

Infinitary Logic and Inductive Definability over Finite Structures*

Anuj Dawar[†]

Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104
a.dawar@swansea.ac.uk

Steven Lindell[‡]

Department of Mathematics
Haverford College
Haverford, PA 19041
slindell@acc.haverford.edu

Scott Weinstein[§]

Department of Philosophy
University of Pennsylvania
Philadelphia, PA 19104
weinstein@cis.upenn.edu

December, 1991

Revised January, 1994

*We would like to thank Yuri Gurevich, Phokion Kolaitis, Moshe Vardi, and three anonymous referees for valuable comments on earlier versions of this paper. This paper has been accepted for publication in *Information and Computation*.

[†]Supported in part by ONR grant N00014-89-J-1725 and ARO grant DAAL 03-89-C-0031. *Current address*: Dept. of Computer Science, University College of Swansea, Swansea SA2 8PP, U.K.

[‡]Supported in part by NSF grant CCR-9003356.

[§]Supported in part by ONR grant N00014-89-J-1725.

Abstract

The extensions of first-order logic with a least fixed point operator (FO + LFP) and with a partial fixed point operator (FO + PFP) are known to capture the complexity classes P and PSPACE respectively in the presence of an ordering relation over finite structures. Recently, Abiteboul and Vianu [Abiteboul and Vianu, 1991b] investigated the relationship of these two logics in the absence of an ordering, using a machine model of generic computation. In particular, they showed that the two languages have equivalent expressive power if and only if $P = PSPACE$. These languages can also be seen as fragments of an infinitary logic where each formula has a bounded number of variables, $L_{\infty\omega}^{\omega}$ (see, for instance, [Kolaitis and Vardi, 1990]). We investigate this logic of finite structures and provide a normal form for it. We also present a treatment of the results in [Abiteboul and Vianu, 1991b] from this point of view. In particular, we show that we can write a formula of FO + LFP that defines an ordering of the $L_{\infty\omega}^k$ types uniformly over all finite structures. One consequence of this is a generalization of the equivalence of FO + LFP and P from ordered structures to classes of structures where every element is definable. We also settle a conjecture mentioned in [Abiteboul and Vianu, 1991b] by showing that FO + LFP is properly contained in the polynomial time computable fragment of $L_{\infty\omega}^{\omega}$, raising the question of whether the latter fragment is a recursively enumerable class.

1 Introduction

In applications of finite model theory in computer science, extensions of first-order logic by various induction operations have received particular attention. Many database query languages are based on such extensions (see, for instance, [Chandra and Harel, 1982], [Vardi, 1982] and [Abiteboul and Vianu, 1991a] and the references therein) and in the area of descriptive complexity, they have been shown to naturally characterize certain complexity classes. In particular, the extensions of first-order logic with a least fixed point operator (FO + LFP) and with a partial fixed point operator (FO + PFP) are known to capture the complexity classes P and PSPACE respectively in the presence of an ordering relation [Immerman, 1986; Vardi, 1982; Abiteboul and Vianu, 1991a]. Recently, Abiteboul and Vianu [Abiteboul and Vianu, 1991b] investigated the relationship of these two logics in the absence of an ordering, using a machine model of generic computation. In particular, they showed that the two languages have equivalent expressive power if and only if $P = PSPACE$.

The languages FO + LFP and FO + PFP can also be seen as fragments of an infinitary logic, $L_{\infty\omega}^\omega$, where each formula has a bounded number of variables. The logic $L_{\infty\omega}^\omega$ was introduced in [Barwise, 1977]. Rubin [Rubin, 1975] showed that, over fixed infinite structures, FO + LFP is a fragment of this language, and Kolaitis and Vardi [Kolaitis and Vardi, 1992b] proved a similar result for FO + LFP and FO + PFP over the class of all finite structures. Kolaitis and Vardi ([Kolaitis and Vardi, 1990], a full version of this paper appears as [Kolaitis and Vardi, 1992b]) undertook a systematic study of the logic $L_{\infty\omega}^\omega$ over finite structures and established a 0-1 law for it, a generalization of the 0-1 law for FO + LFP [Blass *et al.*, 1985]. They also indicate how this logic implicitly underlies earlier work in descriptive complexity theory [Immerman, 1982; Immerman and Lander, 1990; Cai *et al.*, 1989].

Inspired by [Kolaitis and Vardi, 1990], we posed the question whether every finite structure is characterized up to L^k -equivalence by a single sentence of L^k (first-order logic with only k variables). An adaptation of a technique from Scott [Scott, 1965] answers this question in the affirmative. This yields a normal form for $L_{\infty\omega}^k$ that improves results in [Kolaitis and Vardi, 1992b] and [Kolaitis and Vardi, 1992a]. This technique also finds other applications in the study of $L_{\infty\omega}^\omega$ over finite structures. We present a treatment of the results in

[Abiteboul and Vianu, 1991b] from this point of view. In particular, we show that we can write a formula of FO + LFP that defines an ordering of the $L_{\infty\omega}^k$ types in any structure. This is a refinement of the technique in [Abiteboul and Vianu, 1991b], where a distinct ordering was used for every query. The proofs we present make no reference to a particular model of computation and, it is hoped, shed some light on these results.

We also settle a conjecture mentioned in [Abiteboul and Vianu, 1991b] by showing that FO + LFP is properly contained in the polynomial time computable fragment of $L_{\infty\omega}^\omega$ (this conjecture was also independently settled by [Cosmadakis, 1991] and [Abiteboul and Vianu, 1992]). This raises the question of whether the latter fragment is a recursively enumerable class. We give some indication of how this question might be addressed in Section 9.

This paper is organized as follows. In Section 2, we define the logics FO + LFP and FO + PFP. In Sections 3 and 4, we introduce infinitary logic and some related technical tools, including the Scott construction and a normal form for $L_{\infty\omega}^k$. Sections 5 and 6 establish that types in $L_{\infty\omega}^k$ can be uniformly defined and ordered in FO + LFP and some consequences of this fact. This construction is then used in Sections 7 and 8 to investigate the relationship of FO + LFP and FO + PFP, including the proofs of the results of [Abiteboul and Vianu, 1991b].

Definitions and Notation

A *signature* (also sometimes called a *language* or a *vocabulary*) σ is a finite sequence of relation and constant symbols $\langle R_1, \dots, R_m, c_1, \dots, c_n \rangle$. Associated with each relation symbol, R_i is an arity a_i .

A structure over the the signature σ , $\mathfrak{A} = \langle A, R_1^{\mathfrak{A}}, \dots, R_m^{\mathfrak{A}}, c_1^{\mathfrak{A}}, \dots, c_n^{\mathfrak{A}} \rangle$ consists of a set A , the universe of the structure, relations $R_i^{\mathfrak{A}} \subseteq A^{a_i}$ interpreting the relation symbols in σ and distinguished elements $c_1^{\mathfrak{A}}, \dots, c_n^{\mathfrak{A}}$ of A interpreting the constant symbols. Unless otherwise mentioned, all structures we will be dealing with are assumed to have a finite universe. For convenience, we will assume that the universe A is an initial segment of the natural numbers (note that this does not imply that any of the properties of the natural numbers, such as their linear ordering, are available in the logical vocabulary). We will also write $|\mathfrak{A}|$ for the universe of the structure \mathfrak{A} .

An (m -ary) *query* q (also sometimes called a global relation) is a map from structures

(over some fixed signature σ) to (m -ary) relations on the structures, that is closed under isomorphism. That is, if $\langle a_1 \dots a_m \rangle \in q(\mathfrak{A})$, and f is an isomorphism from \mathfrak{A} to \mathfrak{B} , then $\langle f(a_1) \dots f(a_m) \rangle \in q(\mathfrak{B})$. A Boolean query is a collection of structures K closed under isomorphism.

We will write FO , $\text{FO} + \text{LFP}$, *etc.* both to denote logics (*i.e.* sets of formulas) and the classes of queries that are expressible in the respective logics. It will be clear from the context which usage is intended.

2 Inductive Logic

In the context of finite models, the expressive power of first-order logic is known to be extremely limited. Various extensions of first-order logic have been studied that correspond to independently defined complexity classes. One way of increasing the expressive power of first-order logic is by adding some kind of induction operation.

Let $\phi(R, x_1, \dots, x_k)$ be a first-order formula over the signature $\sigma \cup \{R\}$ with free variables x_1, \dots, x_k , where k is the arity of R . For any structure \mathfrak{A} over the signature σ , ϕ defines a mapping, Φ on relations of arity k in the following sense — given a relation $R^{\mathfrak{A}} \subseteq |\mathfrak{A}|^k$, let $\langle \mathfrak{A}, R^{\mathfrak{A}} \rangle$ be the expansion of \mathfrak{A} interpreting R as $R^{\mathfrak{A}}$. Then, $\Phi(R^{\mathfrak{A}}) = \{ \langle a_1, \dots, a_k \rangle \mid \langle \mathfrak{A}, R^{\mathfrak{A}} \rangle \models \phi[a_1, \dots, a_k] \}$

This map Φ is called *monotone* if for any relations R and S such that $R \subseteq S$, $\Phi(R) \subseteq \Phi(S)$. A map that is monotone has a least fixed point, *i.e.* a smallest relation R such that $\Phi(R) = R$. Moreover, this least fixed point can be obtained by the following iterative construction: Let $\Phi^0 = \emptyset$ and $\Phi^{m+1} = \Phi(\Phi^m)$. Then for some m (depending on the structure \mathfrak{A}), $\Phi^{m+1} = \Phi^m =$ the least fixed point of Φ . The least such m is called the *closure ordinal* of Φ on the structure \mathfrak{A} . If n is the size of \mathfrak{A} , then there are n^k k -tuples in \mathfrak{A} and since Φ is monotone, $m \leq n^k$.

A sufficient syntactic condition for the formula ϕ to define a monotone map on all structures is that ϕ be *positive* in R , that is to say that all occurrences of R in ϕ be within the scope of an even number of negations. We can now define the logic $\text{FO} + \text{LFP}$ over signature σ as the smallest set of formulas satisfying:

- if ϕ is a first-order formula over σ , then $\phi \in \text{FO} + \text{LFP}(\sigma)$,

- if ϕ is formed from formulas in $\text{FO} + \text{LFP}(\sigma)$ by conjunction, disjunction, negation and first-order quantification, then $\phi \in \text{FO} + \text{LFP}(\sigma)$, and
- if $\phi \in \text{FO} + \text{LFP}(\sigma \cup \{R\})$, ϕ is positive in R and x_1, \dots, x_k are distinct variables, where k is the arity of R , then $\mathbf{lfp}(R, x_1 \dots x_k)\phi(t_1 \dots t_k) \in \text{FO} + \text{LFP}(\sigma)$ for any terms t_1, \dots, t_k (note that, as there are no function symbols present, the only terms are variables and constants).

The way to read the last clause above is that the operator \mathbf{lfp} binds the second order variable R and the first-order variables x_1, \dots, x_k in ϕ to form a new predicate. This predicate is to be interpreted as the k -ary relation that is the least fixed point of the monotone operator defined by ϕ . This predicate is then evaluated at the elements denoted by the terms t_1, \dots, t_k .

The following normal form result was established in [Immerman, 1986] for formulas of $\text{FO} + \text{LFP}$.

Theorem 1 *In any vocabulary containing at least two constant symbols, every formula in $\text{FO} + \text{LFP}$ is equivalent to a formula $\mathbf{lfp}(R, \bar{x})\phi(\bar{t})$, where ϕ is first-order, over the class of structures in which not all constant symbols are interpreted by the same element.*

For examples of the use of the \mathbf{lfp} operator, see Axioms 4–6 in Section 8.

Alternatively, we can define the language $\text{FO} + \text{IFP}$ which has an operation \mathbf{ifp} (inflationary fixed point) in place of \mathbf{lfp} . In $\mathbf{ifp}(R, x_1 \dots x_k)\phi(t_1 \dots t_k)$, ϕ is not required to be positive in R . The relational expression $\mathbf{ifp}(R, x_1 \dots x_k)\phi$ denotes the fixed point obtained by iterating the operator Φ' given by $\Phi'(R^{\mathfrak{A}}) = \{\langle a_1, \dots, a_k \rangle \mid \langle a_1, \dots, a_k \rangle \in R^{\mathfrak{A}} \text{ or } \langle \mathfrak{A}, R^{\mathfrak{A}} \rangle \models \phi[a_1, \dots, a_k]\} = \Phi(R^{\mathfrak{A}}) \cup R^{\mathfrak{A}}$. This language is equivalent in expressive power to $\text{FO} + \text{LFP}$:

Theorem 2 ([Gurevich and Shelah, 1986]) *A query on finite structures is expressible in $\text{FO} + \text{IFP}$, if and only if, it is expressible in $\text{FO} + \text{LFP}$.*

Immerman [Immerman, 1986] and Vardi [Vardi, 1982] independently showed that when we include a total ordering on the domain as part of the logical vocabulary, the language $\text{FO} + \text{LFP}$ expresses exactly the class of polynomial time computable queries.

Theorem 3 ([Immerman, 1986],[Vardi, 1982]) *$\text{FO} + \text{LFP}$ with ordering = P.*

We saw above how a formula with one free predicate variable defined an operator on relations. This, of course, is true even when the formula is not positive in the predicate variable and the operator, in turn, may or may not be monotone. Moreover, the iterative stages of the operator can still be defined, though they are not guaranteed to converge to a fixed point in the case of non-monotone operators. Let $\phi(R, \bar{x})$ be a formula that defines a (possibly non-monotone) operator Φ . Define the *partial fixed point* of ϕ to be Φ^m for the least m such that $\Phi^{m+1} = \Phi^m$ if such an m exists, and empty otherwise. Because there are only 2^{n^k} sets of k -tuples over a structure of size n , if such an m exists, then $m \leq 2^{n^k}$. We can then define another extension of first-order logic called FO + PFP with a syntax similar to that of FO + LFP except that the **lfp** operation is replaced by **pfp**, which can operate on arbitrary formulas, not just positive ones. We let $\mathbf{pfp}(R, \bar{x})\phi$ denote the partial fixed point of ϕ .

It has been shown in [Abiteboul and Vianu, 1991a] that the language FO + PFP is equivalent to the query language *while* – an extension of first-order logic with an iterative operation. Putting this together with a result of Vardi [Vardi, 1982], we get the following:

Theorem 4 ([Vardi, 1982],[Abiteboul and Vianu, 1991a]) FO + PFP *with ordering* = PSPACE.

3 Infinitary Logic

We first define the syntax of full infinitary logic. This language is denoted $L_{\infty\omega}$, the first subscript indicating that conjunctions and disjunctions can be taken over arbitrary sets of formulas and the second subscript that only finite quantifier blocks are allowed. The formulas of $L_{\infty\omega}$ are defined as for first-order logic, except that conjunction and disjunction are no longer binary operations. Rather, for any set of infinitary formulas Φ , $\bigvee \Phi$ and $\bigwedge \Phi$ are both formulas of $L_{\infty\omega}$. The notions of sub-formula, free and bound variable are defined for this logic in the natural way. Note that any sub-formula of a sentence in $L_{\infty\omega}$ has only finitely many free variables (see [Barwise, 1973]).

$L_{\infty\omega}$ is complete in expressive power in the following sense. Consider any class of finite structures \mathcal{C} such that \mathcal{C} is closed under isomorphism. Since any finite structure \mathfrak{A} is completely characterized up to isomorphism by a first-order sentence, $\phi_{\mathfrak{A}}$, \mathcal{C} is expressed by

the $L_{\infty\omega}$ sentence $\bigvee\{\phi_{\mathfrak{A}} \mid \mathfrak{A} \in \mathcal{C}\}$. Clearly, this language is too strong on finite structures. One restriction of this language, $L_{\infty\omega}^{\omega}$, was introduced by Barwise [Barwise, 1977].

Definition 1 $L_{\infty\omega}^k$ is the collection of formulas of $L_{\infty\omega}$ that have at most k distinct variables (free or bound). $L_{\infty\omega}^{\omega}$ is the collection of formulas of $L_{\infty\omega}$ that have a finite number of distinct variables.

$$L_{\infty\omega}^{\omega} = \bigcup_{k=1}^{\infty} L_{\infty\omega}^k.$$

In what follows, we will assume that any formula in $L_{\infty\omega}^k$ is written so as to use only the variables x_1, \dots, x_k .

The language $L_{\infty\omega}^{\omega}$ is restricted in its expressive power when compared with $L_{\infty\omega}$, yet it is still powerful enough to express properties that are not recursive (see, for instance, [Kolaitis and Vardi, 1992b]). To show that the restriction is real, we need to exhibit some property that cannot be expressed in the former language. To this end, we now present a version of the Ehrenfeucht-Fraïssé games.

The games of Ehrenfeucht and Fraïssé [Ehrenfeucht, 1961; Fraïssé, 1954] are an algebraic means to characterize elementary equivalence of structures. They were generalized by Karp [Karp, 1965] to characterize $L_{\infty\omega}$ -equivalence up to a given quantifier rank, by means of a sequence of sets of partial isomorphisms with the back and forth property. Barwise [Barwise, 1977] defined a notion of a set of partial isomorphisms with the back and forth property through k to characterize $L_{\infty\omega}^k$ -equivalence of structures. Immerman [Immerman, 1982] defined a k -pebble game that captures this equivalence in the case of finite structures, and essentially the same game was described in [Poizat, 1982]. We combine these results in the single result below, which we state and prove in its full generality. We will then consider special cases that are of interest here. We begin with some notation. For a function f , $dom(f)$ denotes its domain, $rng(f)$ its range and $|f|$ its cardinality. For a formula ϕ , $qr(\phi)$ denotes its quantifier rank, defined as:

Definition 2 The quantifier rank of a formula ϕ , written $qr(\phi)$, is defined inductively as follows:

1. if ϕ is atomic then $qr(\phi) = 0$,
2. if $\phi = \neg\psi$ then $qr(\phi) = qr(\psi)$,

3. if $\phi = \bigvee \Phi$ or $\phi = \bigwedge \Phi$ then $qr(\phi) = \sup\{qr(\psi) \mid \psi \in \Phi\}$, and
4. if $\phi = \exists x\psi$ or $\phi = \forall x\psi$ then $qr(\phi) = qr(\psi) + 1$.

In general, the quantifier rank of an infinitary formula is an (infinite) ordinal. The operations sup and addition are to be understood as ordinal operations.

Definition 3 A function f is a partial isomorphism from \mathfrak{A} to \mathfrak{B} , if and only if, the domain of f is a subset of \mathfrak{A} that includes the interpretations of all constants in the language of \mathfrak{A} and f is an isomorphic map over this domain, i.e. $f(c^{\mathfrak{A}}) = c^{\mathfrak{B}}$ for all constants c and for all relation symbols R and a_1, \dots, a_m in the domain of f , $\mathfrak{A} \models R^{\mathfrak{A}}(a_1, \dots, a_m)$ if and only if $\mathfrak{B} \models R^{\mathfrak{B}}(f(a_1), \dots, f(a_m))$.

Definition 4 For any two structures $\mathfrak{A} = \langle A, \dots \rangle$ and $\mathfrak{B} = \langle B, \dots \rangle$ and any ordinal α , a collection of sets of partial isomorphisms $\{I_\beta \mid \beta < \alpha\}$ is said to have the k back and forth property if and only if:

1. Each I_β is non-empty,
2. $I_\beta \supseteq I_{\beta+1}$ (for all $\beta + 1 < \alpha$),
3. If $f \in I_\beta$ ($0 \leq \beta < \alpha$) and $g \subseteq f$ is a partial isomorphism, then $g \in I_\beta$, and
4. For every $f \in I_{\beta+1}$ ($\beta + 1 < \alpha$) such that $|f| < k$ and every $a \in A$ (resp. $b \in B$), there is a $g \in I_\beta$ with $f \subseteq g$ and $a \in \text{dom}(g)$ (resp. $b \in \text{rng}(g)$).

We shall see later, that in the case where the ordinal α is finite, this corresponds to a k pebble Ehrenfeucht-style pebble game. In that context, the collection of sets of partial isomorphisms with the k back and forth property can be viewed as a winning strategy for Player II in the game.

Theorem 5 For any two structures, $\mathfrak{A} = \langle A, \dots \rangle$, $\mathfrak{B} = \langle B, \dots \rangle$ in a purely relational language, the following statements are equivalent:

1. For all sentences $\phi \in L_{\infty\omega}^k$ with $qr(\phi) < \alpha$,

$$\mathfrak{A} \models \phi \text{ iff } \mathfrak{B} \models \phi$$

2. There is a set $\{I_\beta \mid \beta < \alpha\}$ of non-empty sets of partial isomorphisms from \mathfrak{A} to \mathfrak{B} with the k back and forth property.

Proof:

(2 \Rightarrow 1) We show by induction on β that for formulas $\phi(y_1 \dots y_m) \in L_{\infty\omega}^k$, with $qr(\phi) \leq \beta$, if $f \in I_\beta$ and $a_1, \dots, a_m \in \text{dom}(f)$, then $\mathfrak{A} \models \phi[a_1 \dots a_m]$ iff $\mathfrak{B} \models \phi[f(a_1) \dots f(a_m)]$.

Basis:

If $qr(\phi) = 0$ then ϕ is a Boolean combination of atomic formulas and since f is a partial isomorphism, the result follows.

Induction Step:

We now proceed by induction on the structure of the formula ϕ . The cases $\phi = \neg\psi$ and $\phi = \bigwedge_{j \in J} \psi_j$ are trivial. So, we only need to consider the case where $\phi = \exists y_0 \psi(y_0 \dots y_m)$. We assume that y_0 is distinct from all the variables y_1, \dots, y_m , since if $y_0 = y_i$ for some i , then y_i has no free occurrences in ϕ and the value of $f(a_i)$ does not affect our conclusion, provided that $f(a_i) = f(a_j)$, if and only if, $a_i = a_j$. But, this follows simply from the fact that f is a partial isomorphism. Note that $qr(\phi) = \delta + 1$ where $qr(\psi) = \delta$.

Suppose $\mathfrak{A} \models \phi[a_1 \dots a_m]$ for some $a_1, \dots, a_m \in A$ and that $a_1, \dots, a_m \in \text{dom}(f)$ for some $f \in I_{\delta+1}$. Then, there is an $a_0 \in A$ such that $\mathfrak{A} \models \psi[a_0 a_1 \dots a_m]$. By clause 3 of Definition 4, there is an $f' \in I_{\delta+1}$ with $\text{dom}(f') = \{a_1, \dots, a_m\}$ and $f' \subseteq f$. Since $|f'| < k$, by clause 4 there is a $g \in I_\delta$ extending f' such that $a_0 \in \text{dom}(g)$. But then, by the induction hypothesis, $\mathfrak{B} \models \psi[g(a_0)g(a_1) \dots g(a_m)]$, i.e. $\mathfrak{B} \models \phi[g(a_1) \dots g(a_m)]$ and therefore $\mathfrak{B} \models \phi[f(a_1) \dots f(a_m)]$, since f and g agree on a_1, \dots, a_m .

Similarly, if $\mathfrak{B} \models \phi[b_1 \dots b_m]$ and $b_1, \dots, b_m \in \text{rng}(f)$, then $\mathfrak{A} \models \phi[f^{-1}(b_1) \dots f^{-1}(b_m)]$.

(1 \Rightarrow 2) Define the I_β as follows: $f \in I_\beta$ if and only if f is a partial isomorphism from \mathfrak{A} to \mathfrak{B} , with $|f| \leq k$ and for all formulas $\phi \in L_{\infty\omega}^k$ with $qr(\phi) \leq \beta$ and all $a_0 \dots a_m \in \text{dom}(f)$, $\mathfrak{A} \models \phi[a_0 \dots a_m]$ iff $\mathfrak{B} \models \phi[f(a_0) \dots f(a_m)]$.

By definition, $I_\delta \supseteq I_\beta$ for $\delta \leq \beta$. Also, since \mathfrak{A} and \mathfrak{B} agree on all sentences of quantifier rank less than α , the empty partial isomorphism is in I_β for all β , and

therefore all the I_β are non-empty. It is also clear that if $f \in I_\beta$ and $g \subseteq f$ then $g \in I_\beta$. Thus, we only need to show that clause 4 holds.

For contradiction, suppose that there is an $f \in I_{\beta+1}$ with $|f| < k$ and an $a \in A$ such that for all $g \in I_\beta$ with $g \supseteq f$, $a \notin \text{dom}(g)$. Then, for every $b \in B$, there must be a formula $\psi_b[y_0 y_1 \dots y_m]$ with $qr(\psi_b) \leq \beta$ such that $\mathfrak{A} \models \psi_b[aa_1 \dots a_m]$ and $\mathfrak{B} \models \neg \psi_b[bf(a_1) \dots f(a_m)]$ (where $a_1, \dots, a_m \in \text{dom}(f)$). Let $\phi = \exists y \bigwedge_{b \in B} \psi_b[yy_1 \dots y_m]$. But then, $qr(\phi) \leq \beta + 1$, $\mathfrak{A} \models \phi[a_1 \dots a_m]$ and $\mathfrak{B} \models \neg \phi[f(a_1) \dots f(a_m)]$ contradicting the assumption that $f \in I_{\beta+1}$. \blacksquare

Note that any formula of $L_{\infty\omega}^\omega$ of a given finite quantifier rank is equivalent to a finitary formula of the same rank. This is because, in a finite relational language and over a fixed finite collection of free variables, there are only finitely many inequivalent first-order formulas of a given quantifier rank. It follows that two structures \mathfrak{A} and \mathfrak{B} agree on all sentences of $L_{\infty\omega}^\omega$ of quantifier rank less than n if they agree on all first-order sentences of quantifier rank less than n . Moreover, if the two structures are finite, then any chain of sets of partial isomorphisms as above of length ω can be extended to any ordinal length. To see this, note that there are only finitely many maps from subsets of A into B . Thus, one of the sets in the chain must be repeated, and hence, can be repeated indefinitely. Writing L^k for the fragment of first-order logic with at most k variables, we have the following corollary:

Corollary 1 ([Kolaitis and Vardi, 1992b]) *For finite structures \mathfrak{A} and \mathfrak{B} , the following are equivalent:*

- For every sentence $\phi \in L_{\infty\omega}^k$, $\mathfrak{A} \models \phi$ iff $\mathfrak{B} \models \phi$
- For every sentence $\phi \in L^k$, $\mathfrak{A} \models \phi$ iff $\mathfrak{B} \models \phi$

We write $\mathfrak{A} \equiv_k \mathfrak{B}$ to denote that \mathfrak{A} and \mathfrak{B} satisfy the same sentences of L^k , and call them L^k -equivalent.

When the sequence of sets of partial isomorphisms is finite, we can view it in terms of the following two-player pebble game. We have a board consisting of one copy of each of the structures \mathfrak{A} and \mathfrak{B} . There is also a supply of pairs of pebbles $\{\langle a_1, b_1 \rangle, \dots, \langle a_k, b_k \rangle\}$. At each move of the game, Player I picks up one of the pebbles (either an unused pebble, or one that is already on the board) and places it on an element of the corresponding structure

(i.e. she places a_i on an element of A or b_i on an element of B). Player II then responds by placing the unused pebble in the pair on an element of the other structure. Player II loses if the resulting map, f , from \mathfrak{A} to \mathfrak{B} , given by $f(a_j) = b_j, 1 \leq j \leq k$, is not a partial isomorphism. Player II wins the n -move game if she has a strategy to avoid losing in the first n moves, regardless of what moves are made by Player I. We then have the following characterization:

Corollary 2 ([Immerman, 1982; Poizat, 1982]) *Let \mathfrak{A} and \mathfrak{B} be structures over a fixed signature and let \bar{a} and \bar{b} be l -tuples (for $l \leq k$) of elements from the respective structures. Player II has a winning strategy for n moves of the k pebble game on structures \mathfrak{A} and \mathfrak{B} with the pebbles initially on the tuples \bar{a} and \bar{b} , if and only if for every formula $\phi(\bar{x})$ of quantifier rank up to n with at most k distinct variables, $\mathfrak{A} \models \phi[\bar{a}]$ just in case $\mathfrak{B} \models \phi[\bar{b}]$.*

On finite structure, or on a fixed infinite structure, the languages $\text{FO} + \text{LFP}$ and $\text{FO} + \text{PFP}$ (see Section 2) can be viewed as fragments of $L_{\infty\omega}^\omega$. Consider any formula $\phi \equiv \mathbf{lfp}(S, x_1, \dots, x_n)\psi(S)$, where ψ is a first-order formula positive in S having at most k variables. We can define the m^{th} iterative stage of ϕ by a first-order formula ψ^m with at most $k + n$ variables (see [Kolaitis and Vardi, 1992b]). Then, ϕ is equivalent to the formula $\bigvee_{m=0}^{\infty} \psi^m$. Similarly, $\mathbf{pfp}(S, x_1, \dots, x_n)\psi(S)$ is equivalent to $\bigvee_{m=0}^{\infty} (\psi^m(x_1 \dots x_n) \wedge \forall x_1 \dots \forall x_n (\psi^m(x_1 \dots x_n) \leftrightarrow \psi^{m+1}(x_1 \dots x_n)))$. Writing $(\text{FO} + \text{LFP})^k$ (resp. $(\text{FO} + \text{PFP})^k$) for the fragment of $\text{FO} + \text{LFP}$ (resp. $\text{FO} + \text{PFP}$) with only k variables, we can use a semantic argument to establish the following corollary. This refinement of the containment of $\text{FO} + \text{LFP}$ and $\text{FO} + \text{PFP}$ in $L_{\infty\omega}^\omega$ also appears in [Kolaitis and Vardi, 1992a]¹.

Corollary 3 *On finite structures, $(\text{FO} + \text{LFP})^k \subseteq (\text{FO} + \text{PFP})^k \subseteq L_{\infty\omega}^k$.*

Proof:

Let $\psi(S, \bar{x})$ be an S -positive formula with at most k variables, and let \mathfrak{A} be any structure. We show by induction on m that the query defined by ψ^m is closed under the equivalence relation \equiv_k on tuples. The construction in the next section shows that every query that is closed under \equiv_k is definable in $L_{\infty\omega}^k$ on finite structures.

¹The proof of this refinement is based on a syntactic argument and shows that each ψ^m is equivalent to a formula of L^k .

Basis:

ψ^0 defines the empty relation.

Induction Step:

Let \mathfrak{A} and \mathfrak{B} be structures and \bar{a} and \bar{b} be tuples of elements from these structures such that $\langle \mathfrak{A}, \bar{a} \rangle \equiv_k \langle \mathfrak{B}, \bar{b} \rangle$. The relation defined by ψ^{m+1} on \mathfrak{A} (\mathfrak{B}) is the relation defined by $\psi(S)$ on the structure $\langle \mathfrak{A}, S^{\mathfrak{A}} \rangle$ ($\langle \mathfrak{B}, S^{\mathfrak{B}} \rangle$) where $S^{\mathfrak{A}}$ ($S^{\mathfrak{B}}$) is the relation defined by ψ^m . By Corollary 2, Player II has a winning strategy in the k -pebble game played on the structures $\langle \mathfrak{A}, \bar{a} \rangle$ and $\langle \mathfrak{B}, \bar{b} \rangle$. But, this winning strategy must also be a winning strategy for Player II in the k -pebble game played on the structures $\langle \mathfrak{A}, S^{\mathfrak{A}}, \bar{a} \rangle$ and $\langle \mathfrak{B}, S^{\mathfrak{B}}, \bar{b} \rangle$. If this were not the case, then after some number of moves, the pebbles are on tuples, \bar{a}' and \bar{b}' that differ on an atomic formula that involves S . This contradicts the inductive hypothesis that the query ψ^m is closed under \equiv_k . Thus $\langle \mathfrak{A}, S^{\mathfrak{A}}, \bar{a} \rangle \equiv_k \langle \mathfrak{B}, S^{\mathfrak{B}}, \bar{b} \rangle$ and ψ^{m+1} must be closed under \equiv_k . ■

The result that FO + LFP is a fragment of $L_{\infty\omega}^\omega$ over fixed infinite structures is due to Rubin [Rubin, 1975] and appeared in [Barwise, 1977]. Kolaitis and Vardi [Kolaitis and Vardi, 1992b] showed that FO + LFP and FO + PFP are contained in $L_{\infty\omega}^\omega$ over the class of finite structures.

4 Characterizing Structures up to L^k -equivalence

It is clear that for every finite structure \mathfrak{A} , we can write a first-order sentence $\phi_{\mathfrak{A}}$ such that any structure that satisfies $\phi_{\mathfrak{A}}$ is isomorphic to \mathfrak{A} . A simple application of Theorem 5 shows that not all such sentences are in L^k for any given k . This raises the question of whether there is a sentence $\phi_{\mathfrak{A}}^k$ of L^k associated with \mathfrak{A} such that any structure satisfying it is L^k -equivalent to \mathfrak{A} . In this section, we answer this question in the affirmative. This result was implicit in [Poizat, 1982].

The proof is adapted from the proof of Scott's theorem [Scott, 1965] as presented in [Barwise, 1973]. This theorem is as follows:

Theorem 6 ([Scott, 1965]) *For every countable structure \mathfrak{A} , there is a sentence, $\phi \in L_{\omega_1\omega}$, such that for any countable structure \mathfrak{B} , $\mathfrak{B} \models \phi$ if and only if $\mathfrak{A} \cong \mathfrak{B}$.*

A similar construction also appeared in [Läuchli, 1968] and was used in [Shelah, 1975] and

[Gurevich, 1979].

For the purpose of this section, we will assume that there are no constants in the language being considered. The results can be easily generalized to the case where constants are present.

Let A be the universe of \mathfrak{A} and let $S = A^{\leq k}$ be the set of sequences of elements of A of length less than or equal to k . For $s \in S$ and $a \in A$, let $s \cdot \langle a \rangle$ denote the sequence obtained by extending s by the single element a .

We define a formula ϕ_s^m for each $s \in S$ and each $m \in \mathbf{N}$. The formula has free variables x_1, \dots, x_l , where l is the length of s . We want it to be the case that $\mathfrak{A} \models \phi_s^m[s]$ and that this formula characterizes s completely up to equivalence on formulas with k variables and quantifier rank m . In order to define these formulas, we first define formulas ϕ_σ^m for each assignment map $\sigma : X \rightarrow A$, where $X \subseteq \{x_1 \dots x_k\}$.

For all $\sigma : X \rightarrow A$,

ϕ_σ^0 is the conjunction of all atomic and negated atomic formulas θ with free variables among X such that $\mathfrak{A} \models \theta[\sigma]$;

if $\text{card}(\sigma) < k$ then,

$$\phi_\sigma^{m+1} = \phi_\sigma^m \wedge \tag{1}$$

$$\bigwedge_{a \in A} \exists x_i \phi_{\sigma \cup \{(x_i, a)\}}^m \wedge \tag{2}$$

$$\forall x_i \bigvee_{a \in A} \phi_{\sigma \cup \{(x_i, a)\}}^m \tag{3}$$

for the least i such that $x_i \notin X$;

if $\text{card}(\sigma) = k$ then,

$$\phi_\sigma^{m+1} = \phi_\sigma^m \wedge \bigwedge_{i=1 \dots k} \phi_{\sigma_i}^{m+1} \tag{4}$$

where $\sigma_i = \sigma \setminus \{(x_i, \sigma(x_i))\}$.

For any sequence $s = \langle a_1 \dots a_l \rangle$, let ϕ_s^m denote ϕ_σ^m , for $\sigma = \{(x_i, a_i) \mid i = 1, \dots, l\}$.

Example 1 We illustrate the construction of the formulas ϕ_s^m on the simple graph $\langle V, E \rangle$, with $V = \{0, 1, 2\}$ and $E = \{(0, 1), (1, 2)\}$. Taking $k = 2$,

$$\begin{aligned} \phi_{\langle 0 \rangle}^1 &= \phi_{\langle 0 \rangle}^0 \wedge \bigwedge_{a \in \{0, 1, 2\}} \exists x_2 \phi_{\langle 0, a \rangle}^0 \\ &\quad \wedge \forall x_2 \bigvee_{a \in \{0, 1, 2\}} \phi_{\langle 0, a \rangle}^0; \end{aligned}$$

where

$$\begin{aligned} \phi_{\langle 0 \rangle}^0 &= x_1 = x_1 \wedge \neg E x_1 x_1; \\ \phi_{\langle 0, 0 \rangle}^0 &= x_1 = x_2 \wedge \neg E x_1 x_2 \wedge \neg E x_2 x_1; \\ \phi_{\langle 0, 1 \rangle}^0 &= x_1 \neq x_2 \wedge E x_1 x_2 \wedge \neg E x_2 x_1; \\ \phi_{\langle 0, 2 \rangle}^0 &= x_1 \neq x_2 \wedge \neg E x_1 x_2 \wedge \neg E x_2 x_1. \end{aligned}$$

Observe that $\phi_{\langle 0 \rangle}^1$ identifies the element 0 uniquely.

We state Lemma 1 and Theorem 7 only for the case of finite structures, since that is the case that is of interest here. However, similar results can be derived for the case where the structures may be infinite. In the latter case, the conjunction in (2) and the disjunction in (3) could be infinitary. Thus, the formulas constructed are no longer first-order, but they are in $L_{\infty\omega}^k$.

Lemma 1 Let $s = \langle a_1 \dots a_l \rangle \in S$ be a sequence of elements from the finite structure \mathfrak{A} with $l \leq k$. For any structure $\mathfrak{B} = \langle B, \dots \rangle$ and $b_1, \dots, b_l \in B$, $\mathfrak{B} \models \phi_s^m[b_1 \dots b_l]$ if and only if there is a sequence of sets of partial isomorphisms $I_0 \supseteq \dots \supseteq I_m$ with the k back and forth property and $f = \{\langle a_1, b_1 \rangle \dots \langle a_l, b_l \rangle\} \in I_m$.

Proof:

\Leftarrow This follows immediately from the proof of Theorem 5 since the existence of such a sequence implies that for any ϕ of quantifier rank m , $\mathfrak{B} \models \phi[b_1 \dots b_l]$ if $\mathfrak{A} \models \phi[a_1 \dots a_l]$.

Clearly, $qr(\phi_s^m) = m$ and $\mathfrak{A} \models \phi_s^m[s]$.

\Rightarrow The proof is by induction on m .

Basis Let $I_0 = \{g \mid g \subseteq f\}$. f is a partial isomorphism, because $\mathfrak{B} \models \phi_s^0[f(s)]$. Even if $s = \langle \rangle$ and f is the empty map, I_0 is non-empty.

Induction Step There are two cases to be considered:

Case: $l < k$

Let $I_{m+1} = \{g \mid g \subseteq f\}$.

By induction hypothesis and (1), there is a sequence $I_0^s \dots I_m^s$ with the k back and forth property and $f \in I_m^s$.

Furthermore, by (2) and the induction hypothesis, for every $a \in A$, there is a $b \in B$ and a sequence $I_0^{s \cdot \langle a \rangle} \dots I_m^{s \cdot \langle a \rangle}$ with the k back and forth property such that $\{\langle a_1, b_1 \rangle \dots \langle a_l, b_l \rangle, \langle a, b \rangle\} \in I_m^{s \cdot \langle a \rangle}$.

Let $I_j = I_j^s \cup \bigcup_{a \in A} I_j^{s \cdot \langle a \rangle}$ (for $0 \leq j \leq m$). Note that, in general, the k back and forth property is preserved under this kind of element-wise union. Thus, we need to verify that $I_{m+1} \subseteq I_m$ and that every element of I_{m+1} is extensible in I_m to arbitrary elements of A and B . The former follows from the fact that $I_{m+1} \subseteq I_m^s$ and the latter follows from (2) and (3) respectively.

Case: $l = k$

By the argument for the case above, there are sequences $I_0^i \dots I_{m+1}^i$ corresponding to each of the partial isomorphisms, f_i , obtained by dropping the pair $\langle a_i, b_i \rangle$ from f . Since $\mathfrak{B} \models \phi_s^m[b_1 \dots b_k]$, f itself is a partial isomorphism.

Let $I_j = \{f\} \cup \bigcup_{i=1 \dots k} I_j^i$ for $0 \leq j \leq m+1$. Each of the I_j is still closed under restrictions, because if $g \subseteq f$, then either $g = f$ or $g \subseteq f_i$ for some i . Since $|f| = k$, extensibility of f is not required, and we are done. \blacksquare

Intuitively, we can also see why Lemma 1 holds in terms of pebble games. If $\mathfrak{B} \models \phi_a^m[\bar{b}]$, then Player II has a winning strategy for m moves of the k -pebble game starting from $\langle \mathfrak{A}, \bar{a} \rangle$ and $\langle \mathfrak{B}, \bar{b} \rangle$. This is clear for the case when $m = 0$ since a partial isomorphism between $\langle \mathfrak{A}, \bar{a} \rangle$ and $\langle \mathfrak{B}, \bar{b} \rangle$ exists just in case \bar{a} and \bar{b} satisfy the same atomic formulas. If $\mathfrak{B} \models \phi_a^{m+1}[\bar{b}]$ then any one move from the position \bar{a}, \bar{b} leads to a position from which Player II has a winning strategy for m more moves. Consider the two cases. If $\text{length}(\bar{b}) < k$, there are fewer than k pairs of pebbles on the board. In this case, (1) guarantees that Player II has a winning strategy for m moves from the initial position, (2) guarantees that if Player I places a new pebble on structure \mathfrak{A} , Player II can respond, resulting in a position from which she has a winning strategy for m moves and (3) guarantees that if Player I places

a new pebble on the structure \mathfrak{B} , Player II can respond similarly. If $\text{length}(\bar{b}) = k$, all the pebbles are on the board, and Player I must pick up a pebble in order to move, and therefore the game is equivalent to an $m + 1$ move game with one of the pairs of pebbles removed. This is expressed in (4).

For a given sequence s of length l , let $X_s^m = \{s' \in S \mid \mathfrak{A} \models \phi_s^m[s']\}$. Each X_s^m is a set of l -tuples of A and $X_s^m \supseteq X_s^{m+1}$. Since A is finite, there must be an m_s such that $X_s^{m_s} = X_s^m$ for all $m > m_s$. Let $m^* = \max\{m_s \mid s \in S\}$. Now, define the *Scott sentence* ϕ as follows:

$$\phi \equiv \phi_{\langle \rangle}^{m^*} \wedge \bigwedge_{s \in S} \forall x_1 \dots \forall x_k (\phi_s^{m^*} \rightarrow \phi_s^{m^*+1})$$

Note that $\phi \in L^k$ and that $\mathfrak{A} \models \phi$. We now show that this sentence characterizes the structure \mathfrak{A} up to L^k -equivalence.

Theorem 7 *For every finite structure \mathfrak{A} and any k , there is a sentence, $\phi \in L^k$ such that for any structure \mathfrak{B} , $\mathfrak{B} \models \phi$ if and only if $\mathfrak{A} \equiv_k \mathfrak{B}$.*

Proof:

Let ϕ be as defined above. We only need to show that if $\mathfrak{B} \models \phi$, then $\mathfrak{A} \equiv_k \mathfrak{B}$. Let F be the set of maps $\{\langle a_1, b_1 \rangle, \dots, \langle a_l, b_l \rangle\}$ such that $\mathfrak{B} \models \phi_{\langle a_1, \dots, a_l \rangle}^{m^*+1}[b_1 \dots b_l]$. The set F is non-empty since $\mathfrak{B} \models \phi_{\langle \rangle}^{m^*+1}$. By Lemma 1, for each $f \in F$, there is a sequence $I_0^f \dots I_{m^*+1}^f$ with the k back and forth property. Let $I_i = \bigcup_{f \in F} I_i^f$ and let $I_m = I_{m^*+1}$ for all $m > m^* + 1$. We claim the infinite sequence $I_0 \supseteq \dots \supseteq I_m \dots$ has the k back and forth property. We will establish the extensibility of every element of I_{m^*+2} . The rest then follows.

Consider any $f \in I_{m^*+2}$ with $|f| < k$ and any $a \in A$. By definition, $f \in I_{m^*+1}$. Since we know that the sequence through I_{m^*+1} has the k back and forth property, there is a $g \in I_{m^*}$ such that $f \subseteq g$ and $a \in \text{dom}(g)$. Let $\text{dom}(g) = \{a_1, \dots, a_l\}$. Then, by the other direction of Lemma 1, $\mathfrak{B} \models \phi_{\langle a_1, \dots, a_l \rangle}^{m^*}[\langle g(a_1) \dots g(a_l) \rangle]$ and therefore, by the implication in ϕ , $\mathfrak{B} \models \phi_{\langle a_1, \dots, a_l \rangle}^{m^*+1}[\langle g(a_1) \dots g(a_l) \rangle]$ and so $g \in F$. Since $g \in I_{m^*+1}^g$, it follows that $g \in I_{m^*+1}$ and we are done. \blacksquare

There are some points about the above construction that are noteworthy. First of all, we could have, alternatively, defined m^* as the smallest m such that $X_s^m = X_s^{m+1}$ for all s . To see this, just observe that this is the only property of m^* used in the above proof. It is easily seen that if $X_s^m = X_s^{m+1}$ for all sequences of length k , then $X_s^m = X_s^{m+1}$ for all

s. Since the relation $s' \in X_s^m$ is an equivalence relation on tuples, the sets X_s^m determine a partition of the sequences of a given length. Given that there are n^k sequences of length k in a structure of size n , we can derive the bound $m^* \leq n^k$. This gives us the following:

Corollary 4 *If \mathfrak{A} is a structure of size n and \mathfrak{B} a structure such that \mathfrak{A} and \mathfrak{B} agree on all sentences of L^k of quantifier rank up to $n^k + k + 1$, then $\mathfrak{A} \equiv_k \mathfrak{B}$.*

It follows immediately from Theorem 7 that any class, K , of finite structures that is closed under L^k -equivalence (that is, if $\mathfrak{A} \in K$ and $\mathfrak{A} \equiv_k \mathfrak{B}$ then $\mathfrak{B} \in K$) is definable in $L_{\infty\omega}^k$. This result is due to Kolaitis and Vardi [Kolaitis and Vardi, 1992b].

The following normal form result is also immediate. It improves normal forms in [Kolaitis and Vardi, 1992b] (Corollary 2.20) and [Kolaitis and Vardi, 1992a]. The normal form result in [Kolaitis and Vardi, 1992a] was derived from a weaker form of Theorem 7, where the sentence ϕ , characterizing a structure up to equivalence in \equiv_k was shown to be in $(\text{FO} + \text{LFP})^l$, for some $l \geq k$.

Corollary 5 *If K is definable in $L_{\infty\omega}^k$, it is defined by a single (countable) disjunction of sentences of L^k .*

Proof:

Suppose K is definable in $L_{\infty\omega}^k$. By Corollary 1, K is closed under L^k -equivalence. If we write $\phi_{\mathfrak{A}}$ for the sentence of L^k that characterizes a structure \mathfrak{A} up to L^k -equivalence, it follows that K is defined by the sentence $\bigvee\{\phi_{\mathfrak{A}} \mid \mathfrak{A} \in K\}$. ■

Since the complement class of a class of structures that is closed under L^k -equivalence is itself closed under L^k -equivalence we also have the following dual normal form:

Corollary 6 *If K is definable in $L_{\infty\omega}^k$, it is defined by a single (countable) conjunction of sentences of L^k .*

Finally, it is not only structures that are characterized up to L^k -equivalence in the above proof, but also sequences of elements.

Definition 5 *For any sequence $s = \langle a_1 \dots a_l \rangle$ of elements in a structure \mathfrak{A} , with $l \leq k$, define the L^k -type of s , denoted $\text{Type}_k(s)$, to be the set of formulas, $\phi \in L^k$ with free variables among x_1, \dots, x_l , such that $\mathfrak{A} \models \phi[a_1 \dots a_l]$.*

Then, using the formulas $\phi_s^{m^*}$, we get the following from the proof of Theorem 7:

Corollary 7 *For every finite structure \mathfrak{A} , for every $l \leq k$ and sequence a_1, \dots, a_l of elements from \mathfrak{A} , there is a formula, $\phi \in \text{Type}_k(\langle a_1, \dots, a_l \rangle)$ such that for any structure \mathfrak{B} , and elements $b_1 \dots b_l \in B$, if $\mathfrak{B} \models \phi[b_1 \dots b_l]$, then $\text{Type}_k(\langle a_1, \dots, a_l \rangle) = \text{Type}_k(\langle b_1, \dots, b_l \rangle)$.*

5 Inductively Ordering the Types

Having seen how, for a particular L^k type, we can write a formula that characterizes it, we now turn to writing a formula that will define a total ordering of these types. We will show that this can be done uniformly in FO + LFP, *i.e.* a single formula will define the ordering on all structures. From this result we will derive Abiteboul and Vianu's result [Abiteboul and Vianu, 1991b] that PSPACE = P if and only if FO + LFP = FO + PFP. The technique used for defining the ordering is inspired by a color-refinement algorithm in [Immerman and Lander, 1990].

Looking again at the definitions of the ϕ_s^m in the last section, we can see that these formulas were defined by a simultaneous induction on first-order formulas – simultaneous in all the sequences s . This, along with the observation that the basis of this induction is finite, in the sense that there are, up to equivalence, only finitely many quantifier free formulas, using k variables, in a finite relational language, suggests that we could accomplish the entire process with a single formula of FO + LFP. We formalize this intuition below.

We first construct a formula of FO + LFP which defines, on any structure \mathfrak{A} , an equivalence relation on k -tuples of elements such that two tuples are equivalent if and only if they have the same L^k -type.

Definition 6 *For any structure \mathfrak{A} and elements $a_1 \dots a_l \in |\mathfrak{A}|$, the basic L^k -type of $a_1 \dots a_l$ is the set of atomic formulas, ϕ , of L^k in l free variables such that $\mathfrak{A} \models \phi[a_1 \dots a_l]$.*

Note that for a given finite signature, σ , there are only finitely many distinct basic types. Furthermore, each basic type is characterized by a single quantifier free formula of L^k .

Let $\alpha_1(x_1 \dots x_k), \dots, \alpha_q(x_1 \dots x_k)$ be a fixed enumeration of these quantifier free formulas of L^k in k free variables characterizing all the basic types in some signature σ . Then,

define ϕ_0 as follows:

$$\phi_0(x_1 \dots x_k y_1 \dots y_k) \equiv \bigvee_{1 \leq i \neq j \leq q} (\alpha_i(\bar{x}) \wedge \alpha_j(\bar{y}))$$

where $\alpha_i(\bar{y})$ is obtained from $\alpha_i(\bar{x})$ by replacing every x_j by y_j . It should be clear that for any tuples $\bar{a}, \bar{b} \in |\mathfrak{A}|^k$, $\mathfrak{A} \models \phi_0[\bar{a}\bar{b}]$ if and only if the basic types of \bar{a} and \bar{b} are different.

Now we define formulas ϕ and ψ as follows. The least fixed point of ϕ expresses the *inequivalence* of L^k -types, and therefore, ψ expresses the equivalence of types.

$$\begin{aligned} \phi(R, x_1 \dots x_k y_1 \dots y_k) \equiv & \phi_0(\bar{x}\bar{y}) \vee \bigvee_{1 \leq i \leq k} \exists x_i \forall y_i R(x_1 \dots x_k y_1 \dots y_k) \\ & \vee \bigvee_{1 \leq i \leq k} \exists y_i \forall x_i R(x_1 \dots x_k y_1 \dots y_k); \end{aligned}$$

$$\psi(z_1 \dots z_{2k}) \equiv \neg \mathbf{lfp}(R, \bar{x}, \bar{y}) \phi(z_1 \dots z_{2k}).$$

Claim 1 For any structure \mathfrak{A} on signature σ , $\mathfrak{A} \models \psi[a_1 \dots a_k a'_1 \dots a'_k]$ if and only if \bar{a} and \bar{a}' have the same L^k -type.

Proof:

To establish this claim, we need to show that $\mathbf{lfp}(R, \bar{x}, \bar{y}) \phi[\bar{a}\bar{a}']$ expresses the *inequivalence* of the two k -tuples. Picture the k -pebble game being played on two isomorphic copies of \mathfrak{A} , and at some stage the pebbles are placed on $\langle a_1, a'_1 \rangle, \dots, \langle a_k, a'_k \rangle$. By Corollary 2, if the two tuples have the same L^k -type, then Player II can play indefinitely from this point on without losing. We claim that if $\phi^r[\bar{a}\bar{a}']$ (the r^{th} iterative stage of ϕ), then Player I can win in r moves or less. Clearly, if $r = 0$, by the definition of ϕ_0 , \bar{a} and \bar{a}' differ on a quantifier free formula and hence the map from one to the other is not a partial isomorphism. If $r = m + 1$, then the definition of ϕ tells us that we can, in one move, get to a configuration that is in ϕ^m . Conversely, it is easily seen from the definition of ϕ that if \bar{a} and \bar{a}' are such that Player I wins in r moves or less, then $\phi^r[\bar{a}\bar{a}']$. \blacksquare

We will henceforth use the symbol \sim_k in infix notation to denote the relation defined by ψ . Kolaitis and Vardi [Kolaitis and Vardi, 1992a] have independently shown that the relation \sim_k is definable in $\text{FO} + \text{LFP}$. The following theorem will give an inductive definition of an ordering relation on the equivalence classes defined by this relation. In order to state the theorem, we require the following definition.

Definition 7 *The relation P is a strict linear pre-order on D , if and only if, for every $x, y, z \in D$,*

1. $P(x, y) \rightarrow \neg P(y, x)$;
2. $(P(x, y) \wedge P(y, z)) \rightarrow P(x, z)$;
3. $(\neg P(x, y) \wedge \neg P(y, x)) \rightarrow ((P(x, z) \leftrightarrow P(y, z)) \wedge (P(z, x) \leftrightarrow P(z, y)))$

If P is a strict linear pre-order on D the relation $E = \{\langle a, b \rangle \mid a, b \in D \wedge \neg P(a, b) \wedge \neg P(b, a)\}$ is the equivalence relation induced by P on D .

Theorem 8 *For every $k \in \omega$, there is a $2k$ -ary global relation P_k on the class of finite structures such that*

1. P_k is definable in $\text{FO} + \text{LFP}$;
2. for all finite structures \mathfrak{A} , $P_k(\mathfrak{A})$ is a strict linear pre-order of $|\mathfrak{A}|^k$;
3. the global equivalence relation induced by P_k is \sim_k .

Proof:

Fix $k \in \omega$. The relation P_k is defined by an induction that can be seen to parallel the induction defining the inequivalence relation among L^k -types presented above. Initially, the basic types are ordered, and at each inductive stage we refine this to an ordering of the equivalence classes under the equivalence relation obtained through that stage. At any given stage, the symmetric closure of the ordering relation is the same as the inequivalence relation at that stage.

Suppose that at stage r of the induction determined by ϕ above, we have the equivalence classes t_1, \dots, t_m and they have been ordered in that order. Let \bar{a} and \bar{b} be two tuples such that $\phi^{r+1}[\bar{a}\bar{b}]$ but not $\phi^r[\bar{a}\bar{b}]$, i.e. \bar{a} and \bar{b} are in the same equivalence class at stage r , but are differentiated at stage $r + 1$. By the definition of ϕ , this means that for some i , there is at least one tuple that differs from \bar{a} (resp. \bar{b}) only in the i^{th} place such that in its equivalence class there is no tuple that differs from \bar{b} (resp. \bar{a}) only in the i^{th} place. For the smallest such i , let S be the set of all such tuples, and let t_j be the smallest (in the ordering that has been defined so far) equivalence class that is represented in S . There

are two possibilities: either all the elements of S that are in t_j are one move away from \bar{a} or they are all one move away from \bar{b} . There cannot be two tuples in S that are one move away from \bar{a} and \bar{b} respectively and are both in t_j , because then, by definition they would not be in S . Thus, depending on whether t_j is represented in S by a tuple that is one move away from \bar{a} or \bar{b} , we order \bar{a} to be smaller or greater than \bar{b} . In this way, the new collection of equivalence classes obtained from the inequivalence relation ϕ^{r+1} is totally ordered. We now construct a formula that will do this. In the following, for ease of reading, we will use the notation $x_1 \dots x \dots x_k$ to indicate a sequence of variables in which x has been substituted for x_i when the particular i is clear from the context.

Let R be a $2k$ -ary relation symbol. Define the following formulas for each $1 \leq i \leq k$:

$$\begin{aligned} \beta_i(\bar{x}\bar{y}) &\equiv \forall x_i \exists y_i (\neg R(x_1 \dots x_k y_1 \dots y_k) \wedge \neg R(y_1 \dots y_k x_1 \dots x_k)) \wedge \\ &\quad \forall y_i \exists x_i (\neg R(x_1 \dots x_k y_1 \dots y_k) \wedge \neg R(y_1 \dots y_k x_1 \dots x_k)); \\ \delta_i(\bar{x}\bar{y}) &\equiv \bigwedge_{j < i} \beta_j \wedge \neg \beta_i. \end{aligned}$$

R is the predicate variable for the induction we are going to define. Thus, at any given stage, r , R defines an ordering on the equivalence classes defined so far. So, the formula $\neg R(\bar{x}\bar{y}) \wedge \neg R(\bar{y}\bar{x})$ is true of pairs of tuples that are in the same equivalence class. As explained above, for a pair of tuples that becomes inequivalent at stage $r + 1$, we wish to pick the first position in the tuples that can be used to establish this inequivalence. δ_i holds of a pair of tuples just in case i is that position.

Let

$$\theta_0(x_1 \dots x_k y_1 \dots y_k) \equiv \bigvee_{1 \leq i < j \leq k} (\alpha_i(\bar{x}) \wedge \alpha_j(\bar{y})).$$

That is, θ_0 defines a total ordering on the basic L^k -types. To refine this ordering by induction, define for each i ($1 \leq i \leq k$) the following pair of formulas:

$$\begin{aligned} \sigma_1^i(x, \bar{x}\bar{y}) &\equiv \forall y (R(x_1 \dots x \dots x_k y_1 \dots y \dots y_k) \vee R(y_1 \dots y \dots y_k x_1 \dots x \dots x_k)); \\ \sigma_2^i(y, \bar{x}\bar{y}) &\equiv \forall x (R(x_1 \dots x \dots x_k y_1 \dots y \dots y_k) \vee R(y_1 \dots y \dots y_k x_1 \dots x \dots x_k)). \end{aligned}$$

Define the set of tuples $move_i(\langle a_1 \dots a_k \rangle) = \{ \langle a_1 \dots a \dots a_k \rangle \mid a \in A \}$, *i.e.* the tuples obtained by replacing the i^{th} element. If two tuples \bar{a} and \bar{b} are inequivalent at stage $r + 1$ in the induction determined by ϕ , then, for some i , there is a tuple in $move_i(\bar{a})$ which is inequivalent to every tuple in $move_i(\bar{b})$ (or *vice versa*) at stage r . The formula σ_1^i

(parametrized by the tuples \bar{a} and \bar{b}) picks out the elements a such that $a_1 \dots a \dots a_k$ is such a tuple (similarly for σ_2^i). The following formula would then order the tuples \bar{a} and \bar{b} as desired, unless they had already been ordered otherwise.

$$\theta(R, \bar{x}\bar{y}) \equiv \theta_0(\bar{x}\bar{y}) \vee \left(\neg R(\bar{y}\bar{x}) \wedge \bigvee_{1 \leq i \leq k} (\delta_i \wedge \exists x(\sigma_1^i(x, \bar{x}\bar{y}) \wedge \forall y(\sigma_2^i(y, \bar{x}\bar{y}) \rightarrow R(x_1 \dots x \dots x_k y_1 \dots y \dots y_k))) \right).$$

We cannot define the least fixed point of the above formula, since it is not positive in R . However, the inflationary fixed point gives us the required ordering.

$$\psi(z_1 \dots z_{2k}) \equiv \mathbf{ifp}(R, \bar{x}\bar{y})\theta(z_1 \dots z_{2k}).$$

It only remains to establish that if we let P_k be the global relation over the class of finite structures determined by ψ , then P_k satisfies conditions 1-3 of the theorem. By Theorem 2, there is a formula of FO + LFP equivalent to ψ over the class of finite structures, hence P_k satisfies condition 1. In order to show that P_k satisfies conditions 2 and 3, we consider the global relation $\sim_{k,r}$ defined as follows: for every structure \mathfrak{A} and k -tuples \bar{a} and \bar{b} of elements from \mathfrak{A} , $\bar{a} \sim_{k,r} \bar{b}$, if and only if, \bar{a} and \bar{b} satisfy the same formulas of L^k of quantifier rank $\leq r$ in \mathfrak{A} . Let $\theta^r(\mathfrak{A})$ be the r^{th} stage of the inflationary induction determined by the above defined θ over the structure \mathfrak{A} . We leave to the reader the easy induction which establishes that for all \mathfrak{A} and r

- i. $\theta^r(\mathfrak{A})$ is a strict linear pre-order of $|\mathfrak{A}|^k$ and
- ii. the global equivalence relation induced by θ^{r+1} is $\sim_{k,r}$.

It now follows immediately that P_k satisfies conditions 2 and 3 of the theorem. ■

Abiteboul and Vianu [Abiteboul and Vianu, 1992] have independently shown that the L^k types can be uniformly ordered using a formula of FO + LFP. They establish this by means of a computational approach. By contrast, our approach exploits the link between the equivalence of types and the pebble games.

We will write $\bar{x} <_k \bar{y}$ for $\psi(\bar{x}\bar{y})$. Using the formulas just defined, it is possible to define the L^k -equivalence and the corresponding pre-order relation on tuples shorter than k .

Yuri Gurevich [Gurevich, 1992] has pointed out that the formulas constructed in this section that define the inequivalence and the ordering also work on infinite structures, where

they define the inequivalence under $L_{\infty\omega}^k$. In this case, the use of the inflationary fixed-point cannot be eliminated, since Theorem 2 does not apply.

6 Rigid Structures

Consider the pre-order $<_{\frac{1}{k}}$ on single elements, *i.e.* tuples of length 1. Clearly, if there is at most one element of any L^k -type in a structure, then $<_{\frac{1}{k}}$ defines a total ordering on the universe of the structure. Since this ordering is definable in FO + LFP, and since FO + LFP expresses all of P in the presence of ordering (see Theorem 3), this implies that FO + LFP expresses all of P on these structures. We now formalize this argument. Recall that a structure \mathfrak{A} is called *rigid* if the only automorphism on \mathfrak{A} is the identity.

Definition 8 *Call a structure \mathfrak{A} k -rigid if no two elements of \mathfrak{A} have the same L^k -type.*

Note that a structure is k -rigid if and only if $<_{\frac{1}{k}}$ defines a total ordering on the structure. Since there is a sentence of FO + LFP that defines $<_{\frac{1}{k}}$ and says that it is a total order, we have the following result:

Lemma 2 *For all k , the class of k -rigid structures is definable in FO + LFP.*

Combining this with the argument above, showing that FO + LFP expresses all of P on this class, we get:

Theorem 9 *If K is a class of k -rigid structures, then every polynomial time query on K is definable by formula of FO + LFP.*

Observe that any structure with a linear ordering, $<$, is 2-rigid. There is a formula $\alpha_i(x) \in L^2$ which defines the i^{th} element in the ordering uniquely. For instance,

$$\alpha_3(x) \equiv \exists y(y < x \wedge \exists x(x < y \wedge \forall y(\neg y < x))) \wedge \neg \exists y(y < x \wedge \exists x(x < y \wedge \exists y(y < x))).$$

Hence Theorem 9 generalizes Theorem 3. The generalization is proper in the sense that there are classes K (for instance, let $K = \{\{\{0, \dots, n-1\}, S\}\}$, where S is the successor relation) in which every structure is k -rigid, but which do not have a first-order definable order.

Clearly, every k -rigid structure is rigid. Conversely,

Theorem 10 *Every finite rigid structure \mathfrak{A} is k -rigid for some k .*

Proof:

Towards a contradiction, assume that \mathfrak{A} is a rigid structure that is not k -rigid for any k . Then for each k there are distinct elements a_1^k, a_2^k in \mathfrak{A} which have the same L^k -type. Since \mathfrak{A} is finite, this implies that there are distinct a_1, a_2 such that for infinitely many k , a_1 and a_2 have the same L^k -type. But, two elements that share their L^k -type share their L^l -type for all $l < k$. Hence, a_1 and a_2 have the same first-order type. Now, expand the vocabulary by a constant symbol c , and consider the expanded structures $\langle \mathfrak{A}, a_1 \rangle$ and $\langle \mathfrak{A}, a_2 \rangle$. These structures are elementarily equivalent, since a_1 and a_2 have the same first-order type over \mathfrak{A} . But any two finite structures that are elementarily equivalent are isomorphic. Hence there is an automorphism of \mathfrak{A} mapping a_1 to a_2 which contradicts the hypothesis that \mathfrak{A} is rigid. ■

Yuri Gurevich has pointed out that this theorem fails to hold for infinite structures. In particular, there is a countable structure that is rigid, but not k -rigid for any k [Gurevich, 1992].

7 Reduction to an Ordered Structure

Even on non-rigid structures, $<_k$ defines a pre-order, or alternatively, a total ordering on the L^k equivalence classes. We can look at this as the basis for a reduction of the structure \mathfrak{A} onto a totally ordered structure in which each of the equivalence classes is collapsed to a point. This is similar to the construction defined by Abiteboul and Vianu [Abiteboul and Vianu, 1992]. They define a map from equivalence classes of tuples to natural numbers and show that this map can be constructed in FO + LFP. They use this map to establish normal forms for FO + LFP and FO + PFP and to show that these two logics are equivalent if and only if P = PSPACE. We show in this section, how this latter result can be derived from our construction.

Such a translation to an ordered structure is interesting from the following point of view – consider any formula $\phi \equiv \mathbf{lfp}(R, \bar{x})\psi$ (or $\mathbf{pfp}(R, \bar{x})\psi$) with only k variables. Then, not only is the relation defined by ϕ on \mathfrak{A} closed under L^k -equivalence, but so is every iterative stage of ϕ . This raises the possibility that we can describe ϕ as an induction on the L^k

equivalence classes of tuples.

In the rest of the section, we will assume that we are dealing with a signature $\sigma = \langle R_1 \dots R_t \rangle$ in which all the relations have arity at most k . For other signatures, we replace each relation symbol R (of arity $m > k$) with a collection of relation symbols of arity k , one for each way that an m -tuple can be formed from a sequence of k elements. This encoding suffices, since we are only considering formulas with at most k variables.

More formally, for any structure $\mathfrak{A} = \langle A, R_1, \dots, R_t \rangle$, let

$$E_k(\mathfrak{A}) = \langle A^k / \sim_k, <_k, =', R'_1, \dots, R'_t, X_1, \dots, X_k, P_{s_1}, \dots, P_{s_n} \rangle, \text{ where } n = k^k.$$

be the structure defined as follows:

- The universe of $E_k(\mathfrak{A})$ is A^k / \sim_k , *i.e.* the equivalence classes of tuples from A of length k under the equivalence relation \sim_k . We will write $[\bar{a}]$ to denote the equivalence class that includes the tuple $\bar{a} = \langle a_1 \dots a_k \rangle$.
- $<_k$ is the total ordering on the universe of $E_k(\mathfrak{A})$ defined in Section 5.
- $='$ is a unary relation such that $='([\bar{a}])$ holds if and only if $\bar{a} = \langle a_1, a_2 \dots a_k \rangle$ and $a_1 = a_2$. This relation is well-defined, since a tuple in which the first two elements are distinct cannot be equivalent to one in which they are identical, since they differ on a basic type.
- For each relation R_i in \mathfrak{A} , of arity $m \leq k$, we have a unary relation R'_i in $E_k(\mathfrak{A})$ such that $[\langle a_1 \dots a_k \rangle] \in R'_i$ holds if and only if $\langle a_1 \dots a_m \rangle \in R_i$. Again, these relations are clearly well-defined, for the same reason as above.
- X_i – the i^{th} substitution relation – is a binary relation such that $X_i([\bar{a}], [\bar{a}'])$ holds if the tuples \bar{a} and \bar{a}' differ *at most* on their i^{th} element.
- P_s – the sequence relation – is a binary relation for every sequence $s = \langle i_1 \dots i_k \rangle$ of integers from $\{1, \dots, k\}$ such that for any tuple $\langle a_1 \dots a_k \rangle$, $([\langle a_1 \dots a_k \rangle], [\langle a_{i_1} \dots a_{i_k} \rangle]) \in P_s$.

We will also write $E_k(\sigma)$ to denote the signature of $E_k(\mathfrak{A})$, when σ is the signature of \mathfrak{A} .

Lemma 3 *For every first-order formula ϕ with m free variables in the language $E_k(\sigma)$, there is an FO + LFP formula ϕ' with km free variables in the language σ such that for any structure \mathfrak{A} , $\mathfrak{A} \models \phi'[\bar{a}_1 \dots \bar{a}_m]$ if and only if $E_k(\mathfrak{A}) \models \phi[[\bar{a}_1] \dots [\bar{a}_m]]$.*

Proof:

All the relations on $E_k(\mathfrak{A})$, including equality, are definable in FO + LFP on \mathfrak{A} . Moreover, these definitions are uniform, *i.e.* for each $R \in E_k(\sigma)$ there is a single FO + LFP formula defining it for all \mathfrak{A} . So, we obtain ϕ' by substituting this definition for each occurrence of the relation symbol in ϕ . This *includes* substituting the definition of \sim_k for each occurrence of the identity symbol. For each variable in ϕ , we substitute k new variables and for each quantifier, a block of k quantifiers. ■

Let $\phi(R)$ be a first-order formula in the language $E_k(\sigma) \cup \{R\}$ where R is a relation symbol of arity m and let $\langle E_k(\mathfrak{A}), \mathcal{R} \rangle$ be a structure for this language. Let $\mathfrak{A}' = \langle \mathfrak{A}, S^{\mathfrak{A}} \rangle$ be a structure interpreting the language $\sigma \cup S$ with $S^{\mathfrak{A}} = \{ \langle a_1 \dots a_{km} \rangle \mid \mathcal{R}([\bar{a}_1] \dots [\bar{a}_m]) \}$. By the proof of Lemma 3, there is an FO + LFP formula, ϕ' , in $\sigma \cup S$ such that $\mathfrak{A}' \models \phi'[\bar{a}_1 \dots \bar{a}_m]$ if and only if $\langle E_k(\mathfrak{A}), \mathcal{R} \rangle \models \phi[[\bar{a}_1] \dots [\bar{a}_m]]$. This gives us the following result:

Lemma 4 *For every FO + LFP (respectively FO + PFP) formula ϕ , with m free variables, in the language $E_k(\sigma)$, there is an FO + LFP (respectively FO + PFP) formula ϕ' , with km free variables, in the language σ such that for any structure \mathfrak{A} , $\mathfrak{A} \models \phi'[\bar{a}_1 \dots \bar{a}_m]$ if and only if $E_k(\mathfrak{A}) \models \phi[[\bar{a}_1] \dots [\bar{a}_m]]$.*

Note that if ϕ has one free variable, then the relation defined by ϕ' is closed under the L^k -equivalence relation.

We now establish a translation of formulas in the other direction. Let ϕ be a formula of L^k in the language σ . We will define, by induction on the structure of ϕ , a first-order formula ϕ^* in the language $E_k(\sigma)$. In the translation, every sub-formula of ϕ with free variables among $x_1 \dots x_k$ is translated into a sub-formula of ϕ^* with exactly one free variable with the property that $E_k(\mathfrak{A}) \models \phi^*[[\bar{a}]]$ if and only if $\mathfrak{A} \models \phi[\bar{a}]$. That is to say, we will treat ϕ as defining a k -ary relation over \mathfrak{A} even if ϕ has fewer than k free variables. The relation $\{ \langle a_1 \dots a_k \rangle \mid \mathfrak{A} \models \phi[a_1 \dots a_k] \}$ is closed under L^k -equivalence, since $\phi \in L^k$.

The translation is defined as follows:

- If $\phi \equiv x_i = x_j$, then $\phi^*(x) \equiv \exists y (P_s(x, y) \wedge ='(y))$

where s is a sequence chosen so that $s = \langle i, j \dots \rangle$.

- If $\phi \equiv R_j(x_{i_1}, \dots, x_{i_m})$, then $\phi^*(x) \equiv \exists y(P_s(x, y) \wedge R'_j(y))$
where s is a sequence chosen so that $s = \langle i_1, \dots, i_m, \dots \rangle$.
- If $\phi \equiv \neg\psi(\bar{x})$, then $\phi^*(x) \equiv \neg\psi^*(x)$
- If $\phi(\bar{x}) \equiv \psi_1(\bar{x}) \wedge \psi_2(\bar{x})$, then, $\phi^*(x) \equiv \psi_1^*(x) \wedge \psi_2^*(x)$
- If $\phi(\bar{x}) \equiv \exists x_i \psi(\bar{x})$ then, $\phi^*(x) \equiv \exists y(X_i(x, y) \wedge \psi^*(y))$

Claim 2 $\mathfrak{A} \models \phi[\bar{a}]$, if and only if, $E_k(\mathfrak{A}) \models \phi^*[[\bar{a}]]$.

Proof:

We proceed by induction on the formulas:

Basis: if $\phi \equiv x_i = x_j$, $\mathfrak{A} \models \phi[\bar{a}]$ implies that $a_i = a_j = a$, which gives $E_k(\mathfrak{A}) \models P_{\langle i, j, \dots \rangle}([\bar{a}], [\bar{a}'])$ for some \bar{a}' such that $\bar{a}' = \langle a, a, \dots \rangle$, by definition of P_s . So, \bar{a}' witnesses $E_k(\mathfrak{A}) \models \phi^*[[\bar{a}]]$ since $E_k(\mathfrak{A}) \models \equiv'([\bar{a}'])$. On the other hand, if for some \bar{b} , $E_k(\mathfrak{A}) \models P_{\langle i, j, \dots \rangle}([\bar{a}], [\bar{b}]) \wedge \equiv'([\bar{b}])$ then some $\bar{a}' \in [\bar{a}]$ satisfies $a'_i = b'_1$ and $a'_j = b'_2$ for some $\bar{b}' \in [\bar{b}]$. Since $b_1 = b_2$, this means $\mathfrak{A} \models \phi[\bar{a}']$, and since \bar{a} and \bar{a}' have the same L^k -type, $\mathfrak{A} \models \phi[\bar{a}]$.

The argument for the other atomic formulas is similar.

Induction step: The cases of negation and conjunction are trivial. If $\phi \equiv \exists x_i \psi$ and $\mathfrak{A} \models \phi[\bar{a}]$ then there is a \bar{b} which differs from \bar{a} in at most the i^{th} element such that $\mathfrak{A} \models \psi[\bar{b}]$. By induction hypothesis, $E_k(\mathfrak{A}) \models \psi^*[[\bar{b}]]$ and so $[\bar{b}]$ witnesses $E_k(\mathfrak{A}) \models \phi^*[[\bar{a}]]$ since $E_k(\mathfrak{A}) \models X_i([\bar{a}], [\bar{b}])$. In the other direction, if for some \bar{b} , $E_k(\mathfrak{A}) \models (X_i([\bar{a}], [\bar{b}]) \wedge \psi^*[[\bar{b}]])$, then some $\bar{b}' \in [\bar{b}]$ differs from some $\bar{a}' \in [\bar{a}]$ is at most the i^{th} position. But then, there must be a $\bar{b}'' \in [\bar{b}]$ which differs in at most the i^{th} position from \bar{a} . To see this, observe that since \bar{a} and \bar{a}' have the same L^k -type, Player II can indefinitely play the k -pebble game on two copies of the structure \mathfrak{A} with the pebbles initially on these tuples. But then, if we can get to \bar{b}' in one move from \bar{a}' there must be a \bar{b}'' equivalent to \bar{b}' one move away from \bar{a} . Now, by induction hypothesis, $\mathfrak{A} \models \psi[\bar{b}'']$ and hence b''_i witnesses $\mathfrak{A} \models \phi[\bar{a}]$. \blacksquare

In particular, if ϕ is a sentence, then $\mathfrak{A} \models \phi$, if and only if, $E_k(\mathfrak{A}) \models \exists x \phi^*$. Moreover, this is true even if ϕ is in an expanded language $\sigma \cup \{R\}$ as long as the interpretation of R on \mathfrak{A} is closed under L^k -equivalence and $E_k(\mathfrak{A})$ is expanded to interpret R' in the obvious way. This gives us the following result:

Lemma 5 *For every FO + LFP (respectively FO + PFP) formula ϕ in the language σ such that ϕ has at most k distinct variables, there is an FO + LFP (respectively FO + PFP) formula ϕ^* , with one free variable, in the language $E_k(\sigma)$ such that for any structure \mathfrak{A} , $\mathfrak{A} \models \phi[\bar{a}]$ if and only if $E_k(\mathfrak{A}) \models \phi^*[[\bar{a}]]$.*

Proof:

Since ϕ has at most k distinct variables, every iteration of every induction operator in ϕ defines a relation closed under L^k -equivalence. ■

We are now in a position to prove the following result from [Abiteboul and Vianu, 1991b]:

Theorem 11 $\text{FO} + \text{LFP} = \text{FO} + \text{PFP}$ if and only if $\text{P} = \text{PSPACE}$.

Proof:

\Rightarrow This follows immediately from the fact that $\text{FO} + \text{LFP} = \text{P}$ and $\text{FO} + \text{PFP} = \text{PSPACE}$ on ordered structures. (Theorems 3 and 4 respectively).

\Leftarrow Suppose $\text{P} = \text{PSPACE}$. Let ϕ be a formula in FO + PFP over signature σ and let the number of distinct variables in ϕ be k . Take ϕ^* to be the corresponding formula of FO + PFP in the language $E_k(\sigma)$ obtained as in Lemma 5. Since ϕ^* is in FO + PFP, it is computable in PSPACE and hence in P, by hypothesis. Since the structures $E_k(\mathfrak{A})$ have a total ordering on their elements, by Theorem 3, there is a formula ψ , of FO + LFP, equivalent to ϕ^* . Then, by Lemma 4 there is a ψ' in FO + LFP over σ that is equivalent to ϕ . ■

8 Complete Binary Trees

It is easy to see that the size of the structures $E_k(\mathfrak{A})$ is bounded by a polynomial over all structures \mathfrak{A} (see the proof of Theorem 15 in Section 9). Over some classes of structures, it can be considerably smaller. For instance, if we consider the class of complete graphs, there is a bound on the size of the structures $E_k(\mathfrak{A})$ which depends only on k . Another class of structures for which the size of $E_k(\mathfrak{A})$ is much smaller than that of \mathfrak{A} is the class of complete binary trees. This yields some interesting results concerning logical expressibility.

Complete binary trees are graphs, *i.e.*, structures $\langle V, E \rangle$ with one binary relation E , satisfying the following axioms:

$$1. \forall x(\forall y(\neg Exy) \vee \exists y\exists z(y \neq z \wedge Exy \wedge Exz \wedge \forall w(Exw \rightarrow w = y \vee w = z)))$$

This says that every vertex has exactly 0 or 2 children.

$$2. \forall x(\forall y(\neg Eyx) \vee \exists y(Eyx \wedge \forall z(Exz \rightarrow z = y)))$$

This says that every vertex has exactly 0 or 1 parent.

$$3. \exists x(\forall y(\neg Eyx) \wedge \forall z(\forall y(\neg Eyx) \rightarrow x = z))$$

This says that there is exactly one vertex (the root) that has no parent.

$$4. \forall x\forall y\mathbf{lfp}(R, x, y)(x = y \vee \exists z(Rxz \wedge Ezy) \vee \exists z(Ryz \wedge Ezx))(x, y)$$

This says that the graph is connected, *i.e.* every pair of vertices is in the reflexive, transitive and symmetric closure of the edge relation.

$$5. \forall x\neg(\mathbf{lfp}(R, x, y)(Exy \vee \exists z(Rxz \wedge Ezy))(x, x))$$

This says that there are no cycles, *i.e.* no vertex is connected to itself by the transitive closure of the edge relation.

$$6. \forall x\forall y((\forall z(\neg Exz) \wedge \forall z(\neg Eyz)) \rightarrow \delta(x, y))$$

where,

$$\delta \equiv \mathbf{lfp}(R, x, y)((\forall z(\neg Exz) \wedge \forall z(\neg Eyz)) \vee \exists w\exists z(Rwz \wedge Ewx \wedge Ezy))(x, y)$$

This says that all leaves are at the same distance from the root (δ defines an equivalence relation that relates vertices at the same depth).

If we let CBT denote the set of complete binary trees, then by the above definition $CBT \in \text{FO} + \text{LFP}$. Moreover, since we used only four distinct variables, $CBT \in L_{\infty\omega}^4$.

Define the formulas α_n recursively as follows:

$$\alpha_0(x) \equiv \forall y\neg Eyx$$

$$\alpha_{n+1}(x) \equiv \exists y(Eyx \wedge \exists x(x = y \wedge \alpha_n(x)))$$

Then, for $T \in CBT$, $T \models \alpha_d[v]$ just in case v is a vertex of depth d in T . So, if T_d is a complete binary tree, it has depth d if and only if $T_d \models \exists x(\alpha_d) \wedge \neg\exists x(\alpha_{d+1})$. Note that

each α_n contains only two distinct variables. Since any two complete binary trees of the same depth are isomorphic, we can conclude the following:

Lemma 6 *If T_1 and T_2 are two complete binary trees such that $T_1 \equiv_2 T_2$, then $T_1 \cong T_2$.*

Combining this with the axiomatization above, we get the following result:

Lemma 7 *If q is any query in the language of graphs consisting only of complete binary trees, then q is definable in $L_{\infty\omega}^4$.*

Define the class, \mathcal{T} , of labeled binary trees as the class of structures over the vocabulary $\{E, U\}$ which satisfy, in addition to the above six axioms, the following one:

$$7. \forall x \forall y (\delta(x, y) \rightarrow ((Ux \wedge Uy) \vee (\neg Ux \wedge \neg Uy)))$$

That is, all vertices at the same depth are either labeled or unlabeled.

Observe that the propositions shown above for complete binary trees apply equally well to labeled binary trees.

We also define the class, \mathcal{B} , of binary strings as structures over the same vocabulary $\{E, U\}$ that make true Axioms 2 through 5 above, as well as:

$$1'. \forall x (\forall y (\neg Exy) \vee \exists y (Exy \wedge \forall z (Exz \rightarrow z = y)))$$

That is every vertex has exactly 0 or 1 children.

There is a natural correspondence between labeled binary trees and binary strings. In some sense, they encode the same information, with the i^{th} bit of the binary string corresponding to the i^{th} level of the tree. While we give formal definitions below, it will be instructive to keep this intuitive picture in mind and we will make appeal to it to simplify the presentation.

Definition 9 *If $B \in \mathcal{B}$ and $T \in \mathcal{T}$, then $B \triangleright T$ if and only if, for all d :*

- $B \models \exists x \alpha_d$ if and only if $T \models \exists x \alpha_d$, and
- $B \models \forall x (\alpha_d \rightarrow Ux)$ if and only if $T \models \forall x (\alpha_d \rightarrow Ux)$.

Note that if $B \triangleright T$ and the size of B is n , then the size of T is $2^n - 1$

Definition 10 For any queries $q_B \subseteq \mathcal{B}$ and $q_T \subseteq \mathcal{T}$, define:

$$\begin{aligned} h(q_B) &= \{T \mid B \triangleright T \text{ for some } B \in q_B\} \\ h^{-1}(q_T) &= \{B \mid B \triangleright T \text{ for some } T \in q_T\} \end{aligned}$$

It should be clear that $h^{-1}(h(q_B)) = q_B$.

Lindell[Lindell, 1991] used this correspondence between binary strings and labeled binary trees to show that FO + LFP does not express all the polynomial-time queries on binary trees.

Lemma 8 If $q_B \in \text{DTIME}[2^{O(n)}]$ then $h(q_B) \in \text{P}$.

Proof:

Given an input T , we can verify that it is a labeled binary tree in polynomial time, since $T \in \text{FO} + \text{LFP}$. We can also extract from it a B such that $B \triangleright T$ in $\text{DSPACE}[\log(n)]$. We then pass B as the input to the acceptor for q_B which runs in time $2^{O(d)}$, where d is the size of B , but this is only polynomial in the size of T which is $2^d - 1$. ■

Lemma 9 If $q_T \in \text{FO} + \text{LFP}$, then $h^{-1}(q_T) \in \text{FO} + \text{LFP}$.

The proof of this lemma is based on a syntactic translation similar to the one given in Section 7. The key element of Lindell's construction is that k -tuples of vertices from the tree can be encoded as fixed length tuples in the corresponding binary string. This is because a complete set of invariants (up to automorphism) for a tuple on a complete binary tree is the sequence of depths of the least common ancestors of pairs of elements in the tuple. We refer to [Lindell, 1991] for details of the translation.

Given that there are queries on strings in $\text{DTIME}[2^{O(n)}]$ that are not in P [Hartmanis and Stearns, 1965], we conclude the following:

Theorem 12 ([Lindell, 1991]) *There is a $q_T \subseteq \mathcal{T}$ such that $q_T \in \text{P}$, but $q_T \notin \text{FO} + \text{LFP}$.*

Since we observed above that for every q such that $q \subseteq \mathcal{T}$, $q \in L_{\infty\omega}^4$, we conclude that:

Corollary 8 $\text{FO} + \text{LFP} \subset L_{\infty\omega}^\omega \cap \text{P}$.

This result has been independently proved by Cosmadakis [Cosmadakis, 1991] and in [Abiteboul and Vianu, 1992].

Define the class $\text{FO} + \text{PFP}|\text{P}$ of queries expressed by a formula of $\text{FO} + \text{PFP}$ with the property that there is a polynomial p such that every occurrence of the **pfp** operation in the formula closes in $p(n)$ steps in any structure of size n . Any query in $\text{FO} + \text{PFP}|\text{P}$ is clearly computable in polynomial time. Also, since the operator **lfp** can be seen as an instance of **pfp** that always closes in polynomially many steps, we get

$$\text{FO} + \text{LFP} \subseteq \text{FO} + \text{PFP}|\text{P} \subseteq L_{\infty\omega}^{\omega} \cap \text{P}$$

It had been conjectured that these three classes are, in fact, equal. We have shown above that the first and the third can be separated. Abiteboul and Vianu [Abiteboul and Vianu, 1991b] have recently shown that the first and the second are equal if and only if $\text{P} = \text{PSPACE}$. They prove this result using a padding technique similar to the one above. We encoded binary strings of size n as trees of size 2^n . For the purpose of the next result, we will need to encode them into trees of size 2^{n^k} . To this end, we introduce, for every k the class of structures \mathcal{T}_k over the signature $\{E, U, L\}$. The trees in \mathcal{T}_k have depth n^k with the first n levels labeled by the unary relation L . Formally, \mathcal{T}_k is the class of structures which in addition to the Axioms 1 through 7, satisfy:

$$8. \forall x \forall y (\delta(x, y) \rightarrow ((Lx \wedge Ly) \vee (\neg Lx \wedge \neg Ly)))$$

That is, all vertices at the same depth are either in L or not.

$$9. \forall x \forall y ((Lx \wedge Eyx) \rightarrow Ly)$$

If a vertex is in L , then so is its parent.

10. The depth of the tree is n^k , where n is the number of levels labeled by L . This can be stated in $\text{FO} + \text{LFP}$ by defining a k -ary induction on the levels in L that is an ordering of length n^k on k -tuples.

The binary string encoded by a tree in \mathcal{T}_k of depth n^k can be extracted by looking at the topmost n levels (the levels labeled by L) and looking at the string defined by the relation U on these levels. We can formalize this as before with a map h_k from queries on binary strings to queries on \mathcal{T}_k .

Definition 11 If $B \in \mathcal{B}$ and $T \in \mathcal{T}_k$, then $B \triangleright T$ if and only if, for all d :

- $B \models \exists x \alpha_d$ if and only if $T \models \exists x (Lx \wedge \alpha_d)$, and
- $B \models \forall x (\alpha_d \rightarrow Ux)$ if and only if $T \models \forall x (Lx \rightarrow (\alpha_d \rightarrow Ux))$.

Definition 12 For any queries $q_B \subseteq \mathcal{B}$ and $q_T \subseteq \mathcal{T}_k$, define:

$$\begin{aligned} h_k(q_B) &= \{T \mid B \triangleright T \text{ for some } B \in q_B\} \\ h_k^{-1}(q_T) &= \{B \mid B \triangleright T \text{ for some } T \in q_T\} \end{aligned}$$

We can define a syntactic translation of formulas that corresponds to the map h_k :

Definition 13 Given a formula ϕ in the language of binary strings, let ϕ' be defined inductively as follows:

- if ϕ is $x = y$ then ϕ' is $\delta(x, y)$ where δ is as defined in Axiom 6.
- if ϕ is $\neg\psi$ or $\psi_1 \wedge \psi_2$ then ϕ' is $\neg\psi'$ or $\psi'_1 \wedge \psi'_2$ respectively.
- if ϕ is $\exists x \psi$ then ϕ' is $\exists x (Lx \wedge \psi')$.

Suppose for some formula ϕ and some $B \in \mathcal{B}$, $B \models \phi[b_1 \dots b_n]$ and for some sequence of integers $d_1 \dots d_n$, $B \models \alpha_{d_i}[b_i]$, i.e. the depths of the points b_i are given by the d_i . Also, let $T \in \mathcal{T}_k$ be such that $B \triangleright T$. Then, if $t_1 \dots t_n$ are any points in T such that $T \models \alpha_{d_i}[t_i]$, then $T \models \phi'[b_1 \dots b_n]$. This can be verified by an easy induction on the structure of the formula. One consequence of this is the following result.

Lemma 10 If $q \subseteq \mathcal{B}$ is a query in FO + LFP (respectively FO + PFP), then $h_k(q)$ is in FO + LFP (respectively FO + PFP).

Proof:

Let ϕ be the sentence that expresses q and let χ be the conjunction of Axioms 1 through 10. Then, $\phi' \wedge \chi$ expresses $h_k(q)$. ■

Another consequence is that if $B \triangleright T$ then the closure ordinal of any occurrence of **fp** (or **lfp**) in ϕ over B is the same as the closure ordinal of the corresponding occurrence in ϕ' over T . There may be additional inductions in ϕ' which were introduced when we substituted the formula δ for the identity, but all these are defined in **lfp**. Thus all inductions in ϕ'

close in a number of steps polynomial in the size of T . Moreover, χ is defined in $\text{FO} + \text{LFP}$, so all inductions that occur there are also polynomial. We can now prove the following.

Theorem 13 ([Abiteboul and Vianu, 1991b]) $\text{FO} + \text{PFP} \mid \text{P} = \text{FO} + \text{LFP}$ if and only if $\text{PSPACE} = \text{P}$.

Proof:

One direction follows immediately from Theorem 11. In the other direction, suppose $\text{FO} + \text{PFP} \mid \text{P} = \text{FO} + \text{LFP}$. Let S be a language in PSPACE and hence in $\text{DTIME}[2^{n^k}]$ for some k . Let $q_B \subseteq \mathcal{B}$ be the collection of structures corresponding to strings in S . Since an ordering is easily (in $\text{FO} + \text{LFP}$) definable on structures in \mathcal{B} , $q_B \in \text{FO} + \text{PFP}$. Hence $h_k(q_B) \in \text{FO} + \text{PFP}$, by Lemma 10 and as we argued above, all inductions in the sentence expressing $h_k(q_B)$ are polynomial in the size of T . Thus, $h_k(q_B) \in \text{FO} + \text{PFP} \mid \text{P}$. By hypothesis, then, $h_k(q_B) \in \text{FO} + \text{LFP}$ and by an application of Lemma 9, $q_B \in \text{P}$. ■

This result is remarkable in that it reduces the separation of P and PSPACE to the separation of two classes that are properly contained in P .

9 L^k Canonical Structures

In this section, we examine the question of whether the properties in the class $L_{\infty\omega}^\omega \cap \text{P}$ are recursively indexable. Can we enumerate a set of Turing machines, for instance, each of which accepts a property in this class and such that every property in the class is accepted by some machine in the set. We know that the class $\text{FO} + \text{LFP}$ is recursively indexable, since there is an effective way to construct, from a sentence of $\text{FO} + \text{LFP}$, a machine that accepts all models of the sentence. On the other hand, it is not known if the class P , of polynomial time computable queries, is recursively indexable.

Suppose we have a Turing machine C_k , for every $k \in \mathbf{N}$, which computes a function F_k of the input with the property that $F_k(\mathfrak{A}) \equiv_k \mathfrak{A}$ and if $\mathfrak{A} \equiv_k \mathfrak{B}$ then $F_k(\mathfrak{A}) = F_k(\mathfrak{B})$. We say that C_k computes an L^k canonical structure or an L^k -canon of its input. Suppose further that each of the C_k computes in polynomial time. If this is indeed the case, then the class $L_{\infty\omega}^\omega \cap \text{P}$ is recursively indexable. To see this, consider an enumeration of all polynomial time Turing machines M_1, \dots, M_i, \dots . We can then enumerate all machines of

the form $C_k \rightarrow M_i$ which accepts input \mathfrak{A} if and only if M_i accepts $F_k(\mathfrak{A})$. This is an indexing of the class $L_{\infty\omega}^\omega \cap P$.

The situation is similar in the case of the class P . If we could canonically label a structure in polynomial time, then the class P would be recursively indexable. However, in this case, even the problem of testing equivalence (*i.e.* the isomorphism problem) is not known to be in P . We can, however, test the equivalence of two structures under the relation \equiv_k in polynomial time. We can do this by computing the map E_k on the two structures and comparing the result. We show below that $E_k(\mathfrak{A})$ and $E_k(\mathfrak{B})$ are isomorphic just in case $\mathfrak{A} \equiv_k \mathfrak{B}$. Because $E_k(\mathfrak{A})$ and $E_k(\mathfrak{B})$ are ordered structures, if they are isomorphic, they are represented by identical bit-strings.

We now give the proof that the map E_k does indeed compute an L^k -invariant structure and that this computation can be done in polynomial time.

Theorem 14 *For any two structures \mathfrak{A} and \mathfrak{B} , $\mathfrak{A} \equiv_k \mathfrak{B}$ if and only if $E_k(\mathfrak{A}) \cong E_k(\mathfrak{B})$*

Proof:

\Rightarrow If $\mathfrak{A} \equiv_k \mathfrak{B}$ then every L^k -type that is realized in \mathfrak{A} is realized in \mathfrak{B} and vice versa.

To see this, let \bar{a} be a k -tuple from \mathfrak{A} . Recall from Corollary 7 that there is a formula $\phi(x_1 \dots x_k)$ in L^k with k free variables that expresses this type. But then, $\mathfrak{A} \models \exists x_1 \dots x_k \phi$ and therefore $\mathfrak{B} \models \exists x_1 \dots x_k \phi$. This tells us that the structures $E_k(\mathfrak{A})$ and $E_k(\mathfrak{B})$ have the same size.

Let f be the order-preserving map from $E_k(\mathfrak{A})$ to $E_k(\mathfrak{B})$. If $f([\bar{a}]) = [\bar{b}]$, then \bar{a} and \bar{b} have the same L^k -type. This is because the definition of the ordering relation $<_k$ is uniform, that is to say that the same types in different structures are ordered in the same way. As a result, the relations $='$ and R_i are clearly preserved by f . Consider the case $X_i([\bar{a}], [\bar{a}'])$. Let $\phi(x_1 \dots x_k)$ be the L^k formula expressing the L^k -type of \bar{a}' . Then, $\exists x_i \phi$ is in the L^k -type of \bar{a} and hence of any element of $f([\bar{a}])$. It follows that $X_i(f([\bar{a}]), f([\bar{a}']))$. Similarly, f preserves the relation P_s for $s = \langle i_1 \dots i_k \rangle$ because if ϕ is a formula in $\text{Type}_k(\langle a_1 \dots a_k \rangle)$, then the formula ϕ_s obtained by substituting every free occurrence of every x_j with x_{i_j} , with the appropriate renaming of bound variables, is in $\text{Type}_k(\langle a_{i_1} \dots a_{i_k} \rangle)$. Thus, f is an isomorphism.

\Leftarrow Let f be an isomorphism from $E_k(\mathfrak{A})$ to $E_k(\mathfrak{B})$. We show that Player II has a strategy

for playing the k -pebble game on \mathfrak{A} and \mathfrak{B} indefinitely. Suppose that at some stage of the game, the pebbles are on the elements \bar{a} and \bar{b} . (We assume that all k pairs of pebbles are on the board. If not, then just consider any extension of these tuples.) Further suppose, without loss of generality, that Player I moves on \mathfrak{A} resulting in the configuration \bar{a}' . Player II finds a tuple $\bar{b}' \in f([\bar{a}'])$ such that \bar{b}' is one move away from \bar{b} and then plays that move. We need to show that such a \bar{b}' can always be found. Note that we can assume, as an inductive hypothesis that $f([\bar{a}]) = [\bar{b}]$. Suppose that Player I moves the pebble from a_i to a new element, then $X_i([\bar{a}], [\bar{a}'])$ holds. Because f is an isomorphism, $X_i(f([\bar{a}]), f([\bar{a}']))$ holds and we can get from \bar{b} to some $\bar{b}' \in f([\bar{a}'])$ by moving the pebble on b_i ■

Theorem 15 *The map E_k is computable on all structures \mathfrak{A} in time polynomial in the size of the structure \mathfrak{A} .*

Proof:

The number of tuples in A^k is n^k where $n = |A|$. The equivalence relation \sim_k is defined by an FO + LFP formula and hence computable in polynomial time as is the ordering $<_k$. We can get, therefore, in polynomial time, a representation of the universe of $E_k(\mathfrak{A})$. All the other relations are easily defined on \mathfrak{A} (in FO). ■

The most direct approach to constructing an L^k -canon, given a polynomial time algorithm for the translation E_k , would be to try and invert E_k , *i.e.* given an input structure \mathcal{K} , to find an \mathfrak{A} such that $\mathcal{K} = E_k(\mathfrak{A})$. However, this cannot be done in time polynomial in the size of \mathcal{K} . To see this, suppose for contradiction that we have a polynomial time computable E_k^{-1} which acts as a translation from the range of E_k into its domain. Since the range of E_k consists of totally ordered structures, E_k^{-1} is definable in FO + LFP. Composing this with the FO + LFP definition of E_k , we get an FO + LFP translation that yields an L^k canon, and therefore that $L_{\infty\omega}^\omega \cap \text{P} \subseteq \text{FO} + \text{LFP}$, which we know to be false. It is still conceivable that the computation of E_k^{-1} , while not polynomial in the size of the input $E_k(\mathfrak{A})$ is polynomial in the size of \mathfrak{A} , since the former could be much smaller. In fact, it is exactly the case where $E_k(\mathfrak{A})$ is much smaller than \mathfrak{A} that demonstrated that $\text{FO} + \text{LFP} \neq L_{\infty\omega}^\omega \cap \text{P}$.

References

- [Abiteboul and Vianu, 1991a] S. Abiteboul and V. Vianu. Datalog extensions for database queries and updates. *Journal of Computer and System Sciences*, 43:62–124, 1991.
- [Abiteboul and Vianu, 1991b] S. Abiteboul and V. Vianu. Generic computation and its complexity. In *Proceedings of the 23rd ACM Symposium on the Theory of Computing*, 1991.
- [Abiteboul and Vianu, 1992] S. Abiteboul and V. Vianu. Computing with first-order logic. to appear in *Journal of Computer and System Sciences*, 1992.
- [Barwise, 1973] J. Barwise. Back and forth through infinitary logic. In M. D. Morley, editor, *Studies in Model Theory*. Mathematical Association of America, 1973.
- [Barwise, 1977] J. Barwise. On Moschovakis closure ordinals. *Journal of Symbolic Logic*, 42:292–296, 1977.
- [Blass *et al.*, 1985] A. Blass, Y. Gurevich, and D. Kozen. A zero-one law for logic with a fixed point operator. *Information and Control*, 67:70–90, 1985.
- [Cai *et al.*, 1989] J-y. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. In *Proceedings of the 30th IEEE Symposium on Foundations of Computer Science*, pages 612–617, 1989.
- [Chandra and Harel, 1982] A. Chandra and D. Harel. Structure and complexity of relational queries. *Journal of Computer and System Sciences*, 25:99–128, 1982.
- [Cosmadakis, 1991] S. Cosmadakis, 1991. Personal Communication.
- [Ehrenfeucht, 1961] A. Ehrenfeucht. An application of games to the completeness problem for formalized theories. *Fund. Math.*, 49:129–141, 1961.
- [Fraïssé, 1954] R. Fraïssé. Sur quelques classifications des systèmes de relations. *Publications Scientifiques de l'Université d'Algerie, Séries A*, 1:35–182, 1954.
- [Gurevich and Shelah, 1986] Y. Gurevich and S. Shelah. Fixed-point extensions of first-order logic. *Annals of Pure and Applied Logic*, 32:265–180, 1986.
- [Gurevich, 1979] Y. Gurevich. Modest theory of short chains. I. *Journal of Symbolic Logic*, 44(4):481–490, 1979.
- [Gurevich, 1992] Y. Gurevich, 1992. Personal Communication.
- [Hartmanis and Stearns, 1965] J. Hartmanis and R. E. Stearns. On the computational complexity of algorithms. *Transactions of the AMS*, 117:285–306, 1965.

- [Immerman and Lander, 1990] N. Immerman and E. S. Lander. Describing graphs: A first-order approach to graph canonization. In A. Selman, editor, *Complexity Theory Retrospective*. Springer-Verlag, 1990.
- [Immerman, 1982] N. Immerman. Upper and lower bounds for first-order expressibility. *Journal of Computer and System Sciences*, 25:76–98, 1982.
- [Immerman, 1986] N. Immerman. Relational queries computable in polynomial time. *Information and Control*, 68:86–104, 1986.
- [Karp, 1965] C. Karp. Finite quantifier equivalence. In J.W. Addison, L. Henkin, and A. Tarski, editors, *The Theory of Models*, pages 407–412. North Holland, 1965.
- [Kolaitis and Vardi, 1990] Ph. G. Kolaitis and M. Y. Vardi. 0-1 laws for infinitary logics. In *Proceedings of the 5th IEEE Symposium on Logic in Computer Science*, pages 156–167, 1990.
- [Kolaitis and Vardi, 1992a] Ph. G. Kolaitis and M. Y. Vardi. Fixpoint logic vs. infinitary logic in finite-model theory. In *Proc. 7th IEEE Symp. on Logic in Computer Science*, pages 46–57, 1992.
- [Kolaitis and Vardi, 1992b] Ph. G. Kolaitis and M. Y. Vardi. Infinitary logics and 0-1 laws. *Information and Computation*, 98(2):258–294, 1992.
- [Läuchli, 1968] H. Läuchli. A decision procedure for the weak second order theory of linear order. In H.A. Schmidt, K. Schütte, and H.-J. Thiele, editors, *Contributions to Mathematical Logic*, pages 189–197. North Holland, 1968.
- [Lindell, 1991] S. Lindell. An analysis of fixed-point queries on binary trees. *Theoretical Computer Science*, 85(1):75–95, 1991.
- [Poizat, 1982] B. Poizat. Deux ou trois choses que je sais de L_n . *Journal of Symbolic Logic*, 47(3):641–658, 1982.
- [Rubin, 1975] A. Rubin. *Free Algebras in Von Neumann-Bernays-Gödel Set Theory and Positive Elementary Inductions in Reasonable Structures*. PhD thesis, California Institute of Technology, 1975.
- [Scott, 1965] D. Scott. Logic with denumerably long formulas and finite strings of quantifiers. In J.W. Addison, L. Henkin, and A. Tarski, editors, *The Theory of Models*, pages 329–341. North Holland, 1965.
- [Shelah, 1975] S. Shelah. The monadic theory of order. *Annals of Mathematics*, 102:379–419, 1975.
- [Vardi, 1982] M. Y. Vardi. The complexity of relational query languages. In *Proceedings of the 14th ACM Symposium on the Theory of Computing*, pages 137–146, 1982.