mathbf{DTC}^k$–sentence that does not depend on r or t. Thus by Theorem 6.17 we are finished once we have proved:

6.18 Proposition. *The duplicator has a winning strategy for the $(k-1)$–ary t–nested r–move TC–game on $\mathfrak{H}_1$ and $\mathfrak{H}_2$.*

Proof. For all $\boldsymbol{a}, \boldsymbol{b} \in \mathfrak{H}$ we define the *distance* between $\boldsymbol{a}$ and $\boldsymbol{b}$ by $d(\boldsymbol{a}, \boldsymbol{b}) := |\mathrm{row}(\boldsymbol{a}) - \mathrm{row}(\boldsymbol{b})|$. Since corresponding elements will always be in the same row we can extend this definition to two pairs of corresponding elements.

By a simultaneous induction on the number of pebbles that are already placed on the board and the number of TC–moves that have already been made we prove that the duplicator can preserve the following throughout the whole game:

Suppose that $r' \leq r$ pairs of pebbles have already been placed on the board and $t' \leq t$ is the number of TC-moves that have been made yet.
Then

(i) If $(\boldsymbol{a}, \boldsymbol{b})$ is a pair of corresponding elements we have $\mathrm{row}(\boldsymbol{a}) = \mathrm{row}(\boldsymbol{b})$.

(ii) For each pair $\boldsymbol{p} = (\boldsymbol{a}, \boldsymbol{b})$ of corresponding elements there exists an automorphism $f_{\boldsymbol{p}} \in \mathcal{F}^{t'(k-1)}$ of $\mathfrak{H}$ such that $f_{\boldsymbol{p}}(\boldsymbol{a}) = \boldsymbol{b}$.

Furthermore, the automorphisms can be chosen such that if $\boldsymbol{q}$ is another pair of corresponding elements and $d(\boldsymbol{p}, \boldsymbol{q}) < D_k(r - r', t - t')$ then $f_{\boldsymbol{p}} = f_{\boldsymbol{q}}$.

Note that we are finished once we have proved this.

In the beginning we let $f_{(\boldsymbol{c}_i, \boldsymbol{c}_i)} = id$ and $f_{(\boldsymbol{d}_i, \boldsymbol{\mathfrak{c}}(\boldsymbol{d}_i))} = \boldsymbol{\mathfrak{c}}$ (for each $i \leq k$). Then (i) and (ii) hold because $id, \boldsymbol{\mathfrak{c}} \in \mathcal{F}^0$ and $\forall i, j \leq k :\ d(\boldsymbol{c}_i, \boldsymbol{d}_j) = n - 1 = D_k(r, t)$.

$\exists$-move: The spoiler places a pebble on $\boldsymbol{a} \in \mathfrak{H}_1$. Let $\boldsymbol{p}$ be a nearest pair of corresponding elements. The duplicator places the corresponding pebble on $\boldsymbol{b} := f_{\boldsymbol{p}}(\boldsymbol{a})$, and we define $f_{(\boldsymbol{a},\boldsymbol{b})} := f_{\boldsymbol{p}}$.

Then for all pairs $\boldsymbol{q}$ of corresponding elements with $f_{\boldsymbol{q}} \neq f_{\boldsymbol{p}}$ (hence $d(\boldsymbol{q}, \boldsymbol{p}) \geq D_k(r - r', t - t')$ by induction hypothesis) we have

$$d((\boldsymbol{a}, \boldsymbol{b}), \boldsymbol{q}) \geq \left\lfloor \frac{D_k(r - r', t - t')}{2} \right\rfloor = D_k(r - r' - 1, t - t') \ .$$

$\forall$-move: is treated analogously.

The essential part of the duplicator's answer to a TC–move is hidden in Lemma 6.19 below. However, to understand the lemma it is easier to see its application first. Hence we finish our game (and thus the proof of Proposition 6.18) before:

TC-move: The spoiler selects a sequence $\overset{l}{\boldsymbol{a}}_1, \ldots, \overset{l}{\boldsymbol{a}}_m \in \mathfrak{H}_1$ where $\overset{l}{\boldsymbol{a}}_1$ and $\overset{l}{\boldsymbol{a}}_m$ consist of constants or elements that have already been pebbled in previous moves. Furthermore, $l \leq k - 1$ such that $r' + 2l \leq r$.

We show that the duplicator can define a sequence $\overset{l}{\boldsymbol{b}}_1, \ldots, \overset{l}{\boldsymbol{b}}_m \in \mathfrak{H}_2$ such that if the spoiler places pebbles on $\overset{l}{\boldsymbol{b}}_i, \overset{l}{\boldsymbol{b}}_{i+1}$ afterwards she can answer by placing the corresponding pebbles on $\overset{l}{\boldsymbol{a}}_i, \overset{l}{\boldsymbol{a}}_{i+1}$. To preserve (ii) we are also going to define automorphisms $f_{(\boldsymbol{a}_{ij}, \boldsymbol{b}_{ij})}$ and $f_{(\boldsymbol{a}_{i+1j}, \boldsymbol{b}_{i+1j})}$ in $\mathcal{F}^{(t'+1)(k-1)}$.
Formally, we will define suitable automorphisms first; they will determine the $\boldsymbol{b}_{ij}$. That is, we are going to define automorphisms $f_{ij} \in \mathcal{F}^{t'(k-1)+l} \subseteq \mathcal{F}^{(t'+1)(k-1)}$ for $i \leq m, j \leq l$ such that $f_{(\boldsymbol{a}_{ij}, \boldsymbol{b}_{ij})}$ will be f_{ij} and thus $\boldsymbol{b}_{ij} = f_{ij}(\boldsymbol{a}_{ij})$.

Recall, however, that the automorphisms $f_{\boldsymbol{p}}$ needed in (ii) depend on the pair $\boldsymbol{p}$ of elements that will eventually be pebbled (after the duplicator has defined her path), but we want to use them to define the path $\overset{l}{\boldsymbol{b}}_1, \ldots, \overset{l}{\boldsymbol{b}}_m$, i.e. in a situation where we do not know which elements will be pebbled later. This may cause problems, consider the following scenario: Suppose we have defined our f_{ij} and $\boldsymbol{b}_{ij} = f_{ij}(\boldsymbol{a}_{ij})$ for all $i \leq m, j \leq l$ and the spoiler places pebbles on, say, $\overset{l}{\boldsymbol{b}}_i, \overset{l}{\boldsymbol{b}}_{i+1}$. Then the duplicator answers by placing her pebbles on $\overset{l}{\boldsymbol{a}}_i, \overset{l}{\boldsymbol{a}}_{i+1}$ and we let $f_{(\boldsymbol{a}_{ij}, \boldsymbol{b}_{ij})} := f_{ij}$ and $f_{(\boldsymbol{a}_{i+1j}, \boldsymbol{b}_{i+1j})} := f_{i+1j}$.

Now consider another choice of the spoiler, suppose he places his pebbles on $\overset{l}{\boldsymbol{b}}_{i-1}, \overset{l}{\boldsymbol{b}}_i$ (where $\boldsymbol{b}_{i-1j} = f_{i-1j}(\boldsymbol{a}_{i-1j})$). The duplicator answers by placing her pebbles on $\overset{l}{\boldsymbol{a}}_{i-1}, \overset{l}{\boldsymbol{a}}_i$. But there is no reason that we should also let $f_{(\boldsymbol{a}_{ij}, \boldsymbol{b}_{ij})} := f_{ij}$ in this case. (We only need $f_{(\boldsymbol{a}_{ij}, \boldsymbol{b}_{ij})}(\boldsymbol{a}_{ij}) = \boldsymbol{b}_{ij} = f_{ij}(\boldsymbol{a}_{ij})$.) In fact it is necessary to define it in a different way here. Therefore we also have to define automorphisms $f^+_{ij} \in \mathcal{F}^{t'(k-1)+l}$ $(2 \leq i \leq m, j \leq l)$ such that we can define $f_{(\boldsymbol{a}_{ij}, \boldsymbol{b}_{ij})} = f^+_{ij}$ in the second case. Clearly this requires $\forall j \leq l :$ $f_{ij}(\boldsymbol{a}_{ij}) = f^+_{ij}(\boldsymbol{a}_{ij})$.

These considerations lead us to the following:

Claim: There exist automorphisms $f_{ij} \in \mathcal{F}^{t'(k-1)+l}$ $(1 \leq i \leq m-1, j \leq l)$ and $f^+_{ij} \in \mathcal{F}^{t'(k-1)+l}$ $(2 \leq i \leq m, j \leq l)$ of $\mathfrak{H}$ such that

(1) $f_{ij}(\boldsymbol{a}_{ij}) = f^+_{ij}(\boldsymbol{a}_{ij})$ (for all $2 \leq i \leq m-1, j \leq l$)

(2) If $d(\boldsymbol{a}_{ij}, \boldsymbol{a}_{ij'}) < D_k(r-r', t-t'-1)$ then $f_{ij} = f_{ij'}$ and $f^+_{ij} = f^+_{ij'}$.
If $d(\boldsymbol{a}_{ij}, \boldsymbol{a}_{i+1j'}) < D_k(r-r', t-t'-1)$ then $f_{ij} = f^+_{i+1j'}$.
(for all possible i, j, j').[13]
(This item guarantees that (ii) will be preserved among the new pebbled elements.)

[13]Hopefully it is clear which i, j, j' are "possible" here. For example, the statement "If $d(\boldsymbol{a}_{ij}, \boldsymbol{a}_{ij'}) < D_k(r-r', t-t'-1)$ then $f_{ij} = f_{ij'}$" is claimed for all $i \leq m-1, j, j' \leq l$ and the analogous statement for f^+_{ij} is claimed for all $2 \leq i \leq m, j, j' \leq l$.

(3) If $d(\boldsymbol{a}_{ij}, \boldsymbol{a}) < D_k(r-r', t-t'-1)$ for an $\boldsymbol{a}$ that belongs to a pair $(\boldsymbol{a}, \boldsymbol{b})$ of (already) corresponding elements then $f_{ij} = f_{ij}^+ = f_{(\boldsymbol{a},\boldsymbol{b})}$
(for all possible i, j).
(This item guarantees that the relation between new and old corresponding elements preserves (ii).)

We have seen how to give the duplicator's answer to the TC–move once we have proved this:
She defines her path by $\boldsymbol{b}_{ij} := f_{ij}(\boldsymbol{a}_{ij})$ for each $i < m$, $j \leq l$ and $\boldsymbol{b}_{mj} := f_{mj}^+(\boldsymbol{a}_{mj})$ for each $j \leq l$. If the spoiler places pebbles on $\overset{l}{\boldsymbol{b}}_i, \overset{l}{\boldsymbol{b}}_{i+1}$ she answers by placing the corresponding pebbles on $\overset{l}{\boldsymbol{a}}_i, \overset{l}{\boldsymbol{a}}_{i+1}$. We let, for each $j \leq l$, $f_{(\boldsymbol{a}_{ij}, \boldsymbol{b}_{ij})} := f_{ij}$ and $f_{(\boldsymbol{a}_{i+1j}, \boldsymbol{b}_{i+1j})} := f_{i+1j}^+$.

Proof (of the claim): By the induction hypotheses (i) and (ii) we can find integers $h \geq 1$ and

$$1 = N_0 \leq M_1 < N_1 \leq M_2 < \ldots < N_{h-1} \leq M_h < N_h \leq M_{h+1} = n$$

such that

- $N_i - M_i \geq D_k(r-r', t-t')$ $(1 \leq i \leq h)$.
- For each pair p of corresponding elements there exists an $i \in \{0, \ldots, h\}$ such that $\mathrm{row}(\boldsymbol{p}) \in [N_i, M_{i+1}]$.
- On the other hand for each $i \in \{0, \ldots, h\}$ there exists a $\boldsymbol{p}$ such that $\mathrm{row}(\boldsymbol{p}) \in [N_i, M_{i+1}]$.
- If $\boldsymbol{p}$ and $\boldsymbol{p}'$ are pairs of corresponding elements such that $\boldsymbol{p}, \boldsymbol{p}' \in [N_i, M_{i+1}]$ then $f_{\boldsymbol{p}} = f_{\boldsymbol{p}'} =: f_{[N_i, M_{i+1}]}$.

We divide $\{1, \ldots, n\}$ into the non-empty intervals

$$[N_0, M_1],]M_1, N_1[, \ldots,]M_h, N_h[, [N_h, M_{h+1}].$$

These intervals are treated independently in the following way:

For each $h' \leq h$ we consider only those $\boldsymbol{a}_{ij}$ with $N_{h'-1} \leq \mathrm{row}(\boldsymbol{a}_{ij}) \leq M_{h'+1}$. Hence we obtain a sequence of (at most) l–tuples in each interval $[N_{h'-1}, M_{h'+1}]$. Without loss of generality we can assume that it is in fact a sequence of l–tuples.

We apply Lemma 6.19 to this sequence with $s = t'(k-1)$, $M = M_{h'}$, $N = N_{h'}$, $f_* = f_{[N_{h'-1}, M_{h'}]}$, $f^* = f_{[N_{h'}, M_{h'+1}]}$, and $d = D_k(r-r', t-t'-1)$ (Notice the definition of the function h preceding the lemma and the observation that

$$h(D_k(r-r', t-t'-1), k-1) \leq D_k(r-r', t-t').)$$

We obtain functions f_{ij}, f_{ij}^+ with properties (1)–(3) because:

(1) and (2) coincide with the corresponding items in the lemma; and (3) is guaranteed by (3) of the lemma because if $\boldsymbol{p}$ is a pair of (old) corresponding elements in a row between $N_{h'-1} - D_k(r-r', t-t'-1)$ and $M_{h'+1} + D_k(r-r', t-t'-1)$ (so it might be "near" an $\boldsymbol{a}_{ij}$ in question) then either $\mathrm{row}(\boldsymbol{p}) \in [N_{h'-1}, M_{h'}]$ hence $f_{\boldsymbol{p}} = f_{[N_{h'-1}, M_{h'}]} = f_*$ or $\mathrm{row}(\boldsymbol{p}) \in [N_{h'}, M_{h'+1}]$ hence $f_{\boldsymbol{p}} = f_{[N_{h'}, M_{h'+1}]} = f^*$.

We do this for each $h' \leq h$. Hence on the intersections of the intervals $[N_{h'-1}, M_{h'+1}]$ $(1 \leq h' \leq h)$ the definitions of those f_{ij} and f^+_{ij} with $\mathrm{row}(\boldsymbol{a}_{ij})$ in this intersection should coincide. But obviously they do because the intersections are intervals $[N_{h'}, M_{h'+1}]$ $(0 \leq h' \leq h)$ and for $\boldsymbol{a}_{ij}$ with rows in such intervals f_{ij} and f^+_{ij} are equal to $f_{[N_{h'}, M_{h'+1}]}$ anyway (by (3) of the lemma).

Hence the claim is proved (modulo Lemma 6.19).

¬TC-move is treated analogously.

■

Before formulating the missing lemma we define a new function h by $h(d,0) := d+2$ and $h(d,l) := 3(l+1)h(d,l-1)$. An easy induction shows that for all $d, l \geq 0$ we have $h(D_k(r-r', t-t'-1), k-1) \leq D_k(r-r', t-t')$. The intuitive reason behind introducing this additional function to calculate distances is that now the parameter l denoting the length of the tuples in our paths becomes important (whereas in D_k it is estimated by its upper bound $k-1$).

6.19 Lemma. *Let $d \geq 1$, $0 \leq l \leq k-1$, $s \leq t(k-1) - l$, $f_*, f^* \in \mathcal{F}^s$ and $m \geq 2$. Let $\overset{l}{\boldsymbol{a}}_1, \ldots, \overset{l}{\boldsymbol{a}}_m \in \mathfrak{H}$ and $M, N \in \{1, \ldots, n\}$ such that $N - M \geq h(d,l)$ and $\mathrm{row}(\boldsymbol{a}_{1j})$, $\mathrm{row}(\boldsymbol{a}_{mj}) \notin\,]M, N[$ for all $j \leq l$.*

Then there exist automorphisms $f_{ij} \in \mathcal{F}^{s+l}$ $(1 \leq i \leq m-1, j \leq l)$ and $f^+_{ij} \in \mathcal{F}^{s+l}$ $(2 \leq i \leq m, j \leq l)$ of $\mathfrak{H}$ such that

(1) $f_{ij}(\boldsymbol{a}_{ij}) = f^+_{ij}(\boldsymbol{a}_{ij})$ (for all possible i, j)

(2) If $d(\boldsymbol{a}_{ij}, \boldsymbol{a}_{ij'}) \leq d$ then $f_{ij} = f_{ij'}$ and $f^+_{ij} = f^+_{ij'}$.
If $d(\boldsymbol{a}_{ij}, \boldsymbol{a}_{i+1j'}) \leq d$ then $f_{ij} = f^+_{i+1j'}$.
(for all possible i, j, j')

(3) If $\mathrm{row}(\boldsymbol{a}_{ij}) - M \leq d$ then $f_{ij} = f^+_{ij} = f_*$.
If $N - \mathrm{row}(\boldsymbol{a}_{ij}) \leq d$ then $f_{ij} = f^+_{ij} = f^*$.
(for all possible i, j)

Proof. The proof is harder than it seems at first sight, although the idea is relatively simple: Consider for a moment the case $l = 1$. If the path $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_m$ starts in a row below M, then crosses the interval $]M, N[$, and ends up above N, there will be $\boldsymbol{a}_i$ and $\boldsymbol{a}_j$ such that $\boldsymbol{a}_i$ is near

M and $\boldsymbol{a}_j$ is near N. By the homogeneity property of structure $\mathfrak{H}$ we find an automorphism f such that $f(\boldsymbol{a}_i) = f_*(\boldsymbol{a}_i)$ and $f(\boldsymbol{a}_j) = f^*(\boldsymbol{a}_j)$. We define the $f_{i'j'}$ to be f_* if $i' \leq i$, f if $i \leq i' < j$, and f^* if $i' \geq j$. Then the path switches its component on the way from $\boldsymbol{a}_i$ to $\boldsymbol{a}_j$.
Note that the situation may be more complicated. For example, the path may cross the interval $]M, N[$ more than once. To deal with this, we have to cut the path into short, "simple" pieces. The most serious problem, however, is that we are dealing with a path of l–tuples, and each component may behave completely differently. Therefore our proof is by induction on l:

For $l = 0$ there is nothing to prove.

Suppose $l \geq 1$ and the lemma is proved for $l - 1$.
The idea of the following induction step is a (more or less) straightforward generalization of a proof of the statement for $l = 1$. So while checking the proof it might be helpful to keep the situation for $l = 1$ in mind. The idea becomes much clearer in this case because we do not need the disturbing applications of the induction hypothesis. (Formally we need the induction hypothesis $l = 0$ but it is trivial.)

Step 1: Subdividing $[1, m]$ into Simple Intervals

We subdivide $[M, N]$ into three parts having the following idea in mind:
As long as there is at least one element in the middle there can be at most $(l - 1)$ elements in the bottom or in the top part and we can apply the induction hypothesis there. If nothing is in the middle there is no danger because bottom and top are not connected.
We let $M' := M + (l + 1)h(d, l - 1)$ and $N' := N - (l + 1)h(d, l - 1)$ and note that $N' - M' \geq (l + 1)h(d, l - 1)$.

In a sense we are mainly interested in the j functions (mapping $\{1, \ldots, m\}$ to $\{1, \ldots, n\}$) defined by $i \mapsto \operatorname{row}(\boldsymbol{a}_{ij})$.

We subdivide $[1, m]$ into intervals such that either for all i in such an interval there exists a j with $\operatorname{row}(\boldsymbol{a}_{ij})$ in the middle (i.e. in $]M', N'[$) or for no i in the interval there exists such a j. We let

$$\begin{aligned}
x_1 &:= 1 \\
x_{2h} &:= \min \big\{ i \mid x_{2h-1} \leq i < m,\ \forall j \leq l : \operatorname{row}(\boldsymbol{a}_{ij}) \notin]M', N'[, \\
&\qquad\qquad \exists j \leq l : \operatorname{row}(\boldsymbol{a}_{(i+1)j}) \in]M', N'[\ \big\} \\
x_{2h+1} &:= \min \big\{ i \mid x_{2h} < i \leq m,\ \forall j \leq l : \operatorname{row}(\boldsymbol{a}_{ij}) \notin]M', N'[, \\
&\qquad\qquad \exists j \leq l : \operatorname{row}(\boldsymbol{a}_{(i-1)j}) \in]M', N'[\ \big\}
\end{aligned}$$

and h_{max} be the maximal number such that $x_{h_{max}}$ is defined. Thus for all $h < \frac{h_{max}}{2}$ we have:

$$\begin{aligned}
&\forall i \in [x_{2h-1}, x_{2h}]\ \forall j \leq l :\ \operatorname{row}(\boldsymbol{a}_{ij}) \notin]M', N'[\\
&\forall i \in]x_{2h}, x_{2h+1}[\ \exists j \leq l :\ \operatorname{row}(\boldsymbol{a}_{ij}) \in]M', N'[\ .
\end{aligned}$$

We call $[x_{2h-1}, x_{2h+1}]$ a *simple interval.* (see Figure 6.20)

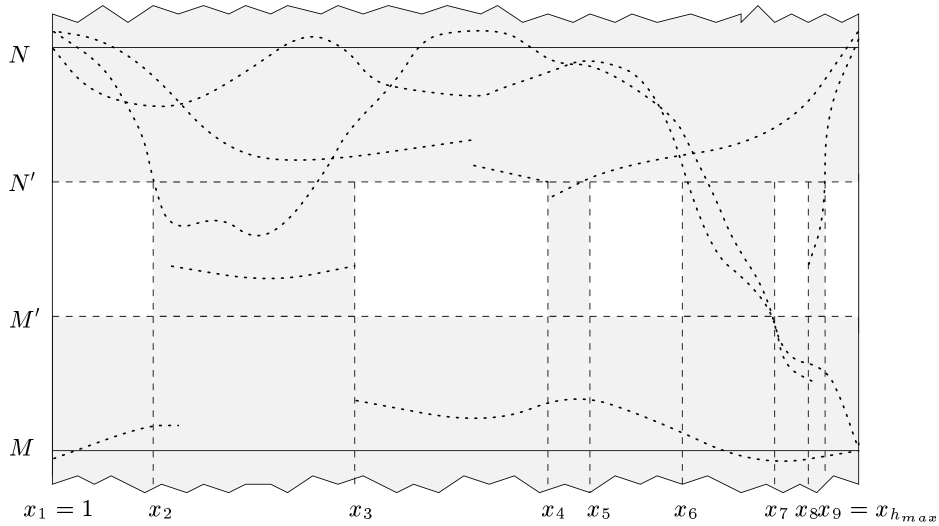

6.20 Figure. *This figure illustrates the functions $i \mapsto \mathrm{row}(\boldsymbol{a}_{ij})$. The range of these functions is somewhere in the grey area (or outside the picture). The dotted lines indicate one possible course of the functions.*
But note that the paths need not be connected at all. On the other hand a reasonably playing spoiler might want to choose paths which are somehow connected.

Before we consider the simple intervals we have to look at the rightmost interval $[x_{h_{max}}, m]$. Note that h_{max} is odd since $\forall j \leq l : \mathrm{row}(\boldsymbol{a}_{mj}) \notin]M, N[\supseteq]M', N'[$. Hence for all $i \geq x_{h_{max}}$, $j \leq l$ we have $\mathrm{row}(\boldsymbol{a}_{ij}) \notin]M', N'[$.
If $x_{h_{max}} < m$ then for each $i \in [x_{h_{max}}, m[$ we let

$$f_{ij} := \begin{cases} f_* & if \quad \mathrm{row}(\boldsymbol{a}_{ij}) \leq M' \\ f^* & if \quad \mathrm{row}(\boldsymbol{a}_{ij}) \geq N' \end{cases}$$

and

$$f^+_{i+1j} := \begin{cases} f_* & if \quad \mathrm{row}(\boldsymbol{a}_{(i+1)j}) \leq M' \\ f^* & if \quad \mathrm{row}(\boldsymbol{a}_{(i+1)j}) \geq N' \end{cases} .$$

We are finished once we have proved:

THE STATEMENT FOR SIMPLE INTERVALS: For each $h \geq 1$ such that $2h + 1 \leq h_{max}$ there exist automorphisms $f_{ij} \in \mathcal{F}^{t'+l}$ $(x_{2h-1} \leq i < x_{2h+1}, j \leq l)$ and $f^+_{ij} \in \mathcal{F}^{t'+l}$ $(x_{2h-1} < i \leq x_{2h+1}, j \leq l)$ of $\mathfrak{H}$ such that (1)–(3) hold and in addition

(4) If $\mathrm{row}(\boldsymbol{a}_{x_{2h-1}j}) \leq M'$ then $f_{x_{2h-1}j}(\boldsymbol{a}_{x_{2h-1}j}) = f_*(\boldsymbol{a}_{x_{2h-1}j})$.
If $\mathrm{row}(\boldsymbol{a}_{x_{2h-1}j}) \geq N'$ then $f_{x_{2h-1}j}(\boldsymbol{a}_{x_{2h-1}j}) = f^*(\boldsymbol{a}_{x_{2h-1}j})$.

(5) If $\mathrm{row}(\boldsymbol{a}_{x_{2h+1}j}) \leq M'$ then $f^+_{x_{2h+1}j}(\boldsymbol{a}_{x_{2h+1}j}) = f_*(\boldsymbol{a}_{x_{2h+1}j})$.
If $\mathrm{row}(\boldsymbol{a}_{x_{2h+1}j}) \geq N'$ then $f^+_{x_{2h+1}j}(\boldsymbol{a}_{x_{2h+1}j}) = f^*(\boldsymbol{a}_{x_{2h+1}j})$.

(4) and (5) make it possible to paste together the intervals $[x_{2h-1}, x_{2h+1}]$ to $[1, m]$ afterwards because they guarantee that for the endpoints x_{2h+1} $(1 \leq h \leq \frac{h_{max}}{2})$ of the simple intervals we have $f_{x_{2h+1}j}(\boldsymbol{a}_{x_{2h+1}j}) = f^+_{x_{2h+1}j}(\boldsymbol{a}_{x_{2h+1}j})$ for all $j \leq l$.
(But note that we do not demand $f_{x_{2h+1}j} = f^+_{x_{2h+1}j}$.)

So we consider an interval $[x_{2h-1}, x_{2h+1}]$. To simplify the notation we let $x := x_{2h-1}$, $y := x_{2h}$ and $z := x_{2h+1}$.

Step 2: The Sparse Part

A place $i \in [x, z]$ is called *dense* if $\forall j \leq l : \mathrm{row}(\boldsymbol{a}_{ij}) \in]M', N'[$. Otherwise it is called *sparse*. Let w be the minimal dense place in $[x, z]$; if no such place exists we let $w := z$. Note that we always have $w > y$ by the definition of $y = x_{2h}$.

Recalling that $M' - M,\ N' - M'\ N - N' \geq (l+1)h(d, l-1)$, by the pigeonhole–principle we can find $M_*, N_*, M'', N'', M^*, N^*$ such that

- $M \leq M_* < M_* + h(d, l-1) \leq N_* \leq M'$
- $M' \leq M'' < M'' + h(d, l-1) \leq N'' \leq N'$
- $N' \leq M^* < M^* + h(d, l-1) \leq N^* \leq N$
- $\forall j \leq l:\ \mathrm{row}(\boldsymbol{a}_{wj}) \notin]M'', N''[\wedge \mathrm{row}(\boldsymbol{a}_{zj}) \notin]M_*, N_*[\cup]M^*, N^*[$.

The pairs M_*, N_*, M'', N'', and M^*, N^* respectively will be taken as the M, N needed to apply the induction hypothesis in the following cases. See Figure 6.21 to get an image of our current situation.

Note that we also have

- $\forall j \leq l : \mathrm{row}(\boldsymbol{a}_{xj}) \notin]M'', N''[$
  (because $\forall j \leq l : \mathrm{row}(\boldsymbol{a}_{xj}) \notin]M', N'[$),
- $\forall j \leq l : \mathrm{row}(\boldsymbol{a}_{wj}) \notin]M_*, N_*[\cup]M^*, N^*[$
  (because $\forall j \leq l : \mathrm{row}(\boldsymbol{a}_{wj}) \in]M', N'[$).

Since all $i \in [x, w[$ are sparse there exists a sequence $(j_i)_{x \leq i \leq w-1}$ such that $\mathrm{row}(\boldsymbol{a}_{ij_i}) \notin]M', N'[\supseteq]M'', N''[$. Without loss of generality we assume that $j_i = l$ for all $i \in [x, w[$.

Now we apply the induction hypothesis to

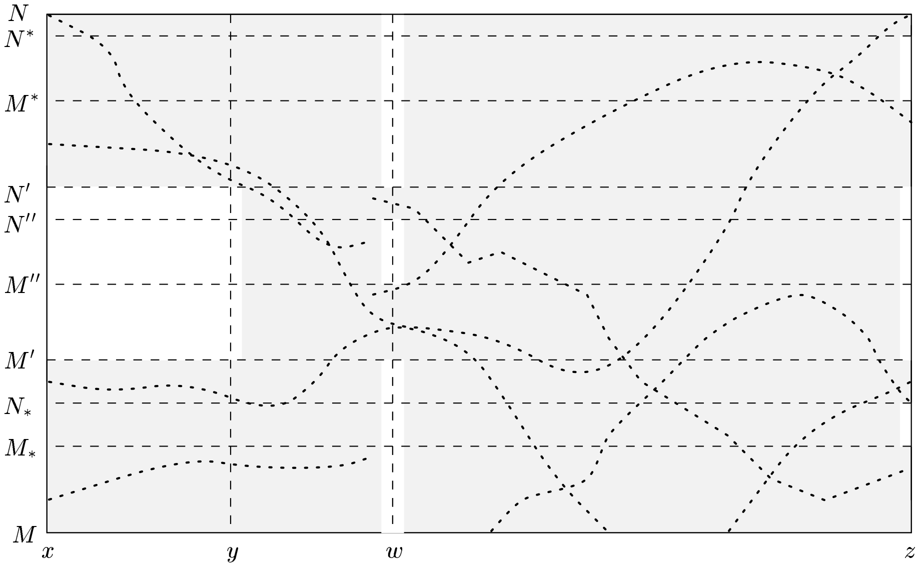

6.21 Figure.

- the path $\overset{l-1}{\boldsymbol{a}}_i$ ($i \in [x, w]$) (i.e. we take $\{x, \ldots, w\}$ instead $\{1, \ldots, m\}$ as index–set)
- the interval $[M'', N'']$ instead of $[M, N]$ (Recall that N'', M'' are chosen in a way that $\forall j \leq l : \mathrm{row}(\boldsymbol{a}_{wj}), \mathrm{row}(\boldsymbol{a}_{xj}) \notin]M'', N''[$ and that $N'' - M'' \geq h(d, l-1)$.)
- d, s, f_*, f^* (unchanged).

We get automorphisms $f_{ij} \in \mathcal{F}^{s+l-1}$ ($x \leq i \leq w-1, j \leq l-1$) and $f^+_{ij} \in \mathcal{F}^{s+l-1}$ ($x+1 \leq i \leq w, j \leq l-1$) of $\mathfrak{H}$ such that

- $f_{ij}(\boldsymbol{a}_{ij}) = f^+_{ij}(\boldsymbol{a}_{ij})$ (for all possible i, j)
- If $d(\boldsymbol{a}_{ij}, \boldsymbol{a}_{ij'}) \leq d$ then $f_{ij} = f_{ij'}$ and $f^+_{ij} = f^+_{ij'}$.
  If $d(\boldsymbol{a}_{ij}, \boldsymbol{a}_{i+1j'}) \leq d$ then $f_{ij} = f^+_{i+1j'}$.
  (for all possible i, j, j').
- If $\mathrm{row}(\boldsymbol{a}_{ij}) - M'' \leq d$ then $f_{ij} = f^+_{ij} = f_*$.
  If $N'' - \mathrm{row}(\boldsymbol{a}_{ij}) \leq d$ then $f_{ij} = f^+_{ij} = f^*$.
  (for all possible i, j).

Furthermore, for each $i \in [x, w[$ we define

$$f_{il} := \begin{cases} f_* & if \quad \mathrm{row}(\boldsymbol{a}_{il}) \leq M'' \\ f^* & if \quad \mathrm{row}(\boldsymbol{a}_{il}) \geq N'' \end{cases}$$

and

$$f^+_{i+1l} := \begin{cases} f_* & if \quad \mathrm{row}(\boldsymbol{a}_{(i+1)l}) \leq M'' \\ f^* & if \quad \mathrm{row}(\boldsymbol{a}_{(i+1)l}) \geq N'' \end{cases} .$$

Note that these automorphisms perfectly fit into the items above , so in fact we have functions f_{ij} and f^+_{ij} for each $j \leq l$ such that the statements above hold.

Thus items (1)–(3) are true for the interval $[x, w]$ (with the functions defined above). Recall that for each $j \leq l$ we have $\mathrm{row}(\boldsymbol{a}_{xj}) \notin]M', N'[$ thus

$$f_{xj} = \begin{cases} f_* & if \quad \mathrm{row}(\boldsymbol{a}_{xj}) \leq M' \\ f^* & if \quad \mathrm{row}(\boldsymbol{a}_{xj}) \geq N' \end{cases}$$

which implies (4).
On the other side we have

$$f^+_{wj} = \begin{cases} f_* & if \quad \mathrm{row}(\boldsymbol{a}_{wj}) \leq M'' \\ f^* & if \quad \mathrm{row}(\boldsymbol{a}_{wj}) \geq N'' \end{cases}$$

(for each $j \leq l$).
If $w = z$ this implies (5) because there exists no $j \leq l$ such that $\mathrm{row}(\boldsymbol{a}_{zj}) \in]M', N'[$; thus we are finished in this case.
So we assume $w < z$ and consider the interval $[w, z]$.

Step 3: The Dense Part

By Theorem 6.8(4) applied to the sets

$$\begin{aligned} A &:= \{\boldsymbol{a}_{ij} \mid i \in \{w, z\}, j \leq l, \mathrm{row}(\boldsymbol{a}_{ij}) \in [N_*, M'']\} \\ B &:= \{\boldsymbol{a}_{ij} \mid i \in \{w, z\}, j \leq l, \mathrm{row}(\boldsymbol{a}_{ij}) \in [N'', M^*]\} \end{aligned}$$

there exists an $f \in \mathcal{F}^{s+1}$ such that

$$f(\boldsymbol{a}_{wj}) = \begin{cases} f_*(\boldsymbol{a}_{wj}) & if \quad \mathrm{row}(\boldsymbol{a}_{wj}) \in]M', M''] \\ f^*(\boldsymbol{a}_{wj}) & if \quad \mathrm{row}(\boldsymbol{a}_{wj}) \in [N'', N'[\end{cases}$$

and

$$f(\boldsymbol{a}_{zj}) = \begin{cases} f_*(\boldsymbol{a}_{zj}) & if \quad \mathrm{row}(\boldsymbol{a}_{zj}) \in [N_*, M'] \\ f^*(\boldsymbol{a}_{zj}) & if \quad \mathrm{row}(\boldsymbol{a}_{zj}) \in [N', M^*] \end{cases} .$$

Here we use that

$$\forall j \leq l : \mathrm{row}(\boldsymbol{a}_{wj}) \in]M', N'[\wedge \mathrm{row}(\boldsymbol{a}_{zj}) \notin]M', N'[$$

to guarantee that there are at most $l \leq (k-1)$ elements of $A \cup B$ in a row.
Moreover, $N'' - M'' > 1$ so for all $\boldsymbol{a} \in A, \boldsymbol{b} \in B$ we have $\mathrm{row}(\boldsymbol{b}) - \mathrm{row}(\boldsymbol{a}) > 1$.

Recall that for each $i \in]y, z[\supseteq [w, z[$ there exists a $j_i \leq l$ such that $\mathrm{row}(\boldsymbol{a}_{ij_i}) \in]M', N'[$ (this is by the definition of $y = x_{2h}$ and $z = x_{2h+1}$). Without loss of generality we can assume that $j_i = l$.

For each $i \in [w, z[$ we let $f_{il} := f$, and for each $i \in]w, z[$ we let $f_{il}^+ := f$.
Furthermore, $\mathrm{row}(\boldsymbol{a}_{zl})$ is either $\leq M_*$ or $\geq N^*$ or contained in one of the intervals $[N_*, M']$, $[N', M^*]$. We let $f_{zl}^+ = f_*$ if $\mathrm{row}(\boldsymbol{a}_{zl}) \leq M_*$, $f_{zl}^+ = f$ if $\mathrm{row}(\boldsymbol{a}_{zl}) \in [N_*, M'] \cup [N', M^*]$, and $f_{zl}^+ = f^*$ if $\mathrm{row}(\boldsymbol{a}_{zl}) \geq N^*$.

Step 4: Final Applications of the Induction Hypothesis

The situation we are in now is the most difficult. To cope with it we have to apply the induction hypothesis twice:
Once to

- the path of (at most) $(l-1)$–tuples consisting of the elements $\boldsymbol{a}_{ij}$, $j \leq l-1$, $i \in [w, z]$ with $\mathrm{row}(\boldsymbol{a}_{ij}) \leq M^*$
- the interval $[M_*, N_*]$ instead of $[M, N]$
- f instead of f^*
- $s+1$ instead of s
- d, f_* (unchanged)

and once to

- the path of (at most) $(l-1)$–tuples consisting of the elements $\boldsymbol{a}_{ij}$, $j \leq l-1$, $i \in [w, z]$ with $\mathrm{row}(\boldsymbol{a}_{ij}) \geq N_*$
- the interval $[M^*, N^*]$ instead of $[M, N]$
- f instead of f_*
- $s+1$ instead of s
- d, f^* (unchanged).

What we get are automorphisms $f_{ij} \in \mathcal{F}^{s+1+l-1}$ $(w \leq i \leq z-1, j \leq l-1)$ and $f_{ij}^+ \in \mathcal{F}^{s+l}$ $(w+1 \leq i \leq z, j \leq l-1)$ of $\mathfrak{H}$ such that

- $f_{ij}(\boldsymbol{a}_{ij}) = f_{ij}^+(\boldsymbol{a}_{ij})$ (for all possible i, j)
- If $d(\boldsymbol{a}_{ij}, \boldsymbol{a}_{ij'}) \leq d$ then $f_{ij} = f_{ij'}$ and $f_{ij}^+ = f_{ij'}^+$.
  If $d(\boldsymbol{a}_{ij}, \boldsymbol{a}_{i+1j'}) \leq d$ then $f_{ij} = f_{i+1j'}^+$.
  (for all possible i, j, j').
- If $\mathrm{row}(\boldsymbol{a}_{ij}) - M_* \leq d$ then $f_{ij} = f_{ij}^+ = f_*$.
  If $N^* - \mathrm{row}(\boldsymbol{a}_{ij}) \leq d$ then $f_{ij} = f_{ij}^+ = f^*$.
  (for all possible i, j).

- If $N_* - d \leq \mathrm{row}(\boldsymbol{a}_{ij}) \leq M^* + d$ then $f_{ij} = f_{ij}^+ = f$
  (for all possible i, j).

Note that actually the two applications of the induction hypothesis might deliver different functions on their "intersection". In this case our functions would not be well–defined. However, writing out the results of both applications separately, the reader can easily check that this is not the case, and hence that our functions are well–defined.
Also note that by the definition of the $f_{il}^{(+)}$ in fact we have functions f_{ij} and f_{ij}^+ for each $j \leq l$ such that the statements above hold.
Clearly these statements guarantee (1)–(3) for the interval $[w, z]$.
Since for each $j \leq l$

$$f_{zj}^+ := \begin{cases} f_* & if \quad \mathrm{row}(\boldsymbol{a}_{zj}) \leq M_* \\ f & if \quad N_* \leq \mathrm{row}(\boldsymbol{a}_{zj}) \leq M^* \\ f^* & if \quad \mathrm{row}(\boldsymbol{a}_{zj}) \geq N^* \end{cases}$$

and, by the definition of f,

$$f(\boldsymbol{a}_{zj}) := \begin{cases} f_*(\boldsymbol{a}_{zj}) & if \quad \mathrm{row}(\boldsymbol{a}_{zj}) \leq M' \\ f^*(\boldsymbol{a}_{zj}) & if \quad \mathrm{row}(\boldsymbol{a}_{zj}) \geq N' \end{cases}$$

we also obtain (5).

Finally we note that our definitions for the intervals $[x, w]$ and $[w, z]$ match ; i.e. for all $j \leq l$ we have

$$f_{wj}(\boldsymbol{a}_{wj}) = f(\boldsymbol{a}_{wj}) = \left\{ \begin{array}{c} \overbrace{f_*(\boldsymbol{a}_{wj})}^{if\ \mathrm{row}(\boldsymbol{a}_{wj}) \leq M''} \\ \underbrace{f^*(\boldsymbol{a}_{wj})}_{if\ \mathrm{row}(\boldsymbol{a}_{wj}) \geq N''} \end{array} \right\} = f_{wj}^+(\boldsymbol{a}_{wj}) \ .$$

(The first equality holds because $\mathrm{row}(\boldsymbol{a}_{wj}) \in]M', N'[$, for the last cf. Step 2.) ■

7 Concluding Remarks

We have seen that the arity hierarchies are strict for fixed–point logics, datalogics, implicitly definable queries and transitive closure logics. Moreover, methods similar to those we use here apply to establish arity hierarchy results for Lindström quantifiers (cf. [GH95]).

The main open problems concerning arities can be arranged according to the following two items:

7.1 Second–order logic. The most important result obtained here is Ajtai's theorem [Ajt83] mentioned in Remark 5.2 which implies that the arity hierarchies of Σ_1^1 and Π_1^1 are strict (but *not* for a uniform signature). Considerable efforts have been made to strengthen Ajtai's theorem in two directions:

- It is still open whether the arity hierarchies of Σ_1^1 or Π_1^1 are strict on any uniform signature. Actually, it is not even known whether one binary existential second–order quantifier suffices to capture whole second–order logic on the class of finite graphs.
- There are no known arity results concerning larger fragments of second–order logic or the whole logic. An approach using diagonalization was taken by Makowsky and Pnueli [MP94], but they could only prove a hierachy result incorporating both arities and alternations of quantifiers.

Let me remark that the method developed here seems to apply to a sublogic $\mathbf{SO}^\omega$ of second–order logic which has recently been introduced by Dawar [Daw95].

7.2 Ordered structures. Ordered finite structures play an important role in descriptive complexity theory, so it would be nice to establish the arity hierarchies being strict on classes of ordered structures. Note that the structures we used in our proofs are far from being ordered because they have many automorphisms.

In fact, the situation looks a bit different on ordered structures. It is not too hard to see that we have $\mathbf{TC} \subseteq \mathbf{PFP}^1$ (cf. [Imh94]) there, so our main lemma fails. On the other hand, Imhof [Imh94] showed that the arity hierarchy of $\mathbf{PFP}$ is strict even on ordered structures.

The analogous question for $\mathbf{LFP}$ and $\mathbf{IFP}$ is harder: Recall that Theorem 6.3 said that $\mathbf{DTC} \subseteq \mathbf{LFP}^2$. Now, if the arity hierarchies are strict, in fact if $\mathbf{LFP} \neq \mathbf{LFP}^2$ on ordered structures, we have $\mathbf{DTC} \neq \mathbf{LFP}$ and hence the complexity classes LOGSPACE and PTIME are distinct (using well–known results from descriptive complexity theory). On the other hand, if the hierarchies collapses for $\mathbf{LFP}$ then $\mathbf{LFP} \neq \mathbf{PFP}$ and hence PTIME$\neq$PSPACE.

It is also still open whether the arity hierarchies of $\mathbf{TC}$ and $\mathbf{DTC}$ are strict on ordered structures.

Let us close this article with a nice application of Büchi's theorem (due to Makowsky) which shows that at least the first step of the hierarchies of $\mathbf{DTC}$, $\mathbf{TC}$, and $\mathbf{LFP}$ is strict:

Büchi's theorem says that on word–structures (orders with unary relations; they can be considerd as words) the class of monadic second–order–definable queries coincides with the regular languages. Now the class of words of the form $a^n b^n$ $(n \geq 1)$ can be easily defined by a $\mathbf{DTC}^2$–formula, but it does not form a regular language. Since $\mathbf{LFP}^1$ is contained in monadic second–order logic this implies $\mathbf{DTC}^2 \not\subseteq \mathbf{LFP}^1$.

Acknowledgements

I would like to thank Heinz-Dieter Ebbinghaus, Bernhard Herwig, Ian Hodkinson, and the anonymous referee for valuable advise.

References

[AC89] F. Afrati and S. Cosmadakis. Expressiveness of restricted recursive queries. In *Proceedings of 21st ACM symposium on Theory of Computing*, pages 113–126, 1989.

[Ajt83] M. Ajtai. Σ_1^1-formulae on finite structures. *Annals of pure and applied logic*, 24:1–48, 1983.

[AS92] N. Alon and J. Spencer. *The Probabilistic Method.* Wiley, 1992.

[AV91] S. Abiteboul and V. Vianu. Datalog extensions for database queries and updates. *Journal of Computer and System Sciences*, 43:62–124, 1991.

[CH82] A. Chandra and D. Harel. Structure and complexity of relational queries. *Journal of Computer and System Sciences*, 25:99–128, 1982.

[CH85] A. Chandra and D. Harel. Horn clause queries and generalizations. *Journal of Logic Programming*, 1:1–15, 1985.

[CM92] A. Calo and J.A. Makowsky. The Ehrenfeucht-Fraïssé game for transitive closure. In A. Nerode and M. Taitsin, editors, *Logical Foundation of Computer Science - Tver'92*, volume 620 of *Lecture Notes in Computer Science*, pages 57–68. Springer-Verlag, 1992.

[Daw95] A. Dawar. A restricted second order logic for finite structures. In D. Leivant, editor, *Logic and computational complexity : international workshop, LCC '94*, volume 960 of *Lecture Notes in Computer Science*, 1995.

[DM89] P. Dublish and S. Maheshwari. Expressibility of bounded arity fixed-point queries. In *Proceedings of 8th ACM symposium on principles of database systems*, pages 324–335, 1989.

[EF95] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory.* Springer-Verlag, 1995.

[EFT94] H.-D. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic.* Springer-Verlag, 2nd edition, 1994.

[Erd59] P. Erdös. Graph theory and probability. *Canadian Journal of Mathematics*, 11:34–38, 1959.

[Grä91] E. Grädel. On transitive closure logic. In *Proceedings of the 5th workshop on computer science logic*, volume 626 of *Lecture Notes in Computer Science*, pages 149–163. Springer-Verlag, 1991.

[Gro93] M. Grohe. Bounded arity hierarchies in fixed–point logics. In *Proceedings of the 7th workshop on computer science logic*, volume 832 of *Lecture Notes in Computer Science*, pages 150–164. Springer-Verlag, 1993.

[Gro94] M. Grohe. *The Structure of Fixed–Point Logics.* PhD thesis, Albert-Ludwigs Universität Freiburg, 1994.

[GH95] M. Grohe and L. Hella. A double arity hierarchy theorem for transitive closure logic, 1995. To appear in *Archive for Mathematical Logic.*

[GS86] Y. Gurevich and S. Shelah. Fixed point extensions of first–order logic. *Annals of pure and applied logic*, 32:265–280, 1986.

[Hel92] L. Hella. Logical hierarchies in PTIME. In *Proceedings of the 7th IEEE Symposium in Logic and Computer science*, pages 360–368, 1992.

[Her95] B. Herwig. Extending partial isomorphisms on finite structures. *Combinatorica*, 15:365–372, 1995.

[Hru92] E. Hrushovski. Extending partial isomorphisms of graphs. *Combinatorica*, 12:411–416, 1992.

[Imh94] H. Imhof. The arity problem for fixed point logic – a complexity theoretical point of view. Research report, Albert-Ludwigs-Universität Freiburg, 1994.

[Imm87] N. Immerman. Languages that capture complexity classes. *SIAM Journal of Computing*, 16:760–778, 1987.

[Kan90] P.C. Kanellakis. Elements of relational database theory. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*. Elsevier Science Publishers B.v., 1990.

[Kol90] Ph.G. Kolaitis. Implicit definability on finite structures and unambiguous computations. In *Proceedings of the 5th IEEE Symposium on Logic in Computer science*, 1990.

[MP94] J.A. Makowsky and Y.B. Pnueli. Arity vs. alternation in second order logic. In *Logical Foundations of Computer Science – St. Petersburg 1994*, volume 813 of *Lecture Notes in Computer Science*, 1994.

Martin Grohe
Insitut für mathematische Logik
Eckerstr. 1
79104 Freiburg
Germany

Email:grohe@sun1.mathematik.uni-freiburg.de