Types and Indiscernibles in Finite Models

Anuj Dawar

Department of Computer Science,
University of Wales Swansea,
Singleton Park, Swansea SA2 8PP, England
a.dawar@swansea.ac.uk

Abstract. We consider L^k — first order logic restricted to k variables, and interpreted in finite structures. The study of classes of finite structures axiomatisable with finitely many variables has assumed importance through connections with computational complexity. In particular, we investigate the relationship between the size of a finite structure and the number of distinct types it realizes, with respect to L^k. Some open questions, formulated as finitary Löwenheim-Skolem properties, are presented regarding this relationship. This is also investigated through finitary versions of an Ehrenfeucht-Mostowski property.

1 Introduction

In this paper, we are concerned only with finite structures. That is, our logical formulas are interpreted in relational structures with finite domain. Interest in finite model theory has largely grown from the fact that there is a close connection between definability of classes of finite structures and their computational complexity. For instance, Fagin [10] showed that a class of finite structures is definable in existential second order logic if, and only if, it is decidable by a nondeterministic Turing machine in polynomial time (i.e. in the class NP). Similarly, many naturally arising complexity classes have been characterised as definability classes in appropriate logics.

These results raised the hope that model theoretic methods could be deployed to attack some of the notoriously open problems in complexity theory. Unfortunately, most methods and results developed in model theory fail to work when only finite structures are considered (see [12]). For instance, the compactness theorem for first order logic trivially fails, as do most of its consequences. Indeed, as we make the transition from arbitrary structures to finite structures, first order logic loses its central role. Two significant reasons can be discerned for this: on the one hand, first order logic is too strong; and on the other hand, first order logic is too weak.

First order logic is too strong in the sense that for every finite structure $\mathfrak{A}$, there is a first order sentence $\varphi_{\mathfrak{A}}$ that describes it up to isomorphism: that is, for any structure $\mathfrak{B}$, if $\mathfrak{B} \models \varphi_{\mathfrak{A}}$, then $\mathfrak{B}$ is isomorphic to $\mathfrak{A}$. This means that the relation $\equiv$ of elementary equivalence is trivial on finite structures — it coincides with the isomorphism relation. Since a large part of model theory can arguably be described as the study of the structure of the elementary equivalence relation and its variants, this raises a serious question about the applicability of model theoretic methods.

First order logic is too weak in the sense that the class of finite models of any first order sentence is of very low computational complexity. It can, in fact, be recognised by a deterministic Turing machine using only a logarithmic amount of workspace (in other words, it is in the complexity class DSPACE[$\log n$]). Moreover, there are easily computed properties in DSPACE[$\log n$] that are not first order definable.

The disparity in the above two statements can be accounted for by contrasting the expressive power of single sentences with that of theories, i.e. sets of sentences. In model theory, we are most often concerned with the limits on the expressive power of sets of sentences. Clearly, two elementarily equivalent structures cannot be distinguished by any such set. However, when we only consider finite structures, *any* isomorphism closed class S of structures is defined by the countable set of sentences[1]:

$$\Sigma = \{\neg\varphi_{\mathfrak{A}} \mid \mathfrak{A} \notin S\},$$

where $\varphi_{\mathfrak{A}}$ is as above. On the other hand, as already mentioned, the expressive power of first order sentences is extremely weak. Another way of stating the same thing is that, while every class of structures is *axiomatisable*, few interesting classes are *finitely axiomatisable*.

One conclusion one can draw from the above is that in order to analyse the expressive power of first order logic (i.e. of first order sentences) and, indeed, of more powerful logics, we need equivalence relations that are coarser than the relation of elementary equivalence. There are two main approaches to coarsening the elementary equivalence relation that have proved useful. The first is to stratify the relation according to quantifier rank. This is useful in analysing first order definability. Indeed, it can be shown that a class of structures is finitely axiomatisable if, and only if, it is closed under the elementary equivalence relation of a fixed quantifier rank.[2] The other approach is provided by a restriction on the number of variables considered in defining the equivalence relation. This way we obtain a notion of axiomatisability intermediate between finite and full axiomatisiability — namely the notion of a class of structures axiomatised with a finite number of variables. This turns out to be a valuable approach in studying several natural extensions of first order logic, particularly extensions by inductive and fixed point operators. For several such extensions, it can be shown that any definable class of structures is, in fact, axiomatisable with a finite number of variables. Moreover, it has been shown that many of the difficult open questions of computational complexity can be recast as questions on the expressive power of these logics.

In the rest of this paper, a signature σ is a finite relational signature, i.e. it is a finite sequence of relation symbols $(R_1, \ldots, R_m)$ and associated arities $(r_1, \ldots, r_m)$, where $r_i \in \omega$. A σ-structure $\mathfrak{A}$ is an interpretation of this signature

[1] Here, to say that Σ defines S means that S is exactly the class of finite models of Σ.

[2] Provided that the vocabulary is purely relational, i.e. it contains no function symbols.

in a finite domain A, i.e. $\mathfrak{A} = (A, R_1^{\mathfrak{A}}, \ldots, R_m^{\mathfrak{A}})$. The *width* of σ, denoted width(σ) is defined as $\max(r_1, \ldots, r_m)$. The domain A of a structure $\mathfrak{A}$ is often denoted $|\mathfrak{A}|$, while the cardinality of a set S is denoted card(S). For a sentence φ (or a set of sentences Σ), Mod(φ) (or Mod(Σ)) denotes the set of finite models of φ (or Σ).

2 Finite Variable Logics

Define L^k to be the fragment of first order logic in which we only use the variables $x_1, \ldots, x_k$. For any sentence φ of L^k, we denote by $qr(\varphi)$ the *quantifier rank* of φ, i.e. the depth of nesting of quantifiers in φ.

The following definition introduces some basic notation used throughout the rest of this paper.

Definition 1. For two structures $\mathfrak{A}$ and $\mathfrak{B}$,

- we write
$$\mathfrak{A} \equiv^k \mathfrak{B}$$
if the two structures agree on all sentences of L^k.
- If s is a tuple of elements in $\mathfrak{A}$ and t is a tuple of elements in $\mathfrak{B}$, with length(s) = length(t) $\leq k$, by abuse of notation, we also write
$$(\mathfrak{A}, s) \equiv^k (\mathfrak{B}, t)$$
to denote that for any formula φ of L^k,
$$\mathfrak{A} \models \varphi[s] \quad \text{if, and only if,} \quad \mathfrak{B} \models \varphi[t].$$
- k-size($\mathfrak{A}$) is the index (i.e. the number of equivalence classes) of the equivalence relation $\equiv^k$ on the set $\{(\mathfrak{A}, s) \mid s \in |\mathfrak{A}|^k\}$.

It can be shown that, if $\mathfrak{A}$ is a finite structure and $\mathfrak{A} \equiv^k \mathfrak{B}$, then
$$k\text{-size}(\mathfrak{A}) = k\text{-size}(\mathfrak{B})$$
(see, for instance, [8]).

The equivalence relation $\equiv^k$ has a natural characterisation in terms of two player pebble games in the style of Ehrenfeucht-Fraïssé games. This characterisation is essentially due to Barwise [4] (see also [13, 18]).

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 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. 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 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. Duplicator has a strategy to avoid losing for q moves, starting with the initial position $(\mathfrak{A}, s)$ and $(\mathfrak{B}, t)$ if, and only if, $(\mathfrak{A}, s)$ and $(\mathfrak{B}, t)$ cannot be distinguished by any formula of L^k of quantifier rank q or less. Hence, if Duplicator has a strategy to play the game indefinitely without losing, then $(\mathfrak{A}, s) \equiv^k (\mathfrak{B}, t)$.

The following fact regarding the relations $\equiv^k$ was established in [8]:

Theorem 2. *For every finite structure $\mathfrak{A}$ and tuple s of elements of $\mathfrak{A}$, with* $\mathrm{length}(s) \leq k$, *there is a formula φ of L^k such that $\mathfrak{B} \models \varphi[t]$ if, and only if, $(\mathfrak{A}, s) \equiv^k (\mathfrak{B}, t)$. Moreover, $qr(\varphi) \leq k\text{-size}(\mathfrak{A}) + k$.*

In particular, it follows from Theorem 2 that the set of L^k sentences true in any finite model is, in fact, finitely axiomatisable within L^k.

Another useful notion is that of the L^k-type of a tuple in a structure. The following definition is from [8].

Definition 3. Let s be an l-tuple of elements from a structure $\mathfrak{A}$ for $l \leq k$. The *L^k-type* of s in $\mathfrak{A}$, denoted $\mathrm{Type}_k(\mathfrak{A}, s)$, is the set of formulas φ of L^k with free variables among $x_1, \ldots, x_l$, such that $\mathfrak{A} \models \varphi[s]$.

Thus, $(\mathfrak{A}, s) \equiv^k (\mathfrak{B}, t)$ can be seen as shorthand for $\mathrm{Type}_k(\mathfrak{A}, s) = \mathrm{Type}_k(\mathfrak{B}, t)$, and Theorem 2 states that every L^k-type is completely determined by one of its elements.

The interest in finite variable logics stems largely from the fact that a number of logics that have been considered in finite model theory — particularly in the characterisation of computational complexity classes — have the property that the definable classes of structures can all be axiomatised with a finite number of variables.

Thus, for instance, LFP is the closure of first order logic under an operation for forming the least fixed points of positive formulas. Similarly, PFP closes first order logic under an operation forming the iterative fixed points of arbitrary formulas [2]. It was independently proved by Immerman and Vardi [14, 20] that LFP exactly characterises polynomial time complexity (the class PTIME) on structures equipped with a linear order, while Abiteboul and Vianu [2] show that PFP similarly characterises polynomial space complexity (the class PSPACE) on ordered structures.

By results of Kolaitis and Vardi [15] and Dawar *et al.* [8], we know the following fact:

> For any sentence φ of LFP (or PFP), there is a k and a set Σ of sentences of L^k such that $\mathrm{Mod}(\varphi) = \mathrm{Mod}(\Sigma)$.[3]

The significance of this fact lies in the observation (see [15]) that most results showing that some property cannot be expressed in a logic such as LFP or PFP,

[3] Kolaitis and Vardi [15] show that each such φ is equivalent to a sentence of the infinitary logic $L^{\omega}_{\infty\omega}$, while in [8], it is shown that every sentence of $L^{k}_{\infty\omega}$ is equivalent to a single countable conjunction of sentences of L^k.

actually show that the property cannot be axiomatised in L^k, for any k. The essential tool for establishing this is the pebble game defined above.

What's more, Abiteboul and Vianu [3] show that, even without the restriction to ordered structures,

$$\mathrm{LFP} = \mathrm{PFP} \quad \text{if, and only if,} \quad \mathrm{PTIME} = \mathrm{PSPACE},$$

where, equality for the logics means that they have the same expressive power.

Recent work has shown that this result of Abiteboul and Vianu can be extended to a variety of complexity classes. That is, for a large number of such classes, logics have been found that characterise the class on ordered structures and for which the expressive power is bounded by finite variable axiomatisability (see [1, 6, 7]). Moreover, the inclusion relations between these logics exactly mirror the inclusions among the corresponding complexity classes. Abiteboul *et al.* [1] have shown that definability in these logics corresponds to complexity on a relational model of computation, in which complexity is measured not as a function of the size of the input, but as a function of the k-size of the input.

Translations of complexity theoretic questions into the framework of logics with limited discerning power has not (as yet, anyway) thrown new light on the separation of complexity classes. One reason for this was hinted at above, where it was stated that the main tool for establishing the inexpressibility of a class of structures in a logic such as LFP is the pebble game for L^k, which is used to show that the class is not closed under the relation $\equiv^k$ for any k. But then, the class is not definable in any of the logics we are considering, and this tool is not directly useful in separating the logics in question.

Nevertheless, some finer separation results have been obtained. For instance, in [8] it is shown that there is a polynomial time decidable property that is closed under $\equiv^k$ for a fixed k, and is still not definable in LFP and in [7] two different ways of restricting implicit first order definitions to finite variable axiomatisibility are separated. Both of these results rely on the construction of classes of structures in which, though the k-size is unbounded, it grows much more slowly than the cardinality of the structures themselves. This suggests that a greater understanding of the relationship between the size of a structure and its k-size could be very useful (see [5] for a discussion of some of these issues). Some questions regarding this relationship are investigated in the next section.

3 Finitary Löwenheim-Skolem Properties

We begin this section with a few straightforward remarks about the relationship between the size of a structure and its k-size. First of all, it is immediate from the definition that, if $\mathrm{card}(|\mathfrak{A}|) = n$, then $k\text{-size}(\mathfrak{A}) \leq n^k$, since there are at most n^k tuples in $\mathfrak{A}$. 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 (sometimes called pure sets), 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. On closer examination, we find that this class splits into only finitely many $\equiv^k$ equivalence classes of structures (indeed, this is necessarily true of any class of structures in which the k-size of structures is bounded — see [9]). 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 (recall that the k-size is invariant among structures in a given $\equiv^k$ equivalence class)? Are there any recursive bounds on either of these? This leads us to formulate the following two open questions, which can be seen as downward and upward Löwenheim-Skolem properties for finite structures (over some fixed signature σ, with

$$2 \leq \text{width}(\sigma) \leq k$$

— the case of unary signatures is straightforward).

Open Question 1 Is there a recursive function f_k, for every k, such that, for every finite structure $\mathfrak{A}$ with k-size$(\mathfrak{A}) \leq n$, there is a $\mathfrak{B}$ such that

$$\text{card}(|\mathfrak{B}|) \leq f_k(n) \text{ and } \mathfrak{A} \equiv^k \mathfrak{B}?$$

Open Question 2 Is there a recursive function g_k for every k, such that, if $\mathfrak{B}$ is a finite structure with k-size$(\mathfrak{B}) \leq n$ and card$(|\mathfrak{B}|) \geq g_k(n)$, then there are arbitrarily large finite $\mathfrak{A}$ such that $\mathfrak{A} \equiv^k \mathfrak{B}$?

We first observe that if we drop the word recursive from the above questions, then the answer is yes in both cases. This is based on the fact that for any fixed q, there are only finitely many sentences φ of L^k with $qr(\varphi) \leq q$.[4] Thus, for any n, the relation $\equiv^k$ has finite index on the class of structures $\mathfrak{A}$ such that k-size$(\mathfrak{A}) \leq n$ (by virtue of Theorem 2). We can then define the functions $d_k, e_k : \mathbf{N} \to \mathbf{N}$ by:

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

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

Here, by finite $\equiv^k$ class, we mean a class that contains finitely many structures up to isomorphism. The functions d_k and e_k are well defined, because they involve taking, for each n, the maximum of a finite set of cardinalities.

The Open Questions 1 and 2 posed above can now be seen as asking if there are recursive upper bounds on the functions d_k and e_k. Moreover, since the relation $\equiv^k$ is decidable, if such upper bounds exist, then the functions d_k and e_k are themselves recursive.

It is also a worthwhile line of investigation to establish lower bounds on the functions d_k and e_k. At present, the best known lower bounds on these functions are exponential, as described in the next section. The following open question can thus be added to the two mentioned above:

[4] Once again, we use the fact that our signatures are purely relational.

Open Question 3 Are there super-exponential lower bounds on the functions d_k and e_k, for any k?

By super-exponential, we mean of the order of $2^{f(n)}$ for some super-polynomial function f. Even though the exponential lower bounds established for trees (and discussed in the next section) are of the form $2^{g(n)}$, where the function g is $O(n)$, there are constructions known that raise this to at least 2^{n^2}.

It should also be pointed out that the case of $k = 2$ is solved, as linear lower and upper bounds are essentially established in [17]. Thus, we adopt the following proviso for the rest of the paper:

Wherever k refers to the number of variables, we assume that $k \geq 3$.

4 Trees

The exponential lower bounds on the functions d_k and e_k can be established by considering the class of complete binary trees. We introduce this class through a series of definitions.

Definition 4. – A tree T is a structure $(V, \preceq)$ where,
- $\preceq$ is a partial order on the set of vertices V;
- there is a unique $\preceq$-minimal element of V, called the *root*; and
- every element $v \in V$ (other than the root) has a unique $\preceq$-predecessor, called the *parent* of v.

– Two vertices are *siblings* if they have the same parent. If u is the parent of v, we also call v a *child* of u.
– The *subtree* rooted at a vertex v is the substructure of T formed by the set $\{u \in V \mid v \preceq u\}$.
– The *height* of a tree is the length of its longest maximal chain.
– For any $c \in \omega$, a *c-ary tree* is a tree in which every vertex has at most c $\preceq$-successors.
– A *complete c-ary tree* is one in which every vertex has either 0 or c successors, and all maximal chains are of equal length.

Now, a lower bound on the functions d_k and e_k is obtained from the following three facts. Let T_h denote the (unique up to isomorphism) complete binary tree of height h.

1. T_h is characterised up to isomorphism by a sentence of L^3.
2. $\text{card}(T_h) = 2^h - 1$.
3. For any k, there is a k' such that $k\text{-size}(T_h) \leq h^{k'}$, for all h.

For details of the above, the reader is referred to [8]. Here, we only note that (1) above is based on the following lemma (stated here without proof, details can be found in [5], for instance), which will be of use later:

Lemma 5. *For every $d \in \omega$, there is a formula δ_d of L^2 such that, in any partial order $P = (V, \preceq)$, $P \models \delta_d[u]$ if, and only if, the length of the maximal chain below u is d.*

A natural first attempt to improve on the exponential lower bound is to consider trees with higher than binary branching. Clearly, facts analogous to the above hold for c-ary trees for any constant c, i.e. each complete c-ary tree is described up to isomorphism by a sentence with no more than $\max(c+1, 3)$ variables; the size of the complete c-ary tree of height h is exponential in h, but the k-size is polynomial in h. This, of course, does not improve on the exponential lower bound.

Another approach is to consider trees of uniform arity, whose arity is a function of their height (this approach is considered in [1]). However, this also fails as we can no longer describe such a tree up to isomorphism with a fixed (independent of the height) number of variables. In particular, if T is a complete c-ary tree, then

1. there is a T' such that $T \equiv^k T'$ and the arity of T' is at most k; and
2. if $c > k$, there are infinitely many non-isomorphic T'' such that $T \equiv^k T''$.

The above facts are consequences of the following general property of trees:

Lemma 6. *Let $v_1, \ldots, v_{k+1}$ be siblings in a tree T, let $T_1, \ldots, T_{k+1}$ be the respective subtrees rooted at these nodes, and let v be their common parent. If $T_i \equiv^k T_j$ for $1 \leq i, j \leq k+1$, then*

1. $T \equiv^k T'$, where T' is obtained from T be removing T_{k+1}.
2. $T \equiv^k T''$, where T'' is obtained from T by adding a new subtree $T_{k+2} \equiv^k T_{k+1}$ above v.

Proof: The proof is a straightforward pebble game. We describe Duplicator's winning strategy in the game played on T and T'. The game on T and T'' is analogous.

First of all, if Spoiler plays anywhere (in T or T') other than in the subtrees $T_1, \ldots, T_{k+1}$, Duplicator simply plays on the same node in the other tree. If Spoiler plays in one of the trees T_i (in either T or T') which does not contain a pebble, Duplicator responds by playing in a subtree T_j of the other tree not containing a pebble (not counting the pebble being currently moved). Such a subtree must exist, because there are at least k matching subtrees and only $k-1$ pebbles on the board. The vertex played by Duplicator is determined by her winning strategy in the k pebble game on structures T_i and T_j, given by the fact that $T_i \equiv^k T_j$. If Spoiler plays a pebble in a subtree T_i which already contains a pebble, Duplicator responds with a pebble in the subtree T_j containing the matching pebble. Again, she uses the winning strategy for the game on T_i and T_j to guarantee that the result is a legal position. It is not difficult to verify that this describes a winning strategy for Duplicator on T and T'. ■

It is possible to use Lemma 6 to show an exponential *upper* bound on the functions d_k and e_k on the class of trees. Thus, the lower bound based on complete binary trees cannot be improved as long as we confine ourselves to trees.

In order to prove this upper bound, it is useful to establish some more basic properties of the $\equiv^k$ relation on trees.

Lemma 7. *For any tree $T = (V, \preceq)$, any vertices $u, v \in V$ and any $k \geq 3$, if*

$$(T, u) \equiv^k (T, v)$$

then the following two conditions hold:

1. *$(T, u') \equiv^k (T, v')$, where u' and v' are the parents of u and v respectively; and*
2. *$T_u \equiv^k T_v$, where T_u and T_v are the subtrees of T rooted at u and v respectively.*

Proof: Since $(T, u) \equiv^k (T, v)$, Duplicator has a winning strategy in the k pebble game played on two copies of T that begins with a pebble on u in one copy and on v in the other. In the following we will refer to these two copies as T^1 and T^2 respectively. We will also use superscripts 1 and 2 to distinguish vertices and subtrees of T^1 from those of T^2.

If Spoiler's first move is to place a pebble on u' (in T^1), then Duplicator must respond with a pebble on v' (in T^2), as placing a pebble on any other vertex $v'' \prec v^2$ would allow Spoiler to pebble a vertex in between v^2 and v'' (for instance, v') resulting in a losing position for Duplicator. Now, from the assumption that Duplicator has a winning strategy, we conclude that Duplicator wins the game from the position (T^1, u, u'), (T^2, v, v') and this implies $(T, u') \equiv^k (T, v')$, establishing condition 1.

To establish condition 2, we argue that if, in the game starting from (T, u) and (T, v), Spoiler confines his moves to the subtrees T_u in T^1 and to T_v in T^2, then Duplicator has a winning strategy also confining her moves to these subtrees. Suppose this were not so and, for some move by Spoiler within the tree T_u^1 (without loss of generality), Duplicator is forced to respond with a vertex outside T_v^2. Note that, after the first such move by Duplicator, T^2 contains pebbles both within and outside T_v^2, since there was a pebble on v initially, while all pebbles in T^1 are within T_u. Spoiler now places a pebble on u^1, using a previously unused pebble if there are only two pairs of pebbles currently on the board, or possibly moving a pebble otherwise (though not moving the pebble used in the previous move). In either case, there remain pebbles within and outside T_v^2, not counting the pebble that Duplicator must now place. Duplicator must respond with a vertex v'' in T^2 that is a lower bound (in the partial order $\preceq$) of all the currently pebbled vertices in T^2. Since these pebbled vertices include vertices both within and outside of T_v^2, it must be the case that $v'' \prec v$. This means that u and v'' can be distinguished by a formula δ_d as in Lemma 5, giving Spoiler a winning strategy from this point on. ■

In fact, the converse of Lemma 7 holds as well, but we do not require it for our purposes. One immediate consequence of Lemma 7 that is of interest to us is the following:

Lemma 8. *If $(T, u) \not\equiv^k (T, v)$ and $u \preceq u'$, then there is no v', such that $v \preceq v'$ and $(T, u') \equiv^k (T, v')$.*

In other words, if $(T, u) \not\equiv^k (T, v)$, then the subtrees T_u and T_v realize completely disjoint sets of L^k-types in T.

Also, combining Lemma 6 with Lemma 7 gives us the following:

Lemma 9. *For every tree T there is a tree T', with $T \equiv^k T'$, such that there is no set $v_1, \ldots, v_{k+1}$ of distinct sibling nodes in T' with $(T', v_i) \equiv^k (T', v_j)$ for $1 \leq i, j \leq k+1$.*

Moreover, if T contains a set $v_1, \ldots, v_k$ of distinct sibling nodes such that $(T, v_i) \equiv^k (T, v_j)$ for $1 \leq i, j \leq k$, then there are arbitrarily large T'' such that $T \equiv^k T''$.

In addition, using Lemma 8 we can establish bounds on the size of the tree T' of Lemma 9, as in the following.

Lemma 10. *If T does not contain a set $v_1, \ldots, v_{k+1}$ of distinct sibling nodes such that $(T, v_i) \equiv^k (T, v_j)$ for $1 \leq i, j \leq k+1$ and k-size$(T) \leq n$, then $\mathrm{card}(|T|) \leq k^n$.*

Proof: We first note that, by virtue of Lemma 5, k-size(T) is at least as large as the height of T. Now, we establish the claim by induction on the height of the tree. The basis is trivial. For the induction step, we group the children of the root of T into equivalence classes according to their L^k-type in T. By the hypothesis of the lemma, there are at most k elements in each equivalence class. Let us say there are e equivalence classes: $E_1, \ldots, E_e$. Associate with each E_i, the number n_i of distinct L^k-types that are realized in the subtree rooted at one of the elements of E_i (this is, by Lemma 7, the same for all the elements of E_i).

By the induction hypothesis, the total number of elements in all the subtrees rooted at elements in E_i is at most $k \cdot k^{n_i}$, and thus

$$\mathrm{card}(|T|) \leq 1 + k \sum_{1 \leq i \leq e} k^{n_i}.$$

Since $k \geq 3$, for each i, $k^{n_i} \geq 3$, and therefore $\sum_{1 \leq i \leq e} k^{n_i} < k^{\sum_{1 \leq i \leq e} n_i}$. But, by Lemma 8, $\sum_{1 \leq i \leq e} n_i < n$, and we conclude that $\mathrm{card}(|T|) \leq k^n$. ∎

Note that, in the proof of Lemma 10, it suffices to count the L^k-types of individual elements. The k-size would be larger than this number, and the bound obtained could be made tighter. However, we are only interested in establishing the exponential order of the bound, and the simple argument given above serves that purpose.

Now, Lemmas 9 and 10 immediately yield the upper bounds on the functions d_k and e_k on the class of trees:

Theorem 11. *If T is a tree with k-size$(T) \leq n$, then there is a tree T' with $\mathrm{card}(|T'|) \leq k^n$ such that $T \equiv^k T'$. Moreover, if $\mathrm{card}(|T|) > k^n$, then there are arbitrarily large T'' such that $T \equiv^k T''$.*

The question arises as to whether any of the techniques that work for trees can be generalised to other classes of structures, and perhaps even to obtain a positive resolution of Open Questions 1 and 2. We outline one such attempt at generalisation in the next section, though it is one that fails to resolve the questions, in the general case.

5 Indiscernibles

The following definitions introduce notions of sets and sequences of indiscernibles for finite variable logics, that are analogous to the standard definitions from model theory.

First, recall the following definition:

Definition 12. Given a structure $\mathfrak{A}$ and a sequence of elements $s = (a_1, \ldots, a_k)$ from $\mathfrak{A}$, the *basic equality type* of $(\mathfrak{A}, s)$ is the formula

$$\bigwedge_{(i,j)\in S} (x_i = x_j) \wedge \bigwedge_{(i,j)\in T} \neg(x_i = x_j),$$

where $S = \{(i,j) \mid i < j \text{ and } a_i = a_j\}$, and $T = \{(i,j) \mid i < j \text{ and } a_i \neq a_j\}$.

That is, the basic equality type of a tuple is its complete quantifier free description in the language of identity.

Now we can introduce the definitions of indiscernible sets and sequences.

Definition 13. A set $X \subseteq |\mathfrak{A}|$ is a *k-indiscernible set* in $\mathfrak{A}$, if whenever $s, t \in X^k$ are such that s and t have the same basic equality type, then $(\mathfrak{A}, s) \equiv^k (\mathfrak{A}, t)$.

$X \subseteq |\mathfrak{A}|$ is a *k-indiscernible sequence* in $\mathfrak{A}$, if there is some (not necessarily definable) linear order on X such that for any order increasing $s, t \in X^l$, $l \leq k$, $(\mathfrak{A}, s) \equiv^k (\mathfrak{A}, t)$.

Clearly, any k-indiscernible set X is a k-indiscernible sequence, as any order on X would satisfy the conditions of the definition. However, the converse may fail.

The next definition is also a straightforward analogue, with finitely many variables, of a standard definition from model theory.

Definition 14. A structure $\mathfrak{B}$ is a *k-elementary substructure* of $\mathfrak{A}$, denoted $\mathfrak{B} \preceq^k \mathfrak{A}$, if $|\mathfrak{B}| \subseteq |\mathfrak{A}|$ and for every $t \in |\mathfrak{B}|^k$ and every formula φ of L^k,

$$\mathfrak{B} \models \varphi[t] \quad \text{if, and only if,} \quad \mathfrak{A} \models \varphi[t].$$

By analogy with the Ehrenfeucht-Mostowski stretching technique, we make the following definition.

Definition 15. A class of structures C has the *k-Ehrenfeucht-Mostowski property* (or k-EM property, for short) if, for every structure $\mathfrak{A}$ in C, if $X \subseteq |\mathfrak{A}|$ is a k-indiscernible sequence in $\mathfrak{A}$, with $\operatorname{card}(X) \geq k$, then

1. there is an $\mathfrak{A}' \in C$ with $\mathfrak{A}' \preceq^k \mathfrak{A}$ and $\operatorname{card}(|\mathfrak{A}'| \cap X) \leq k$.
2. for any set Y disjoint from $|\mathfrak{A}|$, there is an $\mathfrak{A}'' \in C$ such that $\mathfrak{A} \preceq^k \mathfrak{A}''$ and $X \cup Y$ is a k-indiscernible sequence in $\mathfrak{A}''$.

We motivated the introduction of indiscernible sets as an attempt at generalising the techniques used in the previous section on trees. We now justify this through the following theorem.

Theorem 16. *The class of trees has the k-Ehrenfeucht-Mostowski property.*

Proof: We prove that if $X = \{a_1, \ldots, a_n\}$ is a k-indiscernible sequence in the tree T with the order $a_i < a_j$ for $i < j$, then there is a collection of disjoint subtrees $T_1, \ldots, T_n$ of T such that:

1. $a_i \in T_i$ for $1 \leq i \leq n$;
2. $T_i \equiv^k T_j$ for $1 \leq i, j \leq n$; and
3. the roots of T_i and T_j are siblings for $1 \leq i, j \leq n$.

The result then follows from Lemma 6.

In order to establish 1–3, we first note that a straightforward modification of the formula δ_d of Lemma 5 yields a formula γ_d of L^3, such that $T \models \gamma_d[a, b]$ if, and only if, the distance from the root to the greatest lower bound of a and b is d. Thus, from the assumption that X is a sequence of indiscernibles, we conclude that there is a single vertex g such that g is the greatest lower bound of any pair a_i, a_j. By taking T_i to be the subtree rooted at a successor of g which contains a_i, properties 1 and 3 are easily established. Furthermore, since the elements of X all have the same L^k-type, property 2 follows immediately from Lemmas 7 and 8. ■

Other classes of structures that have the k-EM property include the class of linear orders (trivially, since no linear order contains a large enough indiscernible sequence), and the class of pure sets, i.e. structures in the empty signature. We shall see below examples of classes of structures where the k-EM property fails.

We will now show that we can establish recursive upper bounds for the functions d_k and e_k for any class of structures having the k-EM property.

Lemma 17. *There is a recursive function ι_k, for every $k \geq 3$, such that for any structure $\mathfrak{A}$, with k-size$(\mathfrak{A}) \leq n$ and card$(|\mathfrak{A}|) \geq \iota_k(n, m)$, there is a set $X \subseteq |\mathfrak{A}|$ which forms a k-indiscernible sequence in $\mathfrak{A}$ with card$(X) \geq m$.*

Proof: The proof is a direct application of the finite Ramsey theorem (see [11]). This guarantees the existence of a function ι_k such that if we take any ordered set A of more than $\iota_k(n, m)$ elements, and partition the set of order-increasing tuples in A^k into at most n sets $P_1, \ldots, P_n$, then there must exist a homogeneous subset X of A with card$(X) \geq m$, i.e. there is an i such that $[X]^k \subseteq P_i$, where $[X]^k$ is the set of order increasing tuples of elements of X.

Now, taking A to be the universe of $\mathfrak{A}$, choosing an arbitrary ordering of the elements, and letting the sets P_i be the $\equiv^k$ equivalence classes on the order increasing tuples in $|\mathfrak{A}|^k$, the resulting set X is guaranteed to be a k-indiscernible sequence in $\mathfrak{A}$. ■

Now, the following theorem is immediate.

Theorem 18. *If a class of structures C has the k-EM property, then there are recursive functions f_k and g_k satisfying the conditions of Open Questions 1 and 2, on C.*

Proof: For the function g_k, just take $g_k(n) = \iota_k(n, k)$.
For f_k, let $f_k(n) = \iota_k(n, k+1) - 1$. Note that any structure $\mathfrak{A}$ with $k-size(\mathfrak{A}) \leq n$ and $\text{card}(\mathfrak{A}) > f_k(n)$ m has an indiscernible sequence of $k+1$ elements, and therefore contains a proper k-elementary substructure. ■

Unfortunately, the k-EM property does not generalise from the class of trees to the class of all finite structures. It does not even generalise to the class of all partial orders, as the following example shows.

Let B_n be the (unique up to isomorphism) partial order induced by a Boolean algebra with n atoms. This structure is characterised up to isomorphism by a sentence of L^3, namely the sentence that is the conjunction of the axioms of Boolean algebras (stated in terms of the partial order relation) with $\exists x \delta_n \wedge (\neg \exists x \delta_{n+1})$, where the formulas δ are from Lemma 5. Moreover, the set A of atoms is a k-indiscernible set in B_n, for any k, since any permutation of A can be extended to an automorphism of B_n. Thus, we obtain structures with arbitrarily large sets of k-indiscernibles, yet which are, up to isomorphism, unique in their $\equiv^k$ equivalence class.

Though the class of Boolean algebras does not have the k-EM property, it certainly does have recursive (even exponential) bounds on the functions d_k and e_k. This suggests that the k-EM property is too strong a requirement. One way to weaken it is to require collapsing and stretching only for sets of k-indiscernibles of size much larger than k itself. In fact, as the above argument shows, we cannot replace k by any constant in this requirement, and get a property that is satisfied by the class of Boolean algebras. We must, instead, replace it by a function of the k-size of $\mathfrak{A}$.

Definition 19. A class of structures C has the *weak k-Ehrenfeucht-Mostowski property* (or weak k-EM property, for short), if there is a recursive function h such that for every structure $\mathfrak{A}$ in C, if k-size$(\mathfrak{A}) \leq n$ and $X \subseteq |\mathfrak{A}|$ is a k-indiscernible sequence in $\mathfrak{A}$, with $\text{card}(X) \geq h(n)$, then

1. there is an $\mathfrak{A}' \preceq^k \mathfrak{A}$ with $\text{card}(|\mathfrak{A}'| \cap X) \leq h(n)$.
2. for any set Y disjoint from $|\mathfrak{A}|$, there is an $\mathfrak{A}''$ such that $\mathfrak{A} \preceq^k \mathfrak{A}''$ and $X \cup Y$ forms a k-indiscernible sequence in $\mathfrak{A}''$.

With this definition, it can be verified that the class of Boolean algebras has the weak k-EM property.

Moreover, the weak k-EM property is still strong enough for the following:

Theorem 20. *If a class of structures C has the weak k-EM property, then there are recursive functions f_k and g_k satisfying the conditions of Open Questions 1 and 2, on C.*

Proof: Just take $g_k(n) = \iota_k(n, h(n))$ and $f_k(n) = \iota_k(n, h(n) + 1) - 1$. ■

However, the weak k-EM property is still too strong to hold of the class of all finite structures. This is a consequence of a recent result due to Luczak [16] (see also [19] for a proof of this result). Luczak shows that the sentence θ_k which is the conjunction of all extension axioms with k variables (see [15] for the definition of θ_k) has infinitely many minimal models with respect to the substructure relation. Since Kolaitis and Vardi [15] showed that all models of θ_k are $\equiv^k$ equivalent, it follows that the class $\mathrm{Mod}(\theta_k)$ does not have the weak k-EM property. Note that this does not negatively resolve Open Questions 1 and 2, but it does show that we cannot hope to replace $\mathfrak{A} \equiv^k \mathfrak{B}$ with the stronger requirement $\mathfrak{B} \preceq^k \mathfrak{A}$ (at least in Question 1).

We end with the observation that we can define a stronger notion of indiscernibility, that allows us to prove collapsing and stretching properties (*à la* Ehrenfeucht-Mostowski) over the class of all structures. This is given by the following.

Definition 21. A set $X \subseteq |\mathfrak{A}|$ is a set of *strong k-indiscernibles* in $\mathfrak{A}$, if whenever $s \in |\mathfrak{A}|^k$, $a, a' \in X$ and a' is not in s, then, $(\mathfrak{A}, s) \equiv^k (\mathfrak{A}, s[a/a'])$, where $s[a/a']$ is the tuple obtained from s by replacing all occurrences of a by a'.

We then have the following:

Theorem 22. *If $X = \{a_1, \ldots, a_{k+1}\} \subseteq |\mathfrak{A}|$ is a set of strong k-indiscernibles in $\mathfrak{A}$, then*

1. *there is an $\mathfrak{A}' \preceq^k \mathfrak{A}$ with $|\mathfrak{A}'| = |\mathfrak{A}| \setminus \{a_1\}$*
2. *for any set Y disjoint from $|\mathfrak{A}|$, there is an $\mathfrak{A}''$, with $|\mathfrak{A}''| = |\mathfrak{A}| \cup Y$, such that $\mathfrak{A} \preceq^k \mathfrak{A}''$ and $X \cup Y$ is a set of strong k-indiscernibles in $\mathfrak{A}''$.*

The proof of the stretching part (part 2) of Theorem 22 appears in [18]. The proof of the first part is similar, and we omit it.

Unfortunately, with this stronger notion of indiscernibility, we do not have an analogue of Lemma 17.

Acknowledgements: I am grateful to Martin Grohe, Henrik Imhof and Eric Rosen for carefully reading an unpolished version of this paper.

References

1. S. Abiteboul, M. Y. Vardi, and V. Vianu. Fixpoint logics, relational machines, and computational complexity. In *Proc. 7th IEEE Symp. on Structure in Complexity Theory*, 1992.
2. S. Abiteboul and V. Vianu. Datalog extensions for database queries and updates. *Journal of Computer and System Sciences*, 43:62–124, 1991.
3. S. Abiteboul and V. Vianu. Computing with first-order logic. *Journal of Computer and System Sciences*, 50(2):309–335, 1995.
4. J. Barwise. On Moschovakis closure ordinals. *Journal of Symbolic Logic*, 42:292–296, 1977.
5. A. Dawar. *Feasible Computation through Model Theory.* PhD thesis, University of Pennsylvania, 1993.

6. A. Dawar. A restricted second order logic for finite structures. In D. Leivant, editor, *Logic and Computational Complexity*, volume 960 of *LNCS*, pages 393–413. Springer-Verlag, 1995.
7. 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-Verlag, 1995.
8. A. Dawar, S. Lindell, and S. Weinstein. Infinitary logic and inductive definability over finite structures. *Information and Computation*, 119(2):160–175, 1995.
9. 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-Verlag, 1996.
10. 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.
11. R. L. Graham, B. L. Rothschild, and J. H. Spencer. *Ramsey Theory.* Wiley, 1980.
12. Y. Gurevich. Toward logic tailored for computational complexity. In M. Richter et al., editors, *Computation and Proof Theory*, pages 175–216. Springer Lecture Notes in Mathematics, 1984.
13. N. Immerman. Upper and lower bounds for first-order expressibility. *Journal of Computer and System Sciences*, 25:76–98, 1982.
14. N. Immerman. Relational queries computable in polynomial time. *Information and Control*, 68:86–104, 1986.
15. Ph. G. Kolaitis and M. Y. Vardi. Infinitary logics and 0-1 laws. *Information and Computation*, 98(2):258–294, 1992.
16. T. Łuczak, 1995. Personal Communication.
17. M. Otto. Ptime canonization for two variables with counting. In *Proc. 10th IEEE Symp. on Logic in Computer Science*, pages 342–352, 1995.
18. B. Poizat. Deux ou trois choses que je sais de L_n. *Journal of Symbolic Logic*, 47(3):641–658, 1982.
19. E. Rosen, S. Shelah, and S. Weinstein. k-Universal finite graphs. In R. Boppana and J. F. Lynch, editors, *Logic and Random Structures*, DIMACS Series in Deiscrete Mathematics and Theoretical Computer Science. American Mathematical Society, 1996. to appear.
20. 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.