Finite Models and Finitely Many Variables

Anuj Dawar
Department of Computer Science
University of Wales Swansea
Swansea, SA2 8PP, U.K.
e-mail: a.dawar@swansea.ac.uk

Abstract

This paper is a survey of results on finite variable logics in finite model theory. It focusses on the common underlying techniques that unite many such results.

1 Introduction

Finite variable logics have come to occupy a very important place in finite model theory. This survey examines a number of the results that have been established and the techniques that have been used in this connection. Taking a broad enough view allows a picture to emerge which shows that essentially the same techniques have re-appeared in differing guises in entirely different contexts. The questions that motivated Poizat's work on classification seem unconnected with McColm's conjectures, which in turn bear only an incidental relationship with the question of Chandra and Harel that motivated Abiteboul and Vianu's theorem and related work on relational complexity. The fact that finite variable logics play an important role in each case supports the view that they are in some way central to the model theory of finite structures. By focussing on the common techniques, this survey aims to expose the underlying connections between a variety of investigations. It is hoped that this will help to explain the importance of finite variable logics, as well as the breadth of applicability of the ideas that have been developed. The paper does not, however, aim to be comprehensive in its coverage of the work on finite variable logics in finite model theory as several strands of this work are omitted for lack of space. Significant among these is the work on finite variable logics and counting which has been covered in the excellent work by Otto [50], and the relation of finite variable logics to modal and temporal logics for which a good starting point is the survey by Hodkinson [34].

One of the central concerns of finite model theory is to study the limits of the expressive power of logical languages on finite structures. It is in this context that questions of a model theoretic nature arise naturally with respect to finite models. A large part of the motivation and the impetus for the study of finite model theory has come from its connections with computational complexity. There is a close connection between the *descriptive complexity* of a property of finite structures — i.e. the complexity of the logical constructs required to define the property — and its *computational complexity*, the resources required on a machine to decide the property. The paradigmatic result relating these two kinds of complexity is the theorem of Fagin [24] which states that a property is definable in existential second order logic if, and only if, it is decidable in NP. Much work in the subject has been generated by the question first posed by Chandra and Harel [13] asking whether one can give a similar logical characterisation for the complexity class P (see Gurevich [31] for a detailed discussion and a precise formulation of this question).

1.1 From Infinite to Finite Structures

Model theoretic questions in the context of finite structures are rather different from those that arise in the classical setting. One such difference that is immediately apparent on examining the literature on finite model theory is that this branch of logic is not centrally concerned with first order predicate logic. It must be said, however, that the model theory of infinite structures is no longer concerned with first order logic to the extent that it used to be.

One reason why first order logic is not central to finite model theory is that the relation of elementary equivalence between structures is trivial on finite structures. A large part of classical model theory can arguably be described as the study of this equivalence relation and the structure of its equivalence classes. Two structures $\mathfrak{A}$ and $\mathfrak{B}$ are elementarily equivalent if, for every first order sentence φ,

$$\mathfrak{A} \models \varphi \text{ if, and only if, } \mathfrak{B} \models \varphi.$$

This is crucial in establishing inexpressibility results. For instance, by proving that all dense linear orders without endpoints are elementarily equivalent, we establish that other properties that might distinguish such orders (such as Dedekind completeness) are not definable.

On finite structures, the elementary equivalence relation is trivial, in that any two elementarily equivalent structures are isomorphic. Indeed, any finite structure is described up to isomorphism by a single sentence. Given a structure $\mathfrak{A} = (A, R_1, \ldots, R_m)$, where A is a set of n elements, we can construct a sentence

$$\delta_{\mathfrak{A}} = \exists x_1 \ldots \exists x_n (\psi \wedge \forall y \bigvee_{1 \leq i \leq n} y = x_i)$$

where, $\psi(x_1, \ldots, x_n)$ is the conjunction of all atomic and negated atomic formulas that hold in $\mathfrak{A}$. Now, for any structure $\mathfrak{B}$, $\mathfrak{B} \models \delta_{\mathfrak{A}}$ if, and only if, $\mathfrak{A} \cong \mathfrak{B}$.

This means that first order logic can make all the distinctions that are to be made between finite structures. Still, the expressive power of first order logic on finite structures is weak: for any first order sentence φ, $\mathrm{Mod}(\varphi)$ can be decided by a deterministic Turing machine with logarithmic work space, where $\mathrm{Mod}(\varphi)$ is the collection of finite models of φ. What accounts for this disparity?

Essentially, in the model theory of infinite structures, as we classify structures by elementary equivalence, we are looking at the expressive power of theories, i.e. possibly infinite sets of sentences. Two structures that are elementarily equivalent cannot be distinguished by any first order theory. In contrast, any isomorphism closed class S of finite structures is defined by the set of negations of the sentences $\delta_{\mathfrak{A}}$, as above, for finite structures $\mathfrak{A}$ not in S. Certainly, this theory may have infinite models, but the collection of its finite models is exactly S. To establish inexpressibility results for first order *sentences*, we need to consider weaker equivalence relations than elementary equivalence.

There are two different weakenings of the elementary equivalence relation that have been extensively studied in the context of finite model theory. One stratifies this relation according to quantifier rank. Thus, we define for each positive integer p, the relation $\equiv_p$ given by: $\mathfrak{A} \equiv_p \mathfrak{B}$ if, and only if, $\mathfrak{A}$ and $\mathfrak{B}$ satisfy the same sentences of quantifier rank p or less. These relations are very useful in characterising definability in first order logic. Indeed, in a finite relational language, a class of structures is definable by a first order sentence if, and only if, it is closed under the relation $\equiv_p$ for some fixed p. This stratification of elementary equivalence has also been used in analysing the expressive power of fragments of second order logic (see [9, 26, 22] *inter alia*).

However, for many questions arising in finite model theory, a more interesting stratification of the elementary equivalence relation is obtained by counting, not the quantifier rank of a sentence, but the number of distinct variables that occur in it. So, we define for each positive integer k, the relation $\equiv^k$ given by: $\mathfrak{A} \equiv^k \mathfrak{B}$ if, and only if, $\mathfrak{A}$ and $\mathfrak{B}$ satisfy the same sentences with k variables or less. There is a corresponding notion of definability with k variables. Axiomatisability with a finite number of variables is a notion intermediate between finite axiomatisability and full axiomatisability.

Finite variable axiomatisability does not in itself capture any interesting complexity class. On the one hand there are undecidable classes of finite structures that are axiomatisable with just two variables (see [41] for examples) while on the other hand, some classes of very low complexity are not axiomatisable with a finite number of variables. Nonetheless, the notion is still interesting from two points of view: 1) many complexity theoretic questions can be reproduced here, as we shall see; and 2) many interesting model theoretic questions arise here. It seems, therefore, to offer fertile ground for the interaction between model theory and complexity theory.

1.2 Notation

In the following, $\mathfrak{A}$ and $\mathfrak{B}$ always denote relational structures of some fixed finite vocabulary σ. The symbols A and B are used to denote the universes of the structures $\mathfrak{A}$ and $\mathfrak{B}$ respectively. In general, these are also assumed to be finite. When we deal with infinite structures, this will be stated explicitly. For any sentence φ, $\mathrm{Mod}(\varphi)$ denotes the collection of *finite* models of φ.

We write L^k to denote the fragment of first order logic in which the only variables allowed are $x_1, \ldots, x_k$. Given two structures $\mathfrak{A}$ and $\mathfrak{B}$ and two tuples $s \in A^l$ and $t \in B^l$ (where $l \leq k$), we write $(\mathfrak{A}, s) \equiv^k (\mathfrak{B}, t)$ to denote that for every formula φ of L^k, $\mathfrak{A} \models \varphi[s]$ if, and only if, $\mathfrak{B} \models \varphi[t]$. Note that this is not the same as saying that the expanded structures $(\mathfrak{A}, s)$ and $(\mathfrak{B}, t)$ agree on all sentences of L^k, since constants denoting the elements of the tuples would be additional to the k variables. For any structure $\mathfrak{A}$ and tuple s, $\mathrm{Type}^k(\mathfrak{A}, s)$ denotes the set $\{\varphi \in L^k \mid \mathfrak{A} \models \varphi[s]\}$.

The infinitary logic $L^k_{\infty\omega}$ is obtained by closing L^k under conjunction and disjunction of arbitrary (possibly infinite) sets of formulas. We write $\equiv^k_{\infty\omega}$ for the corresponding equivalence relation. The logic $L^\omega_{\infty\omega}$ is given by $L^\omega_{\infty\omega} = \bigcup_{k\in\omega} L^k_{\infty\omega}$. That is, it consists of infinitary formulas in which each formula has only finitely many variables.

1.3 Fixed Point Logics

The interest in studying these infinitary logics, from the point of view of finite model theory arises from the fact that they have proved useful in analysing the expressive power of fixed point logics. We next take a look at these.

Consider a formula $\varphi(R, x_1, \ldots, x_l)$, where R is a relational variable of arity l. Given any (finite or infinite) structure $\mathfrak{A}$, this formula defines a map, which takes a relation $S \subseteq A^l$, to a new relation $\varphi^{(\mathfrak{A},S)} = \{s \in A^l \mid \mathfrak{A} \models \varphi[S, s]\}$, i.e. the relation defined by φ in the expanded structure $(\mathfrak{A}, S)$. If all occurrences of the variable R in φ are positive, then the map defined is monotone and is guaranteed to have a least fixed point. Moreover, by well known properties of lattice maps, this least fixed point can be arrived at by iterating the map. Specifically, we define the following transfinite sequence of relations.

$$\begin{aligned} \varphi_0 &= \emptyset \\ \varphi_{\alpha+1} &= \varphi^{(\mathfrak{A},\varphi_\alpha)} \\ \varphi_\lambda &= \textstyle\bigcup_{\beta<\lambda} \varphi_\beta \quad \text{for limit ordinals } \lambda. \end{aligned}$$

There is an ordinal γ such that $\varphi_\gamma = \varphi_{\gamma+1}$, and φ_γ is the least fixed point of the map defined by φ. We call the least such ordinal the closure ordinal of φ on $\mathfrak{A}$. Since the operator is monotone, the sequence of relations defined above is increasing, and therefore the closure ordinal is bounded by κ^l, where $\kappa = \mathrm{card}(A)$. Of course, if $\mathfrak{A}$ is infinite, the exponent l is superfluous. On the other hand, if $\mathfrak{A}$ is finite, the closure ordinal must be finite, and the third clause in the definition above is unnecessary.

Rubin [54] (see also [10]) observed that for every formula φ (of first order logic), there is a k such that for all ordinals α, there is a formula of $L^k_{\infty\omega}$ that defines the relation φ_α on *all* structures. Without fear of confusion, we call this formula φ_α as well. One can further note that when α is finite, then φ_α is given by a finitary formula, i.e. a formula of L^k.

While inductive definitions were extensively studied by logicians in the 1960s and 70s (see [49, 6]), a fixed point logic, which is the closure of first order logic under such definitions was first proposed in the context of finite models[1] by Aho and Ullman [8]. To be precise, the logic LFP is obtained by closing first order logic under the following formula formation rule:

> If φ is a formula, R an l-ary relational variable occurring only positively in φ, $\bar{x}$ is an l-tuple of elementary variables, and $\bar{t}$ an l-tuple of terms, then
>
> $$\mathbf{lfp}_{R,\bar{x}}\varphi(\bar{t})$$
>
> is also a formula, with all occurrences of R and variables in $\bar{x}$ in φ bound, while all occurrences of variables in $\bar{x}$ in $\bar{t}$ remain free.

The intended meaning of the formula is that the tuple denoted by $\bar{t}$ is in the relation that is the least fixed point of the operator defined by $\varphi(R, \bar{x})$.

It follows immediately from Rubin's observation that on a fixed structure $\mathfrak{A}$, any formula of LFP is equivalent to a formula of $L^\omega_{\infty\omega}$. Indeed, the least fixed point of φ is equivalent to $\bigvee_{\alpha<\beta} \varphi_\alpha$, where β is the closure ordinal of φ on $\mathfrak{A}$. By the same argument,

[1] Or more accurately, database theory.

if $\mathcal{C}$ is a class of structures such that the closure ordinal of φ is bounded by some fixed ordinal β on all structures in $\mathcal{C}$, then φ is equivalent, over $\mathcal{C}$ to a formula of $L^{\omega}_{\infty\omega}$. Thus, in particular when $\mathcal{C}$ is the class of all finite structures (or indeed any class of structures with bounded cardinality), then every formula of LFP is equivalent over $\mathcal{C}$ to a formula of $L^{\omega}_{\infty\omega}$. This was pointed out by Kolaitis and Vardi in [41].

A variant of fixed point logic, known as PFP (for partial fixed point logic) was introduced in [3]. Here, we allow the formation of formulas $\mathbf{pfp}_{R,\bar{x}}\varphi(\bar{t})$, even when φ is not R-positive. Since φ may not be monotone, it may not have a least fixed point (or any fixed point, for that matter). Still, the sequence of stages $\varphi_\alpha, \alpha < \omega$ is well-defined and may or may not converge to fixed point. The formula $\mathbf{pfp}_{R,\bar{x}}\varphi(\bar{t})$ is true in a structure $\mathfrak{A}$ if the sequence $\varphi_\alpha, \alpha < \omega$ converges to a fixed point and the tuple denoted by $\bar{t}$ is in that fixed point, and it is false otherwise.

The importance of the fixed point logics LFP and PFP in finite model theory is partly accounted for by the results that relate definability in these logics to computational complexity classes. In particular, consider structures where the vocabulary contains a distinguished binary relation $<$ which is interpreted in every structure as a linear order of the universe. On such structures, the fixed point logics correspond exactly to natural complexity classes:

Theorem 1 (Immerman, Vardi [36, 60]) *On ordered structures* LFP = P.

Theorem 2 (Abiteboul-Vianu [3]) *On ordered structures* PFP = PSPACE.

While these equivalences hold only on ordered structures, Abiteboul and Vianu also proved [5] the remarkable result that the logics LFP and PFP have equal expressive power *on the class of all finite structures* if, and only if, P=PSPACE. We return to this topic in Section 4.

Kolaitis and Vardi [41] noted that as with LFP, on the class of finite structures, every formula of PFP is equivalent to a formula of $L^{\omega}_{\infty\omega}$.

2 Pebble Games and Their Uses

The equivalence relation $\equiv^k$ has an elegant characterisation in terms of Ehrenfeucht-Fraïssé style two player games, essentially given by Barwise [10]. The game board consists of two structures $\mathfrak{A}$ and $\mathfrak{B}$ and a supply of k pairs of pebbles $(a_i, b_i), 1 \leq i \leq k$. The pebbles $a_1, \ldots, a_l$ are initially placed on the elements of an l-tuple s of elements in $\mathfrak{A}$, and the pebbles $b_1, \ldots, b_l$ on a tuple t in $\mathfrak{B}$. There are two players, Spoiler and Duplicator[2]. At each move of the game, Spoiler picks up a pebble (either an unused pebble or one that is already on the board) and places it on an element of the corresponding structure. For instance he[3] might take pebble b_i and place it on an element of $\mathfrak{B}$. Duplicator must respond by placing the other pebble of the pair in the other structure. In the above example, she must place a_i on an element of $\mathfrak{A}$. If at the end of the move the partial map $f : \mathfrak{A} \to \mathfrak{B}$ given by $a_i \mapsto b_i$ is not a partial isomorphism, then Spoiler has won the game, otherwise it can continue for another move.

The result that links this game with the equivalence relation $\equiv^k$ is the following:

[2] Also variously called Player I and Player II, or $\forall$ and $\exists$, or even Abelard and Eloise (see [34]).

[3] By convention, Spoiler is male and Duplicator female.

If Duplicator has a strategy for playing the k-pebble game for p moves starting with the position $(\mathfrak{A}, s), (\mathfrak{B}, t)$, then $(\mathfrak{A}, s)$ and $(\mathfrak{B}, t)$ cannot be distinguished by any formula of L^k with quantifier rank p or less.

This characterisation was formulated in [10] in terms of sequences of sets of partial isomorphisms. These sequences can be transfinite, and they characterise equivalence in $L^k_{\infty\omega}$ up to a transfinite quantifier rank. If $\mathfrak{A}$ and $\mathfrak{B}$ are finite, then we only need to be concerned in the above with finite ordinals p. To see this suppose Duplicator has a winning strategy for p moves with initial position $(\mathfrak{A}, s), (\mathfrak{B}, t)$, for all finite p. Choosing p large enough (greater than n^k, where n is the larger of the cardinalities of A and B), we see that in such a game, some position must be repeated. Thus, Duplicator, in fact, has a strategy for playing the game indefinitely from this initial position, and therefore $(\mathfrak{A}, s) \equiv^k_{\infty\omega} (\mathfrak{B}, t)$. We therefore have the following:

Theorem 3 (Kolaitis-Vardi [41]) *For finite structures $\mathfrak{A}$ and $\mathfrak{B}$, the following are equivalent:*

- $(\mathfrak{A}, s) \equiv^k (\mathfrak{B}, t)$*; and*
- $(\mathfrak{A}, s) \equiv^k_{\infty\omega} (\mathfrak{B}, t)$.

It follows from this that a class of finite structures is closed under $\equiv^k$ if, and only if, it is definable by a sentence of $L^k_{\infty\omega}$. Indeed, in the context of finite structures, we rarely, if ever, need to consider infinitary logic as such. All the results can be formulated equally well in terms of the finitary k-variable logics. To say of a class of structures (or indeed, a global relation) that it is definable in $L^k_{\infty\omega}$ is the same as saying it is closed under $\equiv^k$, and to say that it is definable in $L^\omega_{\infty\omega}$ is to say that it is closed under $\equiv^k$ for some k. We adopt this practice for the remainder of this paper. However, where the result presented was originally formulated in terms of the infinitary logic, that formulation is usually given.

The k-pebble game was also independently presented by Immerman [35] and Poizat [51]. The former presented it explicitly in terms of a two-player game with pebbles, the latter as what he calls "*le va-et-vient de Fraïssé*". Both of them are concerned only with finite structures, and do not explicitly consider infinitary logic or transfinite quantifier rank.

A comparison of the pebble game with the Ehrenfeucht-Fraïssé game for first order logic is instructive. When we apply a game to show that some property P is not expressible in first order logic, we prove a statement of the form:

$$\forall k\, \forall p\, \exists \mathfrak{A}, \mathcal{B}\ (P(\mathfrak{A}) \wedge \neg P(\mathcal{B}) \wedge \mathfrak{A} \equiv^k_p \mathcal{B}).$$

That is, for every k and p, we can find two structures, one of which has the property P and the other does not, yet the two structures cannot be distinguished with only k variables and quantifier rank p. Since any sentence of quantifier rank p can be written with no more than p distinct variables, we can drop the first quantifier, and focus solely on quantifier rank. This is the usual Ehrenfeucht-Fraïssé game.

On the other hand, to show that P is not in $L^\omega_{\infty\omega}$, we prove the following statement:

$$\forall k\, \exists \mathfrak{A}, \mathcal{B}\, \forall p\ (P(\mathfrak{A}) \wedge \neg P(\mathcal{B}) \wedge \mathfrak{A} \equiv^k_p \mathcal{B}).$$

That is, we seek to find two structures, depending only on k which are k-equivalent for all p. Since this simply means that the Duplicator can play an infinite game, the parameter p plays no part, and we can focus solely on the number of variables. Thus, the number of variables becomes a crucial measure when we are dealing with formulas of unbounded quantifier rank, such as the formulas of fixed point logic.

Example 4 As a first example, consider two linear orders L_m and L_n where $n > m$. Spoiler wins the two pebble game on these two structures by placing the first pebble on the first element of L_n, the second on the second element, then moving the first pebble to the third element, and so marching down the linear order. Duplicator will run out of elements in L_m before Spoiler does.

In fact, using the strategy in the above example, one can show that on any class of linearly ordered structures, the expressive power of $L^{\omega}_{\infty\omega}$ is complete, i.e. equivalent to $L_{\infty\omega}$ (see [14]).

Example 5 Consider two structures over the empty signature, one of size k and one of size $k+1$. It is easily seen that Duplicator wins the k pebble infinite game on these two structures. The strategy is simply to respond to any pebble placed on a currently unpebbled element with a pebble on an unpebbled element of the other structure. Similarly a pebble placed on an element already covered by a pebble a is matched by placing a pebble on the element covered by the pebble matching a. As there are no relations on the two structures, any map is a partial isomorphism as long as it preserves equality. Furthermore, as there are at least k elements in both structures, and only k pairs of pebbles, there is always an unpebbled element available.

Since one of these structures has an even number of elements and the other one has an odd number of elements, we conclude that evenness is not closed under $\equiv^k$, for any k, (and hence is not definable, even in $L^{\omega}_{\infty\omega}$).

2.1 Extension Axioms

In order to introduce the next application of the pebble game, we first need some definitions.

An atomic type

$$\tau(x_1, \dots, x_m)$$

is the conjunction of a maximally consistent set of atomic and negated atomic formulas. In other words, τ is a complete description of an m-tuple with respect to the relations in σ. Note, that for a fixed m and σ, there is some fixed number of atomic types.

Let $\tau(x_1, \dots, x_m)$ and $\tau'(x_1, \dots, x_{m+1})$ be two atomic types such that τ' is consistent with τ. The τ, τ'-extension axiom, $\eta_{\tau,\tau'}$ is the sentence:

$$\forall x_1 \dots \forall x_m \exists x_{m+1} (\tau \rightarrow \tau').$$

That is, $\eta_{\tau,\tau'}$ asserts that every tuple of type τ can be extended to a tuple of type τ'.

Fagin [25] proved that every extension axiom has asymptotic probability 1 on the class of finite structures. The asymptotic probability of a sentence φ is defined as the limit $\lim_{n\to\infty} \frac{|\mathrm{Mod}(\varphi)\cap \mathcal{S}_n|}{|\mathcal{S}_n|}$, where $\mathcal{S}_n$ is the collection of all σ-structures with universe

$\{0, \ldots, n-1\}$. Fagin used this to show that first order logic has a 0–1 law, i.e. every first order sentence has an asymptotic probability, and it is 0 or 1 (this result was first proved in [27]). Fagin's proof relies on the fact that the collection of all extension axioms forms a complete first order theory. As is shown in the following example, the collection of extension axioms that contain at most k variables in fact forms a complete theory in L^k.

Example 6 Let θ_k be the set of all extension axioms $\eta_{\tau,\tau'}$ such that τ' has only k variables. Since θ_k is a finite set, we also write θ_k for the sentence of L^k that is the conjunction of the set. If $\mathfrak{A}$ and $\mathfrak{B}$ are such that $\mathfrak{A} \models \theta_k$ and $\mathfrak{B} \models \theta_k$, then it is clear that the extension axioms guarantee a winning strategy for Duplicator in the pebble game, since they state that, whatever the current position, for any way of modifying it, an appropriate element can be found witnessing that modification, in both $\mathfrak{A}$ and $\mathfrak{B}$. Thus, if $\mathfrak{A} \models \theta_k$ and $\mathfrak{B} \models \theta_k$, then $\mathfrak{A} \equiv^k \mathfrak{B}$.

Moreover, since θ_k is the conjunction of finitely many extension axioms, and each one of these has asymptotic probability 1, it follows that θ_k itself has asymptotic probability 1. Since, by Example 6, the models of θ_k form one $\equiv^k$ equivalence class, it follows that for any sentence φ of $L^k_{\infty\omega}$, either all models of θ_k are models of φ, or none are. In the former case, φ has asymptotic probability 1, and in the latter case it has asymptotic probability 0. We thus obtain:

Theorem 7 (Kolaitis-Vardi [41]) *$L^{\omega}_{\infty\omega}$ has a 0–1 law.*

This subsumes earlier results on 0–1 laws for first order logic [27, 25], LFP [11] and PFP [39].

Theorem 7 is a 0–1 law under what is called the uniform probability measure. That is, each structure of size n is assigned equal probability. A variety of other ways of assigning probabilities to structures are obtained, particularly in the case of graphs, by varying *edge probability*, usually as a function of n. The behaviour of $L^{\omega}_{\infty\omega}$ under such variable probability measures has been extensively studied (see [45, 46, 44, 59]). A detailed discussion of this study is omitted from the present survey, as it would take us too far afield.

2.2 Graph Properties

That the conjunction θ_k of all k-variable extension axioms defines a unique $\equiv^k$ equivalence class, as seen in Example 6, was shown by Immerman [35] who used this construction to demonstrate that several natural properties of graphs are not closed under $\equiv^k$ for any k. In particular, he showed that the property of a graph containing a $k+1$-clique, and the property of a graph containing a Hamiltonian cycle are both not closed under $\equiv^k$. As it turns out, we can give rather simpler proofs of both these facts, as direct applications of the k-pebble game, without reference to the extension axioms[4].

Example 8 Consider the graph K_k, the complete graph on k vertices, and K_{k+1}, the complete graph on $k+1$ vertices. It is clear that Duplicator can play an infinite game on

[4] For another proof that Hamiltonicity is not definable with finitely many variables, also using pebble games, see [21].

these two structures (the situation is completely analogous to Example 5). Since K_{k+1} contains a $k+1$ clique, and K_k doesn't, it follows that the property of containing a $k+1$-clique is not closed under $\equiv^k$.

Example 9 Consider the graphs $K_{k,k}$, the complete bipartite graph on two sets of k vertices, and $K_{k,k+1}$, the complete bipartite graph between a set of k vertices and a set of $k+1$ vertices. That is, in each of these two structures the set of vertices is partitioned into two parts (we call them *left* and *right* for ease of reference) such that there are no edges between any two vertices in the same part, and there is an edge between any two vertices in different parts.

It is again straightforward to show that Duplicator has a winning strategy in the k-pebble game played on these two structures. Duplicator matches any move by Spoiler on the left part of either structure with a move on the left of the other structure, and similarly for the right parts, preserving equalities among pebbled points within the parts. It can be verified that this is all that is required.

Finally, we note that, a complete bipartite graph has a Hamiltonian cycle if, and only if, the two parts have equal cardinality. Thus, $K_{k,k}$ is Hamiltonian, and $K_{k,k+1}$ is not.

Immerman also used the extension axioms to show that the problem of Graph Isomorphism is not axiomatisable with a finite number of variables. The problem of graph isomorphism is defined as the class of structures that are disjoint unions of two isomorphic graphs.

More formally, for any two σ structures $\mathfrak{A}$ and $\mathfrak{B}$, we define the disjoint union $\mathfrak{A} \oplus \mathfrak{B}$ as a structure in the vocabulary σ enriched with a new unary relation U, such that the σ substructure induced by the elements in U is $\mathfrak{A}$, the σ substructure of $\mathfrak{A} \oplus \mathfrak{B}$ induced by the elements not in U is $\mathfrak{B}$, and no relation contains a tuple containing both elements of U and its complement.

A general argument using the k-pebble game can be formulated to show the following fact:

Lemma 10 *If* $\mathfrak{A}_1 \equiv^k \mathfrak{A}_2$ *and* $\mathfrak{B}_1 \equiv^k \mathfrak{B}_2$*, then* $\mathfrak{A}_1 \oplus \mathfrak{B}_1 \equiv^k \mathfrak{A}_2 \oplus \mathfrak{B}_2$.

Proof (*sketch*): Duplicator's strategy is simply to match every move by Spoiler in $\mathfrak{A}_1$ by a move in $\mathfrak{B}_1$ using her strategy on that pair of structures and similarly for $\mathfrak{A}_2$ and B_2.

This observation was first recorded by Kolaitis [38].

Example 11 To see that the graph isomorphism problem is not axiomatisable with k variables, we choose two non-isomorphic graphs G_1 and G_2 such that $G_1 \equiv^k G_2$. That such pairs of graphs exist has been established in the foregoing examples. We now note, by Lemma 10, that $G_1 \oplus G_2 \equiv^k G_1 \oplus G_1$.

Finally, as an example that does illustrate the kind of technique used by Immerman, we present a proof from [14] that the class of rigid graphs (graphs that have no non-trivial automorphisms) is not axiomatised with a finite number of variables.

Example 12 It is known that the class of rigid graphs has asymptotic probability 1, (see [12]), as does the class of connected graphs[5]. Since the intersection of any two classes of

[5] This latter fact is easily deduced from the fact that all extension axioms have asymptotic probability 1.

asymptotic probability 1 itself has probability 1, we know we can find two non-isomorphic graphs G_1 and G_2 that are both rigid, connected and models of θ_k. Indeed, almost all graphs satisfy these properties. Now, it is easily seen that $G_1 \oplus G_2$ is rigid but $G_1 \oplus G_1$ is not. Still, $G_1 \oplus G_2 \equiv^k G_1 \oplus G_1$.

2.3 Classification

While the main concern of Immerman in [35] was the definability of problems in NP with a finite number of variables, Poizat [51] was motivated by entirely different concerns to study definability in L^k on finite structures. His concern was with classification theory, and particularly the classification of theories that are complete (in L^k) and that have small models. The classification of theories is usually given *up to interpretation*. A structure $\mathfrak{A}$ is said to be *interpretable* in a structure $\mathfrak{B}$, if the two structures have the same universe, and every relation in $\mathfrak{A}$ is definable (by a formula of L^k) in $\mathfrak{B}$. Two theories T_1 and T_2 are equivalent up to interpretation if every model of T_1 is interpretable in a model of T_2 and *vice versa*. A flavour of the kind of result that Poizat proves is given by the following theorem, which we present without proof.

Theorem 13 *Up to interpretation, there is exactly one complete theory $T_1(k)$ having a model of size k and three complete theories $T_1(k), T_2(k)$ and $T_3(k)$ having models of size $k+1$ which also have larger models.*

Another notable result from [51] is the following, which has also been proved in [34].

Theorem 14 *The Beth definability theorem fails for L^k.*

In the conclusion of [51], Poizat poses three open questions:

1. If a complete L^k theory has two distinct finite models, does it have infinitely many?
2. If a complete L^k theory has sufficiently large finite models, does it have arbitrarily large ones?
3. Do any two L^k-equivalent structures have a common L^k-elementary extension?

Simon Thomas provides a negative answer to the first of these questions in [56] where he shows the following:

Theorem 15 (Thomas) *For each $k \geq 3$, there are complete theories in L^{2k-2} and L^{2k-1} having exactly $k+1$ models.*

Moreover, since all models of the theories constructed in this proof are of the same cardinality, this also provides a negative answer to the third question in the above list. In [57], Thomas also constructs a complete theory in L^3 which has models of distinct finite cardinalities, and still has only finitely many models.

A refutation (or confirmation) of Poizat's second question depends upon an appropriate interpretation of the phrase "sufficiently large". A straightforward reading of the question would be the following: Is there an n, depending only upon k, such that if a complete L^k theory has models larger than n, it has infinitely many models? However, a negative answer to this question is already obtained from an example appearing in [51]. Any linear order is the unique, up to isomorphism, finite model of its L^3 theory. We will consider a somewhat subtler interpretation of this question in Section 5.

3 McColm's Conjectures

In [48, 47], McColm initiated a study of the relationship between first order logic, LFP and $L^{\omega}_{\infty\omega}$ on classes of finite structures. The analysis was in terms of a notion he termed *proficiency.* In the following definition, we write $||\varphi||_{\mathfrak{A}}$ to denote the closure ordinal of φ on $\mathfrak{A}$ (see Section 1.3).

Definition 16 *A class $\mathcal{C}$ of structures is proficient, if there is some positive first order formula φ such that* $\sup(\{||\varphi||_{\mathfrak{A}} \mid \mathfrak{A} \in \mathcal{C}\}) \geq \omega$.

That is, $\mathcal{C}$ is proficient if there is some formula whose closure ordinal is unbounded among structures in $\mathcal{C}$.

McColm in [48] formulated two conjectures, which taken together state that the following three conditions are equivalent for any class of structures $\mathcal{C}$.

1. $\mathcal{C}$ is proficient;
2. There is a formula of LFP that is not equivalent to any formula of first order logic on $\mathcal{C}$.
3. There is a formula of $L^{\omega}_{\infty\omega}$ that is not equivalent to any formula of first order logic on $\mathcal{C}$.

It is clear that (2) implies (1). This is because if a first order formula φ is such that there is a bound m on the closure ordinal of φ on structures in $\mathcal{C}$, then the least fixed point of φ is equivalent to the formula φ_m on $\mathcal{C}$, and φ_m is itself equivalent to a first order formula. McColm showed that (1) implies (3) and Kolaitis and Vardi [40] showed that (3) implies (1), establishing the equivalence of (1) and (3), which was the second of McColm's two conjectures. A proof of this is outlined in Section 3.1 below. The first of McColm's conjectures, the equivalence of (1) and (2) above was refuted by Gurevich *et al.* [32]. Some questions arising from it are discussed in Section 3.2.

3.1 Compactness

The equivalence of (1) and (3) is perhaps best understood in terms of the number of L^k-types that are realised among structures in $\mathcal{C}$. The following definition is from [14, 20].

Definition 17 *We say that a class of structures $\mathcal{C}$ is k-compact if the set $\{Type^k(\mathfrak{A}, s) \mid \mathfrak{A} \in \mathcal{C}, s \in A^k\}$ is finite.*

It turns out that a class $\mathcal{C}$ is k-compact (for all k) if, and only if, it is not proficient. We will see why this is the case through a series of facts which in themselves provide some insight into the behaviour of the equivalence relation $\equiv^k$ on finite structures. The first of these is a theorem proved in [19] (see also [51]):

Theorem 18 *For every finite structure $\mathfrak{A}$ and tuple $s \in A^k$, there is a formula $\varphi \in Type^k(\mathfrak{A}, s)$ such that, if $\mathfrak{B} \models \varphi[t]$, then $(\mathfrak{A}, s) \equiv^k (\mathfrak{B}, t)$.*

In other words each L^k-type that is realised in a finite structure is in fact determined by one of its elements. What's more, and this is the second significant fact, the quantifier rank of the formula φ in Theorem 18 can be bounded (up to an additive constant depending only on k) by the number of distinct types that are realised in $\mathfrak{A}$. The number of types realised in a structure is an invariant of the L^k theory of the structure of sufficient importance that we now give it a name:

Definition 19 *For any structure $\mathfrak{A}$, we write k-size$(\mathfrak{A})$ to denote the index of the equivalence relation $\equiv^k$ on A^k.*

One consequence of Theorem 18 is that if $\mathfrak{A} \equiv^k \mathfrak{B}$, then $\mathfrak{A}$ and $\mathfrak{B}$ realise exactly the same collection of L^k-types, and in particular they also have the same k-size, justifying our earlier statement that this is an invariant of the L^k theory of a structure. We will often refer to the k-size of a theory. Another consequence is that any complete L^k theory which has finite models is finitely axiomatised in L^k.

Our next observation is that there are, up to logical equivalence, only finitely many formulas of L^k of a fixed quantifier rank q. This is proved by an easy induction on q. Combining this with the fact that every complete L^k theory is axiomatised by a sentence whose quantifier rank is determined by its k-size, we see that for any fixed number n, there are up to $\equiv^k$, only finitely many structures with k-size at most n.

We can thus see that the following two statements are equivalent for any class of structures $\mathcal{C}$:

- $\mathcal{C}$ is k-compact.
- There is a bound n such that for all $\mathfrak{A} \in \mathcal{C}$, k-size$(\mathfrak{A}) \leq n$.

To establish the connection with proficiency, we now need to note that if φ is a formula such that all stages φ_m are equivalent to formulas of L^k, then the closure ordinal of φ on $\mathfrak{A}$ is bounded by k-size$(\mathfrak{A})$. It follows that if $\mathcal{C}$ is k-compact for all k it is not proficient. The converse of this also holds. For this, we require the following additional result.

Theorem 20 ([19, 40]) *There is an R-positive first order formula $\zeta(R, \bar{x}, \bar{y})$ such that:*

- *$\mathfrak{A} \models \mathbf{lfp}_{R,\bar{x}\bar{y}}\zeta[s, t]$ if, and only if, $(\mathfrak{A}, s) \not\equiv^k (\mathfrak{B}, t)$; and*
- *$\mathfrak{A} \models \zeta_m[s, t]$ if, and only if, s and t are distinguished in $\mathfrak{A}$ by a formula of L^k of quantifier rank m.*

That is, there is a fixed point formula that defines *uniformly* the inequivalence $\not\equiv^k$, and therefore there is also an LFP formula defining $\equiv^k$. In fact, the formula ζ is nothing more than a definition of the winning positions for Spoiler in the k-pebble game, which are quite naturally described inductively. This is the reason why the stages of the induction of ζ correspond to definitions of the inequivalence relations $\not\equiv^k_m$.

Now, suppose $\mathcal{C}$ is a class that is not proficient. That means that for every formula, and therefore ζ in particular, there is an m such that the least fixed point of ζ is equivalent to ζ_m. By Theorem 20, this implies that the relation $\equiv^k$ is the same as $\equiv^k_m$. Since the latter equivalence relation is always of finite index as there are only finitely many L^k formulas of quantifier rank m, this means that $\mathcal{C}$ is k-compact.

Finally, to see that proficiency is equivalent to condition (3), we first note that if $\mathcal{C}$ is proficient, then by the above argument there are infinitely many distinct L^k types realised among structures in $\mathcal{C}$. Each of these types is definable by a formula of L^k and the uncountably many infinite disjunctions that can be formed are all definable in $L^k_{\infty\omega}$. Since there are only countably many first order formulas, not all formulas of $L^k_{\infty\omega}$ can be equivalent to first order formulas. For the converse, if $\mathcal{C}$ is not proficient, then it is k-compact, and therefore the equivalence relation $\equiv^k$ is the same as $\equiv^k_m$ for some m. Thus, every formula of $L^k_{\infty\omega}$ is equivalent to a finite disjunction of formulas of quantifier rank m.

3.2 LFP on Proficient Classes

As mentioned above, McColm's first conjecture, the equivalence of proficiency with the existence of formulas of LFP that are not equivalent to any first order formula was refuted by Gurevich *et al.* [32]. They construct two distinct classes of structures on which every LFP formula is equivalent to a first order formula but which are nonetheless proficient.

The resolution of McColm's conjectures, one positively and the other negatively still leaves a number of questions open. In particular, Kolaitis and Vardi [43] asked whether similar conjectures could be formulated for a fixed number of variables k. They show that McColm's second conjecture (whose proof is outlined above) can be refined to the following statement: on any class of structures $\mathcal{C}$ and for any k, every formula of $L^k_{\infty\omega}$ is equivalent to a formula of L^k if, and only if, no k variable induction on $\mathcal{C}$ is unbounded. To make this precise, we have to specify what a k variable induction is. It is not sufficient to define these as the inductions defined by positive formulas of L^k. Rather, as Kolaitis and Vardi show, we need to consider *systems of inductive formulas* having no more than k variables all together. They call the logic thus formed LFP^k.

With this formulation, there is a version of McColm's first conjecture that remains open, namely are the following two conditions equivalent for any class of structures $\mathcal{C}$:

1. every formula of $L^k_{\infty\omega}$ is equivalent to a formula of L^k; and

2. every formula of LFP^k is equivalent to a formula of L^k?

This is not refuted by the examples in [32], which are classes of structures where the first condition fails, and every formula of LFP is equivalent to a first order formula, but one with possibly many more variables.

Another question related to McColm's first conjecture that remains open is whether it holds of classes of structures $\mathcal{C}$ where structures are linearly ordered. This form of the question was posed in [40]. It was argued in [20] that a resolution of the question either way would solve outstanding complexity theoretic questions. In particular, if it can be shown that there is an infinite class of ordered structures on which every formula of LFP is equivalent to a first order formula, then P$\neq$PSPACE (this was proved in [17]). On the other hand, if every proficient class of ordered structures admits an LFP formula that is not equivalent to any first order formula, then this is true in particular of the class of structures of the form $(n, +, \times)$, i.e. those structures whose universe is a finite initial segment of the natural numbers equipped with ternary relations for addition and multiplication. It is shown that the separation of LFP from first order logic on this class would imply the separation of the linear time hierarchy from the complexity class

ETIME. The question of whether LFP is more expressive than first order logic on the class of structures $(n, +, \times)$ was also posed as an open question in [32] in an equivalent form about structures $(n, <, \mathrm{BIT})$.

It is also shown in [20] that Kolaitis and Vardi's ordered conjecture would be proved if, for every signature σ, there is a k such that every sentence of first order logic is equivalent, over ordered σ-structures, to a formula of L^k. For unary signatures σ this is true, as in this case every first order sentence is equivalent to one of L^3. While it seems extremely unlikely for other signatures, no counterexample has so far been found. Note that it is the case for all σ that there is a k such that every formula of $L^\omega_{\infty\omega}$ is equivalent over ordered σ structures to a formula of $L^k_{\infty\omega}$ (see Example 4).

It also remains an open question whether or not every first order sentence is equivalent to one in L^3 on the class of structures $(n, +, \times)$. On the classical structure of arithmetic $(\mathbf{N}, +, \times)$ every first order sentence is, in fact, equivalent to a sentence with just three variables. This can be proved using an arithmetic pairing function.

4 Relational Complexity

Abiteboul and Vianu proved the remarkable result that the logics LFP and PFP are equivalent in expressive power on the class of finite structures if, and only if, the complexity classes P and PSPACE coincide:

Theorem 21 (Abiteboul-Vianu [4, 5]) LFP = PFP *if, and only if,* P = PSPACE*.*

This holds despite the fact that the two logics capture the respective complexity classes only on ordered structures, and are otherwise weaker. The result, first appearing in [4] was proved through establishing a normal form for these logics using *loosely coupled generic machines*. This result was later re-analysed in [5, 19] in terms of finite variable logics.

One direction of Theorem 21 follows immediately from Theorems 1 and 2. That is, if the two logics LFP and PFP are equally expressive then, in particular, they are equally expressive on ordered structures, where their expressive power is exactly the respective complexity classes P and PSPACE. Thus, the two complexity classes must coincide. It is in proving the other direction that finite variable logics come into play.

4.1 Invariants

We can define for every structure $\mathfrak{A}$ and every k, a structure $I_k(\mathfrak{A})$ whose universe is the set of $\equiv^k$ equivalence classes of A^k. The structure $I_k(\mathfrak{A})$ is equipped with a linear order on its universe and a variety of other relations that encode the structure of $\mathfrak{A}$. To be precise, suppose $\mathfrak{A}$ is a structure in the signature σ, and $\tau_1, \ldots, \tau_m$ is an enumeration of all atomic types in k variables in the signature σ (recall Section 2.1). We first define the structure $I_k(\mathfrak{A})^-$ as follows:

$$I_k(\mathfrak{A})^- = (A^k/\equiv^k, T_1, \ldots, T_m, X_1, \ldots, X_k)$$

where,

- the universe is $A^k/\equiv^k$;

- for $1 \leq j \leq m$, $T_j([s])$ if, and only if, $\mathfrak{A} \models \tau_j[s]$; and
- for $1 \leq i \leq k$, $X_i([s],[t])$ if, and only if, s and t differ *at most* on their ith element.

It turns out that a structure $I_k(\mathfrak{A}) = (I_k(\mathfrak{A})^-, <)$ which is an expansion of $I_k(\mathfrak{A})^-$ with a linear order is interpretable inside $\mathfrak{A}$ uniformly, using formulas of LFP. That is, there are formulas of LFP, one for every relation symbol in the signature of $I_k(\mathfrak{A})$ which define the corresponding relations in every structure $\mathfrak{A}$. We have already seen as a consequence of Theorem 20 that there is such a definition of the relation $\equiv^k$. A rather more involved construction can be used to show that there is, in fact, a formula of LFP that uniformly orders these equivalence classes. The other relations above are easily first order definable.

One consequence of the LFP interpretability of $I_k(\mathfrak{A})$ in $\mathfrak{A}$ is that every LFP formula φ in the signature of $I_k(\mathfrak{A})$ can be translated into an LFP formula φ' in the signature σ such that $I_k(\mathfrak{A}) \models \varphi$ if, and only if, $\mathfrak{A} \models \varphi'$. This translation is carried out by the simple substitution of the LFP formulas defining the relations in $I_k(\mathfrak{A})$ for the corresponding relation symbols. Thus, a similar translation also yields a PFP formula in the signature σ for every PFP formula on the interpreted structures.

The relations X_i are included in the definition of $I_k(\mathfrak{A})$ to allow for a translation in the other direction. That is, we can take any σ-formula in L^k and translate it into a corresponding formula on the structures $I_k(\mathfrak{A})$. This translation also extends to formulas of LFP and PFP with fewer than k variables.

4.2 Complexity

Now, suppose that P=PSPACE and let φ be an arbitrary PFP formula. Choosing an appropriate k, we can translate φ into a PFP formula φ' on the structures $I_k(\mathfrak{A})$. However, since these are ordered structures, by hypothesis, on this class of structures every PFP formula is equivalent to a formula of LFP. This LFP formula ψ can now be translated back into an LFP formula ψ^* in the signature σ, which is by construction equivalent to φ on all structures. This establishes Theorem 21.

This result shows that, even though the fixed point logics do not themselves capture the corresponding complexity classes, their equivalence turns out to be equivalent to an open complexity theoretic question. This leads to a view of these logics as defining complexity classes on a *relational machine* (defined in [2]), where we measure complexity not in terms of the size of a structure, but in terms of its k-size.

A relational machine is a Turing-like machine which is presented an input not on a tape, but a relational store. This relational store is accessed by the machine through first order queries whose results are stored in other relational registers. It can be shown that the class of structures accepted by a relational machine is always closed under $\equiv^k$ for some k. The value of k may be determined by examining the machine, as it is no greater than the largest number of variables occurring in any formula used to query the relational store. Conversely, any recursive set of structures that is closed under $\equiv^k$ is accepted by some relational machine. The power of resource bounded versions of the machine then corresponds to the expressive power of various fixed point logics. Thus, polynomial time and polynomial space bounded relational machines accept exactly the classes of structures that are definable in LFP and PFP respectively. Here we measure polynomial time and space not in terms of the size of the input structure but in terms of its k-size, since a relational machine is not even capable of determining the size of the input structure. Indeed, for any complexity class $\mathcal{C}$ that includes P, it can be shown that a class of structures D is in the relational version of $\mathcal{C}$ if, and only if, the class of ordered structures $\{I_k(\mathfrak{A}) \mid \mathfrak{A} \in D\}$ is in $\mathcal{C}$. For instance, a polynomial time relational machine is equivalent to a a polynomial time machine that receives as input, instead of the encoding of its input structure $\mathfrak{A}$, an encoding of the invariant $I_k(\mathfrak{A})$, where k depends only on the machine. Note that the size of the structure $I_k(\mathfrak{A})$ is k-size($\mathfrak{A}$). Abiteboul *et al.* give characterisations of a variety of relational complexity classes ranging from P to EXP in terms of fixed point logics, including non-deterministic fixed point and alternating fixed point logic.

This line of investigation has shown that a large number of complexity theoretic questions can be reproduced as questions about fragments of the infinitary logic $L^\omega_{\infty\omega}$ (see [2, 15, 18]). Below, we exhibit one such set of results which concerns an effective fragment of $L^\omega_{\infty\omega}$ that is not obtained by fixed point logics, but by a restricted form of second order quantification. This provides a characterisation of relational complexity classes closer in spirit to Fagin's result concerning NP.

The following definition is from [15]:

Definition 22

- *For an l-ary relation symbol R, and $k \geq l$, we define the second order quantifier $\exists^k R$ to have the following semantics: $\mathfrak{A} \models \exists^k R\,\varphi$ if there is an $X \subseteq A^l$ such that X is closed under the equivalence relation $\equiv^k$ in $\mathfrak{A}$, and $(\mathfrak{A}, X) \models \varphi$. As usual, $\forall^k R$ abbreviates $\neg\exists^k R\neg$.*
- *$\Sigma^{1,\omega}_1$ denotes the class of formulas of the form $\exists^{k_1} R_1 \ldots \exists^{k_m} R_m \varphi$, where φ is first order.*
- *$\Pi^{1,\omega}_1$ denotes the class of formulas of the form $\forall^{k_1} R_1 \ldots \forall^{k_m} R_m \varphi$, where φ is first order.*
- *$\Sigma^{1,\omega}_{n+1}$ denotes the class of formulas of the form $\exists^{k_1} R_1 \ldots \exists^{k_m} R_m \varphi$, where φ is $\Pi^{1,\omega}_n$.*
- *$\Pi^{1,\omega}_{n+1}$ denotes the class of formulas of the form $\forall^{k_1} R_1 \ldots \forall^{k_m} R_m \varphi$, where φ is $\Sigma^{1,\omega}_n$.*
- $\mathrm{SO}^\omega = \bigcup_{n \in \omega} \Sigma^{1,\omega}_n$.

On ordered structures, the existential fragment $\Sigma^{1,\omega}_1$ captures NP. This is, because if $\mathfrak{A}$ is ordered, then every $X \subseteq A^l$ is $\equiv^k$-closed. The correspondence extends all the way up the polynomial hierarchy.

Theorem 23 *On ordered structures, for every $n \in \omega$, $\Sigma^{1,\omega}_n = \Sigma^p_n$ and $\Pi^{1,\omega}_n = \Pi^p_n$.*

$\Sigma^{1,\omega}_1$ fails to capture NP on arbitrary structures because for every formula φ of this logic there is a k such that the global relation defined by φ is closed under $\equiv^k$, and we saw in Section 2 many examples of properties in NP that are not closed under any such relation. In fact, the following containments hold, in terms of expressive power:

$$\mathrm{LFP} \subseteq \Sigma^{1,\omega}_1 \subseteq \mathrm{PFP}$$

Moreover, by an argument similar to the one used to prove Theorem 21, we can establish the following [15]:

Theorem 24 LFP $= \Sigma_1^{1,\omega}$ *if, and only if,* P $=$ NP.

Theorem 25 $\Sigma_1^{1,\omega} =$ PFP *if, and only if,* NP $=$ PSPACE.

Analogous results hold for all levels of the polynomial hierarchy. Details can be found in [15].

The logic $\Sigma_1^{1,\omega}$ is equivalent in expressive power to the *non-deterministic fixed point logic* introduced in [2] specifically to establish results of the form presented in Theorem 24. The paper also introduced an *alternating fixed point logic* equivalent in expressive power to PFP. SO^ω is equivalent to the fragment of that logic given by a bounded number of alternations.

Theorem 21 may be compared with a result due to Tiuryn and Urzyczyn [58] to the effect that two different dynamic logics of programs are equivalent in expressive power if, and only if, P=PSPACE. The two logics in question are dynamic logics formed from program schemes with one stack (LFDS) and program schemes with arrays (LPDA). As with Theorem 21, one direction of the equivalence follows from capturing results on a specialised class of structures. Thus, LFDS expresses exactly the polynomial time computable properties of "Herbrand like structures" and a similar result holds for LFDA and polynomial space computation. Herbrand like structures are structures in which every element is named by a term. This requirement plays a role similar to that of the ordering required in Theorems 1 and 2 in that it eliminates all non-trivial automorphisms of the structure. Indeed, it is not difficult to see that there is an LFP-definable order on Herbrand like structures (as these structures are 1-rigid, see Section 6.1). The other direction of Tiuryn and Urzyczyn's result follows from the observation that formulas of the logics LFDS or LFDA have only limited access to a structure outside of the elements that are named by terms. Thus, the truth value of a formula φ on any structure $\mathfrak{A}$ is determined by the substructure of $\mathfrak{A}$ generated by the interpretation of the free variables of φ, and this is, naturally, a Herbrand like structure.

The results on relational complexity classes above have all been formulated for complexity classes including P. It is not known if similar results can be obtained for lower complexity classes. For instance, is there a logic $\mathcal{L}$ such that $\mathcal{L} =$ LFP if, and only if, L=P? It does not seem likely that the methods used above could be deployed to show this, as it seems unlikely that a formula defining $\equiv^k$ or a formula ordering the $\equiv^k$-classes could be constructed in a logic which can only express properties that are in L. Indeed, Grohe [29] has shown that the relation $\equiv^k$ is P-complete.

5 Relation between Size and k-size

Many of the results in the previous sections have established the importance of k-size as an invariant of a complete L^k theory. For any structure $\mathfrak{A}$, k-size($\mathfrak{A}$) is the size of the structure $I_k(\mathfrak{A})$. It is the measure used for determining relational complexity. It is also of crucial importance in the context of McColm's conjectures. A class of structures $\mathcal{C}$ is proficient in the sense of McColm if, and only if, there is no bound on the k-size of structures in $\mathcal{C}$.

While the results in Section 4 established that a variety of interesting complexity theoretic questions can be translated into questions about finite variable logics, i.e. logics in which definable classes are closed under the $\equiv^k$ equivalence relation, such translations have not necessarily thrown new light on the separation of complexity classes. One reason for this is that the main tool for establishing inexpressibility results for these logics is the pebble game. This can be used to show that a class of structures is not closed under any of the relations $\equiv^k$, and therefore not definable in any of the logics we have been considering. This means, however, that the pebble game cannot be used directly to separate the logics in question. Simply put, such games can be used to show that a class of structures is not definable in LFP, but it shows at the same time that it is not definable in PFP either.

There is, however, one method that has been used to show that certain problems that are closed under $\equiv^k$ are nonetheless not definable in LFP. This is the analysis of complexity in terms of k-size. Thus, a class of structures which is closed under $\equiv^k$, but whose complexity is not polynomial in the k-size of the structures cannot be definable in LFP. In [19], it was shown that there are polynomial time decidable classes of binary trees that are closed under $\equiv^4$ but are not definable in LFP. The crucial fact here is that the size of a binary tree is exponential in its height, but its k-size is only polynomial in its height.

The question naturally arises: how far can the size of structure and its k-size diverge, or more generally, how is the k-size of a structure related to its size? We can begin with a few straightforward observations. First of all, it is immediate from the definition that, if $|A| = n$, then $k\text{-size}(\mathfrak{A}) \leq n^k$, since there are at most n^k k-tuples in A. This bound is actually achieved in the class of linear orders, for all $k \geq 2$.

If we turn the question around and ask how large a structure can we find with a given k-size, at first things appear equally simple. It is trivial to construct arbitrarily large such structures. Thus, if we consider the class of structures in the empty signature (see Example 5), it is clear that the k-size of structures in this class is bounded by a constant, and yet there are arbitrarily large structures in this class. However, this class splits into only finitely many $\equiv^k$ equivalence classes of structures since it is k-compact (see Section 3.1). So, the existence of arbitrarily large structures of a fixed k-size follows simply from the existence of infinite $\equiv^k$ equivalence classes. The more interesting question is: how large can the *smallest* or *largest* structure in a $\equiv^k$ equivalence class be, with respect to its k-size?

To be precise, define the following two functions $d_k, e_k : \mathbf{N} \to \mathbf{N}$ by:

$$d_k(n) = \max\{\text{card}(\mathfrak{A}) \mid k\text{-size}(\mathfrak{A}) \leq n \text{ and } \mathfrak{A} \text{ is smallest in its } \equiv^k \text{ class.}\}$$

$$e_k(n) = \max\{\text{card}(\mathfrak{A}) \mid k\text{-size}(\mathfrak{A}) \leq n \text{ and } \mathfrak{A} \text{ belongs to a finite } \equiv^k \text{ class.}\}$$

These functions are well defined, since there are only finitely many $\equiv^k$ classes of structures with a fixed k-size. This latter fact follows from the fact that every complete L^k theory is axiomatised by an L^k sentence whose quantifier rank is determined by its k-size (see the discussion in Section 3.1, following Theorem 18). By the definition of the functions, we can make the following statements:

Proposition 26 *For every structure $\mathfrak{A}$, with k-size$(\mathfrak{A}) \leq n$,*

- *there is a structure $\mathfrak{B}$ with $\text{card}(\mathfrak{B}) \leq d_k(n)$ such that $\mathfrak{A} \equiv^k \mathfrak{B}$; and*
- *if $\text{card}(\mathfrak{A}) > e_k(n)$ then there are arbitrarily large structures $\mathfrak{B}$ such that $\mathfrak{A} \equiv^k \mathfrak{B}$.*

These can be seen as downward and upward Löwenheim-Skolem properties for finite structures. However, they raise the question of what kind of bounds can we put on the functions d_k and e_k? This is also related to the second of Poizat's questions (see Section 2.3). That is, the second part of Proposition 26 can be seen as an answer to the question as it shows one sense in which if a complete L^k theory has a sufficiently large model, it has arbitrarily large models.

As already indicated above, the class of binary trees provides exponential lower bounds on both the functions d_k and e_k. It was shown in [16] that this bound is in some sense optimal for the class of trees. That is, if the structures $\mathfrak{A}$ in the definitions above are restricted to be trees, there are exponential upper bounds on the functions d_k and e_k.

In [30], Grohe showed the following:

Theorem 27 *For all $k \geq 3$, there is no recursive function f such that $d_k(n) < f(n)$ for all n.*

Grohe's construction relies on a gadget called a *scale*, which is a structure, containing two distinguished sets of elements M and N, which have the same cardinality, but such that no non-trivial relation among them is definable in L^k. The equal cardinality condition is enforced by including in the structure elements and relations that encode bijections between the two sets. If a sufficiently large number of such bijections is included, then only trivial relations are definable.

The crucial step in the argument is that such scales can be constructed, even when the set M can be decomposed into two disjoint definable sets M_1 and M_2. This is an *addition scale*, as it can be seen as witnessing that the cardinality of the set N is the sum of the cardinalities of M_1 and M_2. Such scales are then used to construct structures whose L^k theory enforces constraints on cardinalities of sets, where the constraints are given by an arithmetic equation. Multiplication is easily implemented as a cartesian product. Such a construction can be used to obtain an encoding of Diophantine equations as finite structures, where the k-size of the structure depends only on the size of the equation, but the size of the structure depends on the size of a solution to the equation. Since there is no recursive upper bound on the solution of a Diophantine equation in terms of the size of the equation, Theorem 27 follows.

One consequence of this result is that there is no computable function f which is the inverse of the function I_k. That is, there is no computable function f such that $f(I_k(\mathfrak{A})) \equiv^k \mathfrak{A}$. This is related to a question posed in [19] asking whether the construction of the invariant I_k can be inverted in *time polynomial in the size of the output.* Note that this is not immediately ruled out by the non-existence of a function f as above. However, Grohe [28] has also shown that this is unlikely, as the following problem is NP-complete:

Input: A finite structure $\mathcal{J}$ and a string x.
Question: Is there a structure $\mathfrak{A}$ of size $\leq$ length(x) such that $I_k(\mathfrak{A}) = \mathcal{J}$.

This suggests that though the equivalence relation $\equiv^k$ is polynomial time decidable, it most likely does not admit polynomial time construction of canonical representatives.

6 Other Directions

The previous sections have explored in some detail a few of the research directions related to finite variable logics in finite model theory. This section takes a briefer look at directions not yet covered.

6.1 When is an Order Definable?

Since so many results related to fixed point logics hold specifically on ordered structures only — for instance the theorem of Immerman and Vardi identifying definability in LFP with computability in PTIME, an important question that arises is on what classes of structures is a linear order definable, even when it is not explicitly present. That is, what classes of structures $\mathcal{C}$ have the property that there is a formula $\lambda(x, y)$ (in a logic $\mathcal{L}$) such that the interpretation of λ on any structure in $\mathcal{C}$ is a linear order of the universe? If such a formula λ exists, we say that $\mathcal{C}$ admits an $\mathcal{L}$-definable linear order. It is easy to see that if this holds then every structure in $\mathcal{C}$ must be rigid, that is, it does not admit any automorphisms other than the trivial one.

If a structure $\mathfrak{A}$ has the property that no two elements of $\mathfrak{A}$ have the same L^k-type, we say that $\mathfrak{A}$ is k-rigid. This notion was introduced in [14, 19] where it was also shown that every finite rigid structure is k-rigid for some k. It was also established in [14] that a class $\mathcal{C}$ admits an LFP-definable linear order if, and only if, there is a k such that every structure in $\mathcal{C}$ is k-rigid. Indeed, if there is no such k, then $\mathcal{C}$ does not admit even an $L^{\omega}_{\infty\omega}$-definable linear order. The other direction, showing that if such a bound exists then $\mathcal{C}$ admits an LFP-definable linear order is based on the general result about ordering the equivalence classes of $\equiv^k$ in an arbitrary (not necessarily rigid) structure, which is the basis of Theorem 21 discussed in Section 4.s

Stolboushkin [55] showed that there is a *first order definable* class of structures which admits an LFP-definable linear order but no first order definable order. Gurevich and Shelah [33] constructed a first order definable class of structures such that all structures in the class are rigid, but it does not admit an $L^{\omega}_{\infty\omega}$-definable linear order.

6.2 Infinitary Logic and Formal Languages

As was mentioned above, any recursively enumerable class of structures $\mathcal{C}$ that is closed under $\equiv^k$ is accepted by a relational machine. Moreover, the complement of $\mathcal{C}$ is axiomatised by a *recursive* set of sentences of L^k. Thus $\mathcal{C}$ is recursive and closed under $\equiv^k$ if, and only if, both it and its complement are axiomatised by a recursive set of sentences of L^k. These equivalences were essentially established in [1], in terms of infinite disjunctions of sets of L^k formulas.

Abiteboul *et al.* [1] also showed a correspondence between definability in fixed point logics and axiomatisability by sets of sentences generated by regular languages. This is done by making an analogy between concatenation of strings and composition of formulas. Formally, we are given a finite set Σ of formulas of L^k, each with a distinguished second order variable. The composition $\varphi\psi$ of two formulas $\varphi(R)$ and $\psi(S)$ in Σ is defined as the formula obtained by replacing all occurrences of R in φ with ψ, renaming variables in ψ as necessary for each replacement. Given a finite alphabet Γ and a map h from Γ to Σ, h is extended in a natural way to a map from the set of strings Γ^* into the set of formulas of L^k so as to commute with the operations of concatenation on Γ^* and composition in L^k. We now call a set of sentences of L^k *regular* if it is the image under h of a regular subset of Γ^*. The following is the result proved by Abiteboul *et al.*

Theorem 28 (Abiteboul *et al.* [1]) *The following are equivalent for any class of structures $\mathcal{C}$:*

- *$\mathcal{C}$ is definable by a sentence of* PFP.
- *$\mathcal{C}$ is axiomatised by a regular set of sentences of L^k for some k.*

A similar characterisation of LFP is obtained by considering *one letter end marked languages.* These are regular languages of the form a^*b for some letters a and b. We call a set of sentences of L^k *positive one letter end marked* if it is the image of a one letter end marked language under a map h from $\{a, b\}$ to a pair of *positive* formulas of L^k.

Theorem 29 (Abiteboul *et al.* [1]) *The following are equivalent for any class of structures $\mathcal{C}$:*

- *$\mathcal{C}$ is definable by a sentence of* LFP.
- *$\mathcal{C}$ is axiomatised by a positive one letter end marked set of sentences of L^k for some k.*

These results should be compared with an earlier result due to Engeler [23] which characterised an iterative programming language in terms of formulas of $L_{\omega_1\omega}$. The formulas so obtained are infinite disjunctions of first order formulas, and the equivalence is established over any fixed single structure. Engeler also noted that only finitely many variables were required. The programming language considered by Engeler is a version of the *while* language. As has been shown by Abiteboul and Vianu [3], the logic PFP is equivalent to a relational version of the *while* language.

Another related result that deserves mention is one due to Immerman [37] who showed that *on ordered structures* the complexity class PSPACE is characterised by iterated first order formulas. The iteration is essentially similar to the composition of formulas in [1], and the iterated formulas therefore have a bound on the number of variables that may appear. Indeed, Immerman shows a tight correspondence between the bound on the number of variables required and the space complexity. In a sufficiently rich language[6] the properties definable by iterated first order formulas with $k + 1$ variables are exactly the properties in the complexity class DSPACE[n^k].

6.3 Existential Fragments

Just as the fixed point logics of Section 1.3 can be seen as fragments of $L^\omega_{\infty\omega}$ in terms of expressive power, so the database query language Datalog can be seen as a fragment of an existential version of $L^\omega_{\infty\omega}$. This observation, due to Kolaitis and Vardi [42] has been the basis of a line of investigation on existential finite variable logics.

Let $pL^k(\exists)$ denote the *positive existential* fragment of L^k. That is to say that $pL^k(\exists)$ contains those formulas of L^k that are obtained from atomic formulas through the operations of conjunction, disjunction and existential quantification. $L^k(\exists, \neq)$ is the extension of $pL^k(\exists)$ where we allow negation in front of atomic equalities only, and $L^k(\exists)$ is existential L^k, the extension of $L^k(\exists, \neq)$ obtained by allowing negations of arbitrary atomic

[6] Essentially one which besides the order on the structures also provides arithmetic operations.

formulas. The logics $pL^k_{\infty\omega}(\exists)$, $L^k_{\infty\omega}(\exists,\neq)$ and $L^k_{\infty\omega}(\exists)$ are the respective closures of these logics under infinitary conjunctions and disjunctions; $pL^\omega_{\infty\omega}(\exists)$, $L^\omega_{\infty\omega}(\exists,\neq)$ and $L^\omega_{\infty\omega}(\exists)$ are the respective unions of the latter over all k. It is clear that of all these, the logic $L^\omega_{\infty\omega}(\exists)$ is the strongest in that it subsumes all the others.

One important property of these logics is that all classes of structures definable in any of them are closed under extensions. That is, if φ is a sentence of $L^\omega_{\infty\omega}(\exists)$ and $\mathfrak{A} \models \varphi$ and $\mathfrak{A} \subseteq \mathfrak{B}$, then $\mathfrak{B} \models \varphi$. In addition, properties definable in $L^\omega_{\infty\omega}(\exists,\neq)$ are closed under one-to-one homomorphisms, and those definable in $pL^\omega_{\infty\omega}(\exists)$ are closed under all homomorphisms. Moreover, on ordered structures, these logics are complete: $L^\omega_{\infty\omega}(\exists)$ can express *all* properties closed under extensions, $L^\omega_{\infty\omega}(\exists,\neq)$ can express all properties closed under one-to-one homomorphisms and $pL^\omega_{\infty\omega}(\exists)$ all properties closed under homomorphisms (see [7]).

Datalog is a database language in which queries are defined through a series of Horn clauses, whose inductive closure is used to define the query. Datalog($\neq$) denotes the extension of Datalog in which inequalities are allowed to appear in the bodies of the Horn clauses, while Datalog($\neg,\neq$) is an extension in which negations are allowed on atomic formulas whose relation symbol does not appear in the head of any clause. These query languages can be seen as fixed point logics whose expressive power is bounded by the logics $pL^\omega_{\infty\omega}(\exists)$, $L^\omega_{\infty\omega}(\exists,\neq)$ and $L^\omega_{\infty\omega}(\exists)$ respectively. This fact is used by Afrati *et al.* [7] and Kolaitis and Vardi [42] to establish a number of inexpressibility results for these variants of Datalog. Some of these results (particularly those in [42]) are proved using a pebble game simlar to that in Section 2.

We write $\mathfrak{A} \preceq^k \mathfrak{B}$ to denote that for every sentence φ of $L^k(\exists)$, if $\mathfrak{A} \models \varphi$, then $\mathfrak{B} \models \varphi$. Note that this relation is not symmetric. The relation is characterised by a pebble game just like the one in Section 2 except that Spoiler is restricted to playing his moves in the structure $\mathfrak{A}$. Games for $pL^\omega_{\infty\omega}(\exists)$ (or $L^\omega_{\infty\omega}(\exists,\neq)$) are obtained by modifying the winning conditions in this game so that Duplicator is only required to maintian a partial homomorphism (respectively, one-to-one homomorphism) rather than a partial isomorphism. In [42], Kolaitis and Vardi use these games to show that a variety of fixed subgraph homomorphism problems are not definable in $L^\omega_{\infty\omega}(\exists,\neq)$. In particular consider the class of directed graphs with two distinguished pairs of vertices (s_1,t_1) and (s_2,t_2) such that there disjoint paths from s_1 to t_1 and from s_2 to t_2. It is shown that this class is not definable in $L^\omega_{\infty\omega}(\exists,\neq)$. It remains an open question whether this is definable in $L^\omega_{\infty\omega}$.

Rosen and Weinstein [53] considered preservation theorems for the existential finite variable logics. This study is motivated by the failure of the Łoś-Tarski preservation theorem when restricted to finite models. They showed that the existential preservation property fails for $L^\omega_{\infty\omega}$ both on finite structures and in general. That is, there is a sentence of $L^\omega_{\infty\omega}$ whose models (finite or infinite) are closed under extensions, but which is not equivalent (even on finite models) to a sentence of $L^\omega_{\infty\omega}(\exists)$. Another notable result is that unlike for L^k, the $L^k(\exists)$ theory of a finite structure is not necessarily finitely axiomatisable. That is, there is a finite structure $\mathfrak{A}$ such that for every sentence $\varphi \in L^k(\exists)$ such that $\mathfrak{A} \models \varphi$, there is a structure $\mathfrak{B}$ such that $\mathfrak{B} \models \varphi$ and $\mathfrak{A} \not\preceq^k \mathfrak{B}$. Many of the results are proved using the extension axioms of Section 2.1. Recall that θ_k is the conjunction of all k variable extension axioms (Example 6). It can be verified that if $\mathfrak{A} \models \theta_k$, then for all structures $\mathfrak{B}$, $\mathfrak{B} \preceq^k \mathfrak{A}$. This and related properties of models of the extension axioms are explored in [52].

References

[1] S. Abiteboul, M. Y. Vardi, and V. Vianu. Computing with infinitary logic. *Theoretical Computer Science*, 149:101–128, 1995.

[2] S. Abiteboul, M. Y. Vardi, and V. Vianu. Fixpoint logics, relational machines, and computational complexity. *J. ACM*, 44:30–56, 1997.

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

[4] S. Abiteboul and V. Vianu. Generic computation and its complexity. In *Proc. 23rd ACM Symposium on the Theory of Computing*, 1991.

[5] S. Abiteboul and V. Vianu. Computing with first-order logic. *J. of Computer and System Sciences*, 50(2):309–335, 1995.

[6] P. Aczel. An introduction to inductive definitions. In J. Barwise, editor, *Handbook of Mathematical Logic*. North Holland, 1977.

[7] F. Afrati, S. S. Cosmadakis, and M. Yannakakis. On Datalog vs. polynomial time. *J. Computer and System Sciences*, 51:177–196, 1995.

[8] A.V. Aho and J.D. Ullman. Universality of data retrieval languages. In *Proc. 6th ACM Symposium on Principles of Programming Languages*, pages 110–117, 1979.

[9] M. Ajtai and R. Fagin. Reachability is harder for directed than for undirected graphs. *Journal of Symbolic Logic*, 55:113–150, 1990.

[10] J. Barwise. On Moschovakis closure ordinals. *Journal of Symbolic Logic*, 42:292–296, 1977.

[11] 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.

[12] B. Bollobás. *Random Graphs*. Academic Press, 1985.

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

[14] A. Dawar. *Feasible Computation through Model Theory*. PhD thesis, University of Pennsylvania, 1993.

[15] A. Dawar. A restricted second order logic for finite structures. *Information and Computation*, to appear.

[16] A. Dawar. Types and indiscernibles in finite models. In J.A. Makowsky, editor, *Logic Colloquium '95*, Lecture Notes in Logic. Springer, 1997. to appear.

[17] A. Dawar and L. Hella. The expressive power of finitely many generalized quantifiers. *Information and Computation*, 123(2):172–184, 1995.

[18] A. Dawar, L. Hella, and Ph. G. Kolaitis. Implicit definability and infinitary logic in finite model theory. In Z. Fülöp and F. Gécseg, editors, *ICALP95*, volume 944 of *LNCS*, pages 624–635. Springer, 1995.

[19] A. Dawar, S. Lindell, and S. Weinstein. Infinitary logic and inductive definability over finite structures. *Information and Computation*, 119(2):160–175, 1995.

[20] A. Dawar, S. Lindell, and S. Weinstein. First order logic, fixed point logic and linear order. In H. Kleine-Büning, editor, *Computer Science Logic '95*, volume 1092 of *LNCS*, pages 161–177. Springer, 1996.

[21] M. de Rougemont. Second-order and inductive definability on finite structures. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 33:47–63, 1987.

[22] A. Durand, C. Lautemann, and T. Schwentick. Subclasses of binary NP. Technischer Report 1/96, Institut für Informatik, Johannes Gutenberg-Universität Mainz, 1995.

[23] E. Engeler. *Formal Languages: Automata and Structures.* Markham, 1968.

[24] R. Fagin. Generalized first-order spectra and polynomial-time recognizable sets. In R. M. Karp, editor, *Complexity of Computation, SIAM-AMS Proceedings, Vol 7*, pages 43–73, 1974.

[25] R. Fagin. Probabilities on finite models. *Journal of Symbolic Logic*, 41(1):50–58, March 1976.

[26] R. Fagin, L.J. Stockmeyer, and M.Y. Vardi. On monadic NP vs. monadic co-NP. *Information and Computation*, 120(1):78–92, 1995.

[27] Y. V. Glebskiĭ, D. I. Kogan, M.I. Ligon'kiĭ, and V. A. Talanov. Range and degree of realizability of formulas in the restricted predicate calculus. *Kibernetika*, 2:17–28, 1969.

[28] M. Grohe. Inversion of the L^k-invariants is hard. Unpublished manuscript.

[29] M. Grohe. Equivalence in finite-variable logics is complete for polynomial time. In *Proceedings of the 37th Annual Symposium on the Foundations of Computer Science.* IEEE Computer Society Press, 1996.

[30] M. Grohe. Large finite structures with few L^k-types. In *Proceedings of the 12th Annual IEEE Symposium on Logic in Computer Science.* IEEE Computer Society Press, 1997.

[31] Y. Gurevich. Logic and the challenge of computer science. In E. Börger, editor, *Current Trends in Theoretical Computer Science*, pages 1–57. Computer Science Press, 1988.

[32] Y. Gurevich, N. Immerman, and S. Shelah. McColm's conjecture. In *Proc. 9th IEEE Symp. on Logic in Computer Science*, pages 10–19, 1994.

[33] Y. Gurevich and S. Shelah. On finite rigid structures. *Journal of Symbolic Logic*, 61:549–562, 1996.

[34] I. Hodkinson. Finite variable logics. Revised version of Bull. Europ. Assoc. Theor. Comp. Sci. 51 (Oct 1993) 111–140, 1996.

[35] N. Immerman. Upper and lower bounds for first-order expressibility. *J. of Computer and System Sciences*, 25:76–98, 1982.

[36] N. Immerman. Relational queries computable in polynomial time. *Information and Control*, 68:86–104, 1986.

[37] N. Immerman. DSPACE[n^k]=VAR[$k+1$]. In *Proc. 6th IEEE Symp. on Structure in Complexity Theory*, pages 334–340, 1991.

[38] Ph.G. Kolaitis. On asymptotic probabilities of inductive queries and their decision problem. In R. Parikh, editor, *Logics of Programs*, volume 193 of *LNCS*, pages 153–166. Springer, 1985.

[39] Ph.G. Kolaitis and M.Y. Vardi. The decision problem for the probabilities of higher-order properties. In *Proc. 19th ACM Symposium on Theory of Computing*, pages 425–435, 1987.

[40] 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.

[41] Ph.G. Kolaitis and M.Y. Vardi. Infinitary logics and 0-1 laws. *Information and Computation*, 98(2):258–294, 1992.

[42] Ph.G. Kolaitis and M.Y. Vardi. On the expressive power of Datalog - tools and a case study. *J. Computer and System Sciences*, 51:110–134, 1995.

[43] Ph.G. Kolaitis and M.Y. Vardi. On the expressive power of variable confined logics. In *Proc. 11th IEEE Symp. on Logic in Computer Science*, pages 348–359, 1996.

[44] J.F. Lynch. Infinitary logics and very sparse random graphs. In *Proc. 8th IEEE Symposium on Logic in Computer Science*, pages 191–198, 1993.

[45] M. McArthur. The asymptotic behaviour of $L^k_{\infty\omega}$ on sparse random graphs. In R.B. Boppana and J.F. Lynch, editors, *Logic and Random Structures*, pages 53–63. AMS, 1997.

[46] M. McArthur and J. Spencer. Nonconvergence for sparse random graphs. manuscript.

[47] G. L. McColm. Parametrization over inductive relations of a bounded number of variables. *Annals of Pure and Applied Logic*, 48:103–134, 1990.

[48] G. L. McColm. When is arithmetic possible? *Annals of Pure and Applied Logic*, 50:29–51, 1990.

[49] Y. N. Moschovakis. *Elementary Induction on Abstract Structures.* North Holland, 1974.

[50] M. Otto. *Bounded Variable Logics and Counting*, volume 9 of *Lecture Notes in Logic.* Springer, 1997.

[51] B. Poizat. Deux ou trois choses que je sais de L_n. *J. Symbolic Logic*, 47(3):641–658, 1982.

[52] E. Rosen, Shelah S, and S. Weinstein. k-universal finite graphs. In R.B. Boppana and J.F. Lynch, editors, *Logic and Random Structures*, pages 65–77. AMS, 1997.

[53] E. Rosen and S. Weinstein. Preservation theorems in finite model theory. In D. Leivant, editor, *Logic and Computational Complexity*, volume 960 of *LNCS*, pages 480–503. Springer, 1995.

[54] A. Rubin. *Free Algebras in Von Neumann-Bernays-Godel Set Theory and Positive Elementary Inductions in Reasonable Structures.* PhD thesis, California Institute of Technology, 1975.

[55] A. Stolboushkin. Axiomatizable classes of finite models and definability of linear order. In *Proc. 7th IEEE Symp. on Logic in Computer Science*, 1992.

[56] S. Thomas. Theories with finitely many models. *J. Symbolic Logic*, 51:374–376, 1986.

[57] S. Thomas. Theories with finitely many models II. Unpublished manuscript, 1989.

[58] J. Tiuryn and P. Urzyczyn. Some relationships between logics of programs and complexity theory. *Theoretical Computer Science*, 60:83–108, 1988.

[59] J. Tyszkiewicz. On asymptotic probabilities in logics that capture DSPACE($\log n$). In *Proc. TAPSOFT'93*, volume 668 of *LNCS*. Springer, 1993.

[60] M. Y. Vardi. The complexity of relational query languages. In *Proc. 14th ACM Symposium on the Theory of Computing*, pages 137–146, 1982.