The Role of Decidability in First Order Separations over Classes of Finite Structures

Steven Lindell[*]
Department of Computer Science
Haverford College
Haverford, PA 19041, U.S.A.
slindell@haverford.edu

Scott Weinstein[†]
Graduate Program in Logic and Algorithms
University of Athens
Athens GR-15784, Greece
weinstein@cis.upenn.edu

Abstract

We establish that the decidability of the first order theory of a class of finite structures C *is a simple and useful condition for guaranteeing that the expressive power of FO + LFP properly extends that of FO on* C*, unifying separation results for various classes of structures that have been studied. We then apply this result to show that it encompasses certain constructive pebble game techniques which are widely used to establish separations between FO and FO + LFP, and demonstrate that these same techniques cannot succeed in performing separations from any complexity class that contains DLOGTIME.*

1. Introduction

Descriptive complexity theory studies the connection between logical definability and computational complexity. A basic result of this theory, due to Immerman [7] and Vardi [13], is that first order logic supplemented with a least fixed point operator (FO+LFP) captures exactly the polynomial time computable properties of ordered finite structures. Since first order logic does not contain any overt mechanism for defining relations by recursion, it can be shown that over many classes of finite structures the expressive power of FO+LFP properly extends the expressive power of FO. On the other hand, the exact relation between the expressive power of FO and FO+LFP remains a question of considerable interest. Kolaitis and Vardi [8] conjectured that the expressive power of FO+LFP properly extends that of FO over any infinite class of finite ordered structures (the Ordered Conjecture). As it happens, either a positive or negative resolution of (even special cases of) the Ordered Conjecture would have important complexity theoretic consequences (see, for example, Atserias & Kolaitis [1], Dawar, Lindell, & Weinstein [3], Gurevich, Immerman, & Shelah [6]).

In this paper, we establish that the decidability of the first order theory of a class of finite structures C is a simple and useful condition for guaranteeing that the expressive power of FO + LFP properly extends that of FO on C, unifying separation results for various classes of structures that have been studied, see, *e.g.*, Lindell [9]. We then apply this result to show that it encompasses certain constructive pebble game techniques which are widely used to establish separations between FO and FO + LFP, and demonstrate that these same techniques cannot succeed in performing separations from any class that contains DLOGTIME.

2. Background

We assume the standard definitions of a first order language (or signature) and a structure interpreting it. Unless otherwise mentioned, all structures we will be dealing with are assumed to have finite universe and all signatures are assumed to be finite and relational, that is, to consist of finitely many relation symbols. Indeed, for definiteness, we will assume throughout that the universe of every structure is an initial segment of the natural numbers, unless otherwise noted. We write $\mathcal{F}_\sigma$ to denote the class of all finite structures of signature σ.

An *n–ary query* over a class of structures $\mathcal{C}$ is a map Q sending each structure $A \in \mathcal{C}$ to an n–ary relation over A which satisfies the following condition: for all $A, B \in \mathcal{C}$, if f is an isomorphism from A onto B, then $Q(B) = f[Q(A)]$. We will generally be concerned with *boolean queries* which may then be identified with sets of finite structures which are closed under isomorphism or with the characteristic functions of such sets.

We will write FO to denote both the set of first-order formulas, and also the class of queries that are first order expressible. If C is a collection of structures, we write FO(C) for the collection of first order queries restricted to C. Similar notation will be used for the fixed-point queries after

---

[*] Supported in part by NSF CCR-9820899.

[†] Supported in part by NSF CCR-9820899. On leave from Department of Philosophy, University of Pennsylvania, Philadelphia, PA 19104, U.S.A.

they are defined below.

We write $\mathrm{Th}(\mathsf{C})$ for the first order theory of C, that is, the set of sentences each of which is true in every structure in C; we write $\mathrm{Th}^{\infty}(\mathsf{C})$ for the first order limit theory of C, that is, the set of first order sentences each of which is true in all but finitely many structures in C. The quantifier rank of a first order sentence φ is the maximum depth of nesting of quantifiers in φ. If A and B are structures, we say that A is r-equivalent to B (written $A \equiv_r B$), if and only if, A and B satisfy exactly the same sentences of quantifier rank $\leq r$. Note that since all signatures are finite and relational, for every r, the equivalence relation $\equiv_r$ is of finite index.

Let $\varphi(R, x_1, \ldots, x_k)$ be a first-order formula. On a structure, A, φ defines the operator, $\Phi_A(R^A) = \{\langle a_1, \ldots a_k\rangle | \langle A, R^A\rangle \models \varphi[a_1, \ldots, a_k]\}$. If φ is an R–positive formula, Φ_A is monotone. We may view φ as determining an induction on A the stages of which are defined as follows: $\varphi^0_A = \emptyset$; $\varphi^{m+1}_A = \Phi_A(\varphi^m_A)$. The *closure ordinal* of φ on A, denoted $||\varphi||_A$, is the least m such that $\varphi^m_A = \varphi^{m+1}_A$. The m^{th} stage of the induction determined by φ can be uniformly defined over all structures by a first–order formula which we denote by φ^m. The set inductively defined by φ on A, denoted φ^{∞}_A, is the least fixed point of the operator Φ_A, that is, $\varphi^{\infty}_A = \varphi^m_A$, where $m = ||\varphi||_A$. We write FO+LFP for the extension of first-order logic with the lfp operation which uniformly determines the least fixed point of an R–positive formula. That is, for any R–positive formula φ, $\mathrm{lfp}(R, x_1, \ldots, x_k)\varphi$ is a formula of FO+LFP and $A \models \mathrm{lfp}(R, x_1, \ldots, x_k)\varphi[a_1, \ldots, a_k]$, if and only if, $\langle a_1, \ldots, a_k\rangle \in \varphi^{\infty}_A$.

3. First–Order Definability vs. Fixed–Point Definability

We begin with a brief review of the resolution of a conjecture of McColm which relates the expressive power of FO and FO+LFP. The following definition was introduced by McColm [10].

Definition 1 *A class $\mathcal{C}$ of structures is proficient, if and only if, there is some R-positive formula φ such that* $\sup(\{||\varphi||_A \mid A \in \mathcal{C}\}) \geq \omega$.

Since the stages of a first order induction are uniformly first order definable, it is easy to see that if a class of structures C is not proficient, then FO+LFP collapses to FO over C. McColm [10] conjectured that the converse also holds.

Conjecture 1 (McColm [10]) *If C is proficient, then FO+LFP does not collapse to FO on C.*

While McColm's Conjecture was refuted by Gurevich, Immerman, and Shelah [6], the question whether or not FO+LFP collapses to FO on various collections of proficient classes of structures remains of considerable interest. For example, as remarked above, the status of the Ordered Conjecture posed by Kolaitis and Vardi [8] is a major open problem in the field.

The following theorem provides a *positive* resolution of McColm's Conjecture with respect to an extensive collection of classes of finite structures. Indeed, this result provides a uniform explanation for many of the positive results concerning McColm's Conjecture which have heretofore been treated by heterogeneous methods in the literature.

Theorem 1 *Let C be a proficient class of finite structures. If* $\mathrm{Th}(\mathsf{C})$ *is decidable, then* $\mathrm{FO}(\mathsf{C}) \neq \mathrm{FO{+}LFP}(\mathsf{C})$.

We will obtain Theorem 1 as a corollary of a result of McColm [10]. Before proceeding to the proof, note that Theorem 1 applies to the classic examples of strings (colored finite linear orderings) and trees (two finite successor functions), both of which are proficient classes of finite structures with decidable first-order theories (in fact, their monadic second-order theories are decidable). Further examples include: the class $\{\langle\{0, \ldots, n\}, +\rangle \mid n \in \omega\}$ (since Presburger arithmetic is decidable); the class of finite fields and the class of finite boolean algebras. For details of these and additional examples see [5], [11], and [12]. We note for the purposes of our discussion below the following obvious corollary of Theorem 1 which extends its range of application.

Corollary 1 *Let $\mathcal{C}$ be a proficient class of finite structures with a decidable first–order theory. If $\mathcal{C} \subseteq \mathcal{C}'$, then* $\mathrm{FO}(\mathsf{C}') \neq \mathrm{FO{+}LFP}(\mathsf{C}')$.

We will derive Theorem 1 from the following result of McColm [10]

Theorem 2 (McColm [10]) *If C is recursively enumerable and* $\mathrm{Th}^{\infty}(\mathsf{C})$ *is complete, then C is proficient if and only if* $\mathrm{FO}(\mathsf{C}) \neq \mathrm{FO{+}LFP}(\mathsf{C})$.

Theorem 1 is a corollary of Theorem 2 and Lemma 1 below. The following proposition will be used in the proof of Lemma 1.

Proposition 1 *If* $\mathrm{Th}(\mathsf{C})$ *is decidable, then C is decidable.*

Proof: This is an easy corollary of the fact that for every finite structure A, we can effectively construct a first order sentence φ_A such that for every structure B,

$$A \cong B \iff B \models \varphi_A.$$

It follows that for every collection of finite structures C and for every finite structure A,

$$A \in \mathsf{C} \iff \neg\varphi_A \notin \mathrm{Th}(\mathsf{C}).$$

■

Lemma 1 *Let C be a proficient class of finite structures such that $\mathrm{Th}(\mathsf{C})$ is decidable. Then, there is a $\mathsf{D} \subseteq \mathsf{C}$, such that*

1. *D is recursively enumerable,*
2. *D is proficient, and*
3. *$\mathrm{Th}^{\infty}(D)$ is complete.*

Proof: Let C be a proficient class of finite structures such that $\mathrm{Th}(\mathsf{C})$ is decidable. Suppose the R-positive first order formula φ witnesses the proficiency of C. For each n, let θ_n the first order formula $\exists \overline{x}(\varphi^{n+1}(\overline{x}) \wedge \neg\varphi^{n}(\overline{x}))$. Note that for all structures A, θ_n satisfies the following condition:

$$A \models \theta_n \iff ||\varphi||_A > n.$$

Let ψ_n be an effective enumeration of all first order sentences and let A_n be an effective enumeration of C (observe that Proposition 1 implies the decidability of C). We proceed to effectively enumerate a sequence of structures B_n to satisfy the conditions of the lemma. The idea will be to effectively search the refinement tree of finite extensions of $\mathrm{Th}(\mathsf{C})$ for the leftmost infinite branch consistent with $\{\theta_n \mid n \in \omega\}$. This will lead to an effective enumeration of structures satisfying the conditions of the lemma. To each binary sequence σ of length n, we associate a conjunction χ_σ of the sentences $\psi_1, \ldots, \psi_n$ or their negation: if $\sigma(i) = 0$, then ψ_i is a conjunct, and if $\sigma(i) = 1$, then $\neg\psi_i$ is a conjunct. Now, we call a sequence σ of length $\leq n$ live at stage n of our construction, if $\chi_\sigma \wedge \theta_n$ is consistent with $\mathrm{Th}(\mathsf{C})$. Note that the relation "σ is live at stage n" is decidable, since $\mathrm{Th}(\mathsf{C})$ is decidable. Consider the tree T_n of all live sequences at stage n and let σ_n be its leftmost leaf on a path of maximal length. Note that the heights of the trees T_n are strictly increasing as a function of n, since φ witnesses the proficiency of C. We choose B_n to be the first structure in our enumeration of C which satisfies $\chi_{\sigma_n} \wedge \theta_n$.

It follows directly from the construction that $\mathsf{D} = \{B_n | n \in \omega\}$ satisfies conditions 1 and 2 of the lemma, indeed, φ witnesses the proficiency of D. In order to verify condition 3, we argue as follows. We write $(\sigma \mid n)$ for the initial segment of the sequence σ of lenth n and $\sigma * b$ for the extension of σ by b. It suffices to show that for every n, there is an m such that for every $m' \geq m, (\sigma_{m'} \mid n) = (\sigma_m \mid n)$. We proceed to prove this by induction on n. Suppose, as induction hypothesis, that for every $m' \geq m, (\sigma_{m'} \mid n) = (\sigma_m \mid n)$. Now, suppose that for every $m' \geq m$, $\chi_{(\sigma_m|n)} \wedge \psi_{n+1} \wedge \theta_{m'}$ is consistent with $\mathrm{Th}(\mathsf{C})$. Then, for every $m' \geq m+1, (\sigma_{m'} \mid n+1) = (\sigma_m \mid n) * 0$. Otherwise, for some $k \geq m, \chi_{(\sigma_m|n)} \wedge \psi_{n+1} \wedge \theta_k$ is not consistent with $\mathrm{Th}(\mathsf{C})$. In this case, for every $m' \geq k+1, (\sigma_{m'} \mid n+1) = (\sigma_m \mid n) * 1$. ■

The following proposition indicates that the range of application of Theorem 2 is wider than that of Corollary 1.

Proposition 2 *There is an infinite recursive set of finite structures C with the following properties:*

1. *C is proficient,*
2. *$\mathrm{Th}^{\infty}(\mathsf{C})$ is complete, and*
3. *if $\mathsf{B} \subseteq \mathsf{C}$ is infinite, then $\mathrm{Th}(\mathsf{B})$ is undecidable.*

Proof: We begin by introducing an encoding of finite sets of natural numbers into binary string structures. Let $D = \{n_1, \ldots, n_k\}$ where $n_1 < n_2 < \ldots < n_k$. We code D by a structure A_D. The signature of A_D consists of a binary relation S and a unary predicate P. Let $t = (2k-1) + \sum_{i=1}^{k} n_i$. The universe of A_D is the set $\{1, \ldots t\}$. S is interpreted as the successor relation on the universe and P holds of j unless $j = 2l + \sum_{i=1}^{l} n_i$, for some $1 \leq l < k$. Intuitively, A_D codes D as a tape whose entries are the unary representations of the elements of D in increasing order separated by blanks.

Let $X = \{D_i \mid i \in \omega\}$ be an increasing chain of finite subsets of ω, that is, for all $i, D_i \subseteq D_{i+1}$. We write C_X for $\{A_{D_i} \mid i \in \omega\}$. Observe that since X is a chain, for every m, there is an n such that for every $n' \geq n, D_n \cap \{0, \ldots, m\} = D_{n'} \cap \{0, \ldots, m\}$. It follows from this observation and the Hanf locality theorem (see Fagin, Stockmeyer, & Vardi [4]) that for every r there is an n such that for every $n' \geq n, A_{D_n} \equiv_r A_{D_{n'}}$. It follows at once that $\mathrm{Th}^{\infty}(\mathsf{C}_X)$ is complete. Moreover, it is easy to see that if X is infinite, then C_X is proficient, for example, the induction which defines the transitive closure of S witnesses proficiency.

Let $f : \omega \mapsto \omega$, and let $D(f, n) = \{f(0), \ldots, f(n)\}$. Let $X_f = \{D(f, n) \mid n \in \omega\}$. Clearly, X_f is a chain of finite subsets of ω. Now, suppose f is an injective recursive function whose range is not recursive. We claim that $\mathsf{C} = \mathsf{C}_{X_f}$ satisfies the conditions of the proposition. It is clear from the preceding argument that C is proficient and $\mathrm{Th}^{\infty}(\mathsf{C})$ is complete. It is easy to verify that C is recursively enumerable in increasing order (under a suitable effective encoding), since f is injective and recursive. It follows at once that C is infinite and recursive. It is also easy to verify that for any infinite $\mathsf{B} \subseteq \mathsf{C}$, the complement of the range of f is 1-1 reducible to $\mathrm{Th}(\mathsf{B})$. It follows at once that $\mathrm{Th}(\mathsf{B})$ is undecidable. ■

4. Trivialization and Separability

Intuitively, when $\mathrm{Th}^{\infty}(\mathsf{C})$ is complete, FO is trivial over C, that is, every first order boolean query over C is either finite or cofinite.

Definition 2 1. *A boolean query Q is trivial on $\mathcal{C}$ if either $Q \cap \mathcal{C}$ or $\mathcal{C} - Q$ is finite.*

2. $\mathcal{C}$ trivializes *a collection of queries $\mathcal{Q}$, if and only if, every $Q \in \mathcal{Q}$ is trivial on $\mathcal{C}$.*

3. *A collection of queries $\mathcal{Q}$ is* effectively trivializable *over C, if and only if, there is an infinite r.e. set of structures $\mathsf{C}' \subseteq \mathcal{C}$ which trivializes it.*

One popular technique in finite model theory for establishing that a given logic L does not collapse to FO over a given class of structures C, is to show that some subcollection of C trivializes FO while some L-query is nontrivial over that subcollection. The following theorem shows that this technique is universally applicable.

Definition 3 *C and D are first order inseparable, if and only if, there is no first order sentence φ such that $\mathsf{C} \subseteq \mathrm{Mod}(\varphi)$ and $\mathsf{D} \cap \mathrm{Mod}(\varphi) = \emptyset$.*

Note that a query Q is not first order definable over a collection of structures C, if and only if, $Q \cap \mathsf{C}$ and $\mathsf{C} - Q$ are first order inseparable.

Theorem 3 *The following conditions are equivalent.*

1. *C and D are first order inseparable.*

2. *There are sequences $\{C_i \mid i \in \omega\} \subseteq \mathsf{C}$ and $\{D_i \mid i \in \omega\} \subseteq \mathsf{D}$ such that for all $n \in \omega, C_n \equiv_n D_n$. Moreover, these sequences may be chosen so that for all $m, n \in \omega, n \leq m \rightarrow C_n \equiv_n C_m$.*

Proof: Since our languages contain no function symbols, for each n, the equivalence relation $\equiv_n$ has finite index, and each of its equivalence classes is defined by a single sentence of quantifier rank n. We say an equivalence class of $\equiv_n$ is good, just in case its intersection with both C and D is nonempty. Suppose that for some n, $\equiv_n$ has no good equivalence classes. Let θ be the disjunction of the sentences defining equivalence classes which have nonempty intersection with C. Then, $\mathrm{Mod}(\theta)$ separates C and D. So, for every n, $\equiv_n$ has good equivalence classes. We form a directed tree T from the good equivalence classes of the $\equiv_n$ for $n \in \omega$ as follows: the nodes of the tree are pairs $\langle X, n\rangle$ where X is a good equivalence of $\equiv_n$; $\langle Y, m\rangle$ is a child of $\langle X, n\rangle$, if and only if $m = n + 1$ and $Y \subseteq X$. Observe that T is an infinite, finitely branching tree. Hence, by König's Lemma, T has an infinite path $\{\langle X_n, n\rangle \mid n \in \omega\}$. We may now pick $C_n \in (\mathsf{C} \cap X_n)$ and $D_n \in (\mathsf{D} \cap X_n)$ to satisfy the conditions of the theorem. ■

Corollary 2 *A boolean query Q is not first order definable over a class of structures C, if and only if, there is a $\mathsf{D} \subseteq \mathsf{C}$ such that $\mathrm{Th}^{\infty}(\mathsf{D})$ is complete (*i.e. *D trivializes* FO*), but Q is not trivial on D.*

The preceding corollary suggests that it would be useful to be able to *construct* classes of structures that trivialize FO in order to prove separations of descriptive complexity classes from FO. If such constructions are effective, that is, give rise to effective trivializations of FO, then the applicability of this method is severely limited with regard to complexity theoretic separations, as the following results show.

Theorem 4 *There is* no *infinite r.e. class of finite structures $\mathcal{C}$ such that $\mathcal{C}$ trivializes* DLOGTIME.

Proof: The idea will be to recursively define the computation of a boolean query which alternates its truth values infinitely often over C. Let M be a machine which generates, in tally notation, the size of structures in $\mathcal{C}$, writing over its previous outputs. Let $M(n)$ denote the length of the maximum complete output of M after n steps (*n.b.* that $M(n) < n$). Define the spectral query Q on $\mathcal{C}$, which depends only on size, as follows (here we identify Q with its characteristic function). Let

$$Q(|A|) = \begin{cases} 0 & \text{if } M(log|A|) \text{ is } \emptyset \\ 1 - Q(M(\log|A|)) & \text{otherwise} \end{cases}$$

Observe that Q runs in time $O(\log|A|)$. Now suppose that Q is constant for all $|A| > n'$. Then we know that there is an m' such that $M(m) > n'$ for all $m > m'$. Now simply consider $A' \in \mathsf{C}$ with $|A'| > 2^{m'}$. By construction, $Q(|A'|) \neq Q(M(\log|A'|))$, which contradicts that Q was constant beyond n'. ■

It is well-known that on classes of structures equipped with the "built-in" arithmetic relations BIT and $<$ first–order logic suffices to express all DLOGTIME computable boolean queries, a result due to Barrington, Immerman, and Straubing [2]. In the following theorem, LH refers to the LOGTIME hierarchy, the bottom level of which is DLOGTIME.

Theorem 5 (Barrington *et al.* [2]) $\mathrm{LH} = \mathrm{FO}(\mathrm{BIT}, <)$.

In light of Theorem 5 and the above discussion, Theorem 4 has the following significance. It shows that Theorem 2 *cannot* be used to separate P from FO over classes of structures which have the resources, either explicit or implicit, to interpret enough arithmetic. This holds, in particular, for the "bare" class $\{\langle n, <, \mathrm{BIT}\rangle \mid n \in N\}$. In addition, it shows that Corollary 2 *cannot* be effectivized, that is, the appealing separation technique enshrined in the corollary is no longer universal, if requirements of effectivity are imposed.

Though Theorem 4 does establish a sharp limitation on the use of effective trivialization as a separation technique in descriptive complexity theory, it does not preclude the possibility that more general constructive pebble game techniques could succeed in effecting separations from descriptive complexity classes containing DLOGTIME. The following proposition illustrates this possibility. Let $B_n = \langle n, <, \mathrm{BIT}\rangle$ and let $\mathsf{B} = \{B_n \mid n \in N\}$.

Proposition 3 *There are disjoint infinite recursive sets of finite structures* $\mathsf{C}, \mathsf{D} \subseteq \mathsf{B}$ *with recursive enumerations* $\mathsf{C} = \{C_i \mid i \in N\}$ *and* $\mathsf{D} = \{D_i \mid i \in N\}$ *such that for all* $n \in N, C_n \equiv_n D_n$.

Proof: Construct sequences C_i and D_i simultaneously by stages as follows. At stage n let (k, l) be the least pair of natural numbers in some standard enumeration of pairs such that $k < l, B_k \equiv_n B_l$, and neither B_k nor B_l appear earlier in the construction. (Note that at each stage of the construction there are infinitely many such pairs (k, l), since $\equiv_n$ is an equivalence relation of finite index.) Let $C_n = B_k$ and $D_n = B_l$. It is easy to see that $\mathsf{C} = \{C_i \mid i \in N\}$ and $\mathsf{D} = \{D_i \mid i \in N\}$ so constructed satisfy the requisite conditions. ■

By Theorem 3, the query C constructed in Proposition 3 is not first order definable over B, and this fact can be witnessed by a "constructive pebble game argument" as indicated. On the other hand, C does *not* satisfy the "monotonicity" condition stated in the last sentence of Theorem 3 - if it did, C would trivialize FO, but, as noted above, no infinite recursively enumerable subset of B trivializes FO.

5. Conclusions

In this paper, we have established a useful decidability condition which guarantees the separation of FO+LFP from FO on a given collection of finite structures. We have also investigated the limitations on certain constructive proof techniques to establish separations under circumstances in which this sufficient condition fails to apply.

References

[1] A. Atserias and P. G. Kolaitis. First-order logic vs. fixed-point logic in finite set theory. In *Proc. 14th IEEE Symp. on Logic in Computer Science*, 1999.

[2] D. Barrington, N. Immerman, and H. Straubing. On uniformity within NC^1. *Journal of Computer and System Sciences*, 41:274–306, 1990.

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

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

[5] Y. Gurevich. Monadic second-order theories. In J. Barwise and S. Feferman, editors, *Model-Theoretic Logics*, pages 479–506. Springer-Verlag, New York, 1985.

[6] Y. Gurevich, N. Immerman, and S. Shelah. Mccolm's conjecture. In *Proceedings of the 9th IEEE Symposium on Logic in Computer Science*, pages 10–19, 1994.

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

[8] P. G. Kolaitis and M. Y. Vardi. Fixpoint logic vs. infinitary logic in finite-model theory. In *Proceedings of the 7th IEEE Symposium on Logic in Computer Science*, pages 46–57, 1992.

[9] S. Lindell. An analysis of fixed-point queries on binary trees. *Theoretical Computer Science*, 85(1):75–95, 1991.

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

[11] M. Rabin. Decidable theories. In J. Barwise, editor, *Handbook of Mathematical Logic*, pages 595–630. North Holland, New York, 1977.

[12] W. Thomas. Automata on infinite objects. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science B*, pages 133–192. Elsevier, New York, 1990.

[13] 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.