Fixpoint Logics, Relational Machines, and Computational Complexity

Serge Abiteboul
INRIA, BP 105
78153 Le Chesnay CEDEX
France
abitebou@inria.inria.fr

Moshe Y. Vardi
IBM Almaden Research Center
San Jose, CA 95120-6099
USA
vardi@almaden.ibm.com

Victor Vianu
CSE C-0114 UC San Diego
La Jolla, CA 92093
USA
vianu@cs.ucsd.edu

November 24, 1993

Abstract

We establish a general connection between fixpoint logic and complexity. On one side, we have fixpoint logic, parameterized by the choices of 1st-order operators (inflationary or noninflationary) and iteration constructs (deterministic, nondeterministic, or alternating). On the other side, we have the complexity classes between P and EXPTIME. Our parameterized fixpoint logics capture the complexity classes P, NP, PSPACE, and EXPTIME, but equality is achieved only over ordered structures.

There is, however, an inherent mismatch between complexity and logic – while computational devices work on encodings of problems, logic is applied directly to the underlying mathematical structures. To overcome this mismatch, we develop a theory of relational complexity, which bridges tha gap between standard complexity and fixpoint logic. On one hand, we show that questions about containments among standard complexity classes can be translated to questions about containments among relational complexity classes. On the other hand, the expressive power of fixpoint logic can be precisely characterized in terms of relational complexity classes. This tight three-way relationship among fixpoint logics, relational complexity and standard complexity yields in a uniform way logical analogs to all containments among the complexity classes P, NP, PSPACE, and EXPTIME. The logical formulation shows that some of the most tantalizing questions in complexity theory boil down to a single question: the relative power of inflationary vs. noninflationary 1st-order operators.

1 Introduction

The *computational complexity* of a problem is the amount of resources, such as time or space, required by a machine that solves the problem. Complexity theory traditionally has

focused on the computational complexity of problems. A more recent branch of complexity theory, started by Fagin in [Fag74, Fag75] and developed during the 1980s, focuses on the *descriptive complexity* of problems, which is the complexity of describing problems in some logical formalism [Imm87a]. One of the exciting developments in complexity theory is the discovery of a very intimate connection between computational and descriptive complexity.

This intimate connection was first discovered by Fagin, who showed that the complexity class NP coincides with the class of properties expressible in existential 2nd-order logic [Fag74] (cf. [JS74]). Another demonstration of this connection was shown by Immerman and Vardi, who discovered tight relationships between the complexity class P and inflationary fixpoint logic [Imm86, Var82] and between the class PSPACE and noninflationary fixpoint logic [Var82] (cf. [AV89]).[1] The tight connection between descriptive and computational complexity, typically referred to as the connection between "logic and complexity", was then proclaimed by Immerman [Imm87b], and studied by many researchers [Com88, Goe89, Gra84, Gra85, Gur83, Gur84, Gur88, HP84, Imm89, Lei89a, Liv82, Liv83, Lyn82, Saz80b, Saz80a, TU88].[2] See [Imm89] for a survey.

Although the relationship between descriptive and computational complexity is intimate, it is not without its problems, and the partners do have some irreconcilable differences. While computational devices work on *encodings* of problems, logic is applied directly to the underlying mathematical structures. As a result, machines are able to enumerate objects that are logically unordered. For example, while we typically think of the set of nodes in a graph as unordered, it does become ordered when it is encoded on the tape of a Turing machine. This "impedance mismatch" does not pose any difficulty in the identification of NP with existential 2nd-order logic (in [Fag74]), since the logic can simply assert the existence of the desired order. The relationship between the class P and inflationary fixpoint logic is, however, more complicated as a result of the mismatch. Although inflationary fixpoint logic can describe P-complete problems, there are some very easy problems in P that are not expressible in inflationary fixpoint logic (e.g., checking whether the cardinality of the structure is even [CH82]). It is only when we assume built-in order that we get that P coincides with the class of properties expressible in inflationary fixpoint logic [Imm86, Var82]. Similarly, it is only when we assume a built-in order that we get that PSPACE coincides with the class of properties expressible in noninflationary fixpoint logic [Var82].

A consequence of Fagin's result is that NP=co-NP if and only if existential and universal 2nd-order logic have the same expressive power. This equivalence of questions in computational and descriptive complexity is one of the major features of the connection between the two branches of complexity theory. It holds the promise that techniques from one domain could be brought to bear on questions in the other domain.

Unfortunately, the order issue complicates matters. The results by Immerman and Vardi show that P=PSPACE if and only if inflationary and noninflationary fixpoint logics have the same expressive power over *ordered* structures. We would like, however, to eliminate the restriction to ordered structures for two reasons. The first reason is technical; questions about logical expressiveness over ordered structures are typically much harder than their unordered counterparts; see for example [dR84]. The second reason is more fundamental; the restriction to ordered structures seems to be a technical device rather than an intrinsic feature. Because of this restriction, the above equivalence provides a translation of the P vs. PSPACE question to a descriptive complexity question, but does not provide a translation of the inflationary vs. noninflationary question to a computational complexity question. Thus, we would like to obtain an order-free correspondence between questions in computational

[1] Inflationary and noninflationary fixpoint logics are defined in Section 2.

[2] The focus here is on the connection between finite-model theory and complexity. The connection between logic and complexity has also a proof-theoretic aspect; see [Bus86, GSS90, Lei91].

and descriptive complexity. (See also [IL90] for a discussion of the order issue ¿from another perspective.)

The order issue was partially overcome recently by Abiteboul and Vianu [AV91], who showed that indeed P=PSPACE if and only if inflationary and noninflationary fixpoint logics have the same expressive power. The crux of their result is a clever description of a certain "internal" order in inflationary fixpoint logic. While this order is much weaker than the total order available to computational devices, Abiteboul and Vianu showed that it is sufficient to enable the translation of the P=PSPACE question into a logical one.

Although this result goes a long way in overcoming the order issue, it is not completely satisfying. We now have two very different ways to relate descriptive and computational complexity: Fagin's result, extended later to a correspondence between the polynomial hierarchy and 2nd-order logic [Sto77], enables us to phrase questions about the polynomial hierarchy in terms of second-order logic, and Abiteboul and Vianu's result enables us to phrase the P vs. PSPACE question in terms of fixpoint logic. Because the two formalisms are so different, it is not clear whether they can be combined to yield a descriptive-complexity analog to the P=NP question.

Our goal in this paper is extend the results and techniques of Abiteboul and Vianu into a general connection between fixpoint logics and complexity classes between P and EXPTIME. Our results are based on several fundamental ideas. At the heart of our framework is the view of fixpoint logic as 1st-order logic augmented with iteration. While traditionally fixpoint logic was viewed as the extension of 1st-order logic by recursion (cf. [Mos74]), iteration proved to be a more general extension to 1st-order logic than recursion [AV89, GS86, Lei90]. In both inflationary and noninflationary fixpoint logics, iteration is applied in its simplest form: sequential and deterministic. It turns out, however, that in order to express certain problems in fixpoint logic one seems to require more elaborate forms of iteration, such as nondeterministic or alternating. Indeed, part of this research was motivated by the question whether PSPACE-complete problems such as "quantified Boolean formulas" or "nonuniversality for finite automata" [GJ79] can be described in noninflationary fixpoint logic. By augmenting fixpoint logic with nondeterministic or alternating iteration we find that the connection between fixpoint logic and computational complexity is quite broad; it covers not only the classes P and PSPACE, but also the classes NP and EXPTIME. For example, over ordered structures NP coincides with the class of properties expressible in nondeterministic inflationary fixpoint.

This broad connection between fixpoint logic and computational complexity still suffers from the order mismatch discussed above. To get around this difficulty, we develop a theory of *relational* complexity, which bridges the gap beween standard complexity and fixpoint logic. In relational complexity theory, we model computations on unordered structures by *relational machines*. Relational machines are Turing machines that are augmented by a relational store and the ability to perform relational operations on that store. This idea of extending 1st-order logic with a general computational capability was suggested in [CH80] and pursued in [AV91].[3] Unlike Turing machines, which operate on encodings of problems, relational machines operate directly on the underlying mathematical structures.

For Turing machines, the most natural measure of complexity is in terms of the size of the input. This measure is not the most natural one for relational machines, since such machines cannot measure the size of their input. In fact, relational machines have a limited *discerning* power, i.e., the power to distinguish between different pieces of their input, since they are only able to manipulate their input relationally. The discerning power of relational machines is best understood by viewing them as an effective fragment of a certain infinitary

[3] A closely related idea, of generalizing Turing machines to operate on general structures, goes back to [Fri71] and was investigated extensively in [Lei89a, Lei89b].

logic, studied recently in [KV90a, KV90b]. This view yields a precise characterization of the discerning power of relational machines in terms of certain infinite 2-player pebble games. The characterization can then be used by relational machines in order to determine their own discerning power. This suggests that it is natural to measure the size of the input to a relational machine with respect to the discerning power of the machine. The new measure gives rise to a new notion of computational complexity, which we call *relational complexity*, resulting in classes such as P_r (relational polynomial time) and NP_r (relational nondeterministic polynomial time).

Relational complexity can now serve as a mediator between logical complexity and standard complexity. On one hand, we show that questions about containments among standard complexity classes can be translated to questions about containments among relational complexity classes. On the other hand, the expressive power of fixpoint logic can be *precisely* characterized in terms of relational complexity classes. This tight three-way relationship among fixpoint logics, relational complexity and standard complexity yields in a uniform way logical analogs to all containments among the complexity classes P, NP, PSPACE, and EXPTIME. It also enables us to translate known relationships among complexity classes, such as the equality of PSPACE, NPSPACE, and APSPACE, into results about the expressive power of fixpoint logics. This fulfills the promise of applying results from one domain to another domain and shows that some of the most tantalizing questions in complexity theory – P. vs. PSPACE, NP vs. PSPACE, and PSPACE vs. EXPTIME – boil down to *one fundamental issue*: understanding the relative power of inflationary vs. noninflationary 1st-order operators.

2 Fixpoint Logics

Let $\varphi(x_1, \ldots, x_n, S)$ be a 1st-order formula in which S is a n-ary relation symbol (not included in a vocabulary σ) and let $\mathbf{D}$ be a structure over the vocabulary σ. The formula φ gives rise to an operator $\Phi(S)$ from n-ary relations on the universe D of $\mathbf{D}$ to n-ary relations on D, where $\Phi(T) = \{(a_1, \ldots, a_n) : \mathbf{D} \models \varphi(a_1, \ldots, a_n, T)\}$, for every n-ary relation T on D.

Every such operator $\Phi(S)$ generates a sequence of *stages* that are obtained by iterating $\Phi(S)$. The stages Φ^m (also denoted by φ^m), $m \geq 1$, of Φ on $\mathbf{D}$, are defined by the induction: $\Phi^1 = \Phi(\emptyset)$, $\Phi^{m+1} = \Phi(\Phi^m)$. Intuitively, one would like to associate with an operator $\Phi(S)$ the "limit" of its stages. This is possible only when the sequence Φ^m, $m \geq 1$, of the stages "converges", i.e., when there is an integer m_0 such that $\Phi^{m_0} = \Phi^{m_0+1}$ and, hence, $\Phi^{m_0} = \Phi^m$, for all $m \geq m_0$. Notice that in this case Φ^{m_0} is a fixpoint of $\Phi(S)$, since $\Phi^{m_0} = \Phi^{m_0+1} = \Phi(\Phi^{m_0})$. The sequence of stages may not converge. In particular, this will happen if the formula $\varphi(\mathbf{x}, S)$ has no fixpoints.

2.1 Inflationary Fixpoint Logic

A formula $\varphi(x_1, \ldots, x_n, S)$ is *inflationary in* S if $T \subseteq \Phi(T)$ for any n-ary relation T. In particular, φ is inflationary in S if it is of the form $S(x_1, \ldots, x_n) \vee \psi(x_1, \ldots, x_n, S)$. If φ is inflationary in S, then the sequence Φ^m, $m \geq 1$, of stages is increasing. Thus, it must have a "limit". More precisely, if $\mathbf{D}$ is a finite structure with s elements, then there is an integer $m_0 \leq s^n$ such that $\Phi^{m_0} = \Phi^m$ for every $m \geq m_0$. That is, the sequence of stages of $\varphi(\mathbf{x}, S)$ converges to Φ^{m_0}. We write φ^∞ or Φ^∞ to denote the fixpoint Φ^{m_0} of φ. *Inflationary fixpoint logic* (IFP) is 1st-order logic augmented with the inflationary fixpoint formation rule for inflationary formulas. The canonical example of a formula of inflationary fixpoint logic is provided by the inflationary fixpoint $\varphi^\infty(x, y)$ of the 1st-order formula

$$S(x, y) \vee E(x, y) \vee (\exists z)(S(x, z) \wedge S(z, y)).$$

In this case $\varphi^\infty(x,y)$ defines the *transitive closure* TC of the *edge* relation E. It follows that *connectivity* is a property expressible in inflationary fixpoint logic, but, as is well known (cf. [Fag75, AU79]), not in 1st-order logic.

Remark 2.1: Fixpoints can also be guaranteed to exist when the formula $\varphi(x_1,\ldots,x_n,S)$ is *positive in* S, namely, when every occurrence of S is governed by an even number of negations. In that case, the sequence Φ^m, $m \geq 1$, of stages is also increasing, and consequently has a limit that is a fixpoint. In fact, that limit is the least fixpoint of φ. *Positive fixpoint logic* is 1st-order logic augmented with the least fixpoint formation rule for positive formulas. It is easy to see that IFP is at least as expressive as positive fixpoint logic. Gurevich and Shelah showed that in fact the two logics have the same expressive power [GS86] (see also [Lei90]). ∎

The complexity-theoretic aspects of IFP were studied in [CH82, Imm86, Var82]. (These papers actually focused on positive fixpoint logic, but, as observed above, positive fixpoint logic and IFP have the same expressive power.) It is known that IFP captures the complexity class P, where a logic L is said to *capture* a complexity class C if the following holds:

1. All problems expressible in L are in C, and
2. on ordered structures L can express all problems in C.

(A more natural definition would be to require that a problem is in C iff it is expressible in L. Because of the order mismatch, this requirement would be too stringent.) Note, however, that there are problems in P (e.g., checking whether the cardinality of the structure is even) that are not expressible in IFP [CH82].

2.2 Noninflationary Fixpoint Logic

How can we obtain logics with iteration constructs that are more expressive than inflationary fixpoint logic? A more powerful logic results if one iterates general 1st-order operators, until a fixpoint is reached (which may never happen). In this case we may have non-terminating computations, unlike inflationary fixpoint logic, where the iteration was guaranteed to converge. Let Φ^m, $m \geq 1$, be the sequence of stages of the operator $\Phi(S)$ associated with a formula $\varphi(x_1,\ldots,x_n,S)$. If there is an integer m_0 such that $\Phi^{m_0} = \Phi^{m_0+1}$, then we put $\varphi^\infty = \Phi^\infty = \Phi^{m_0}$; otherwise, we set $\varphi^\infty = \Phi^\infty = \emptyset$.[4] We call φ^∞ the *noninflationary fixpoint* of φ on $\mathbf{D}$. *Noninflationary Fixpoint Logic* (NFP) is 1st-order logic augmented with the noninflationary fixpoint formation rule for arbitrary 1st-order formulas. Note that NFP fixpoint logic is an extension of IFP. While IFP formulas can be evaluated in polynomial time, NFP can express PSPACE-complete problems, such as the *network convergence* problem [Fed81].

Noninflationary fixpoint logic was introduced by Abiteboul and Vianu [AV89] (who called it *partial* fixpoint logic). In particular, they observed that noninflationary fixpoint logic coincides with the query language RQL introduced in [CH82] and studied further in [Var82]. It follows that NFP captures the complexity class PSPACE [Var82], but the problem of even cardinality is not expressible in NFP [CH82].

Clearly, if IFP and NFP have the same expressive power, then P=PSPACE. It is not clear, however, that the converse holds because IFP and NFP need order to "fully capture" P and PSPACE, respectively. Nevertheless, Abiteboul and Vianu showed that P=PSPACE if and only if IFP and NFP have the same expressive power [AV91].

[4] Actually, it is shown in [AV90] that there is no loss of generality in restricting attention to converging 1st-order operators.

2.3 More Fixpoint Logics

IFP and NFP are obtained by iterating inflationary and noninflationary 1st-order operators, respectively. In both cases, however, the iteration is sequential and deterministic. Certain problems, however, seem to defy description by such iterations.

As an example, consider nonuniversality of finite automata over a binary alphabet, which is known to be PSPACE-complete [GJ79]. That is, we are given a finite automaton over the alphabet $\{0,1\}$ and we want to known whether there is a word rejected by the automaton. An instance of this problem can be viewed as a structure $(S, S_0, F, \rho_0, \rho_1)$, where S is the set of states, S_0 is the set of initial states, F is the set of accepting states, ρ_0 is the transition relation for the symbol 0 and ρ_1 is the transition relation for the symbol 1. To check whether there is a word rejected by the automaton, one can simply guess a word and check that it is rejected by the automaton. This can be easily described by a *nondeterministic iteration* of 1st-order operators. Let us define what nondeterministic iteration is.

Let Φ_1 and Φ_2 be 1st-order operators. This pair of operators generates *convergent sequences* of stages that are obtained by successively applying, till convergence is reached, Φ_1 or Φ_2. That is, the pair generates sequences of the form $S_0, S_1, \ldots, S_m$, where $S_0 = \emptyset$, either $S_{i+1} = \Phi_1(S_i)$ or $S_{i+1} = \Phi_2(S_i)$, and $\Phi_1(S_m) = \Phi_2(S_m) = S_m$. We call S_m a *local nondeterministic fixpoint* of the pair Φ_1, Φ_2. Note that the pair Φ_1, Φ_2 can have more than one local nondeterministic fixpoint or none. We define the *nondeterministic fixpoint* of the pair Φ_1, Φ_2 as the union of all local nondeterministic fixpoints of the pair Φ_1, Φ_2. If no local nondeterministic fixpoint exists, then we define the nondeterministic fixpoint to be the empty set.

Nondeterministic fixpoint logic are obtained by augmenting 1st-order logic with the nondeterministic fixpoint formation rule, under the restriction that negation cannot be applied to nondeterministic fixpoints. (This kind of nondeterminism should be contrasted with the *data nondeterminism* of [AV89], where we can nondeterministically choose an element of the underlying domain.)

The readers can now convince themselves that the nonuniversality problem can be expressed by a nondeterministic fixpoint of the noninflationary operators Φ_1 and Φ_2 corresponding to the next input symbol being 0 or 1. The problem, however, does not seem to be expressible by a deterministic iteration. Savitch's Theorem [Sav80] tells us how to convert the nondeterministic polynomial-space algorithm into a deterministic polynomial-space algorithm, but this construction assume a built-in order and does not seem to be describable by a deterministic iteration.

For an example that motivates another extension of fixpoint logic, consider truth of quantified Boolean formulas, also known to be PSPACE-complete [GJ79]. This problem, which can be trivially solved by an alternating polynomial algorithm [CKS81], seems to require an *alternating iteration* of 1st-order operators. Let us define what alternating iteration is.

Let Φ_1 and Φ_2 be 1st-order operators. This pair of operators generates *convergent trees* of stages that are obtained by successively applying, till convergence is reached, either one of Φ_1 and Φ_2 or both of Φ_1 and Φ_2. More formally, a convergent tree is a labeled binary tree such that:

1. The root is labeled by the empty relation,
2. if a node x with label S_x is at an odd level of the tree, the x has one child x' labeled by $\Phi_1(S_x)$ or $\Phi_2(S_x)$,
3. if a node x with label S_x is at an even level of the tree, the x has two children x_1 and x_2 labeled by $\Phi_1(S_x)$ and $\Phi_2(S_x)$, respectively, and
4. if x is a leaf with label S_x, then $\Phi_1(S_x) = \Phi_2(S_x) = S_x$.

We take the intersection of the labels of the leaves of a convergent stage tree to be a *local alternating fixpoint* of the pair Φ_1, Φ_2. Note that the pair Φ_1, Φ_2 can have more than one local alternating fixpoint or none. We define the *alternating fixpoint* of the pair Φ_1, Φ_2 as the union of all local alternating fixpoints of the pair Φ_1, Φ_2. If no local alternating fixpoint exists, then we define the nondeterministic fixpoint to be the empty set.

The discussion so far shows that fixpoint logics can be parameterized along two dimensions: the power of their iteration construct, *d*eterministic vs. *n*ondeterministic vs. *a*lternating, and the power of their 1st-order operators, *i*nflationary vs. *n*oninflationary. This gives rise to six fixpoint logics. We use the notation $FP(\alpha, \beta)$ to refer to a fixpoint logic with iteration of type α and operators of type β. Thus, the logics IFP and NFP will be denoted $FP(d, i)$ and $FP(d, n)$, respectively, and $FP(a, n)$ denotes alternating noninflationary fixpoint logic.

The obvious containments among these logics are described by the following diagram:

$$FP(a,i) \qquad FP(a,n)$$

$$FP(n,i) \qquad FP(n,n)$$

$$FP(d,i) \qquad FP(d,n)$$

Let us now examine the complexity-theoretic aspects of the family of fixpoint logics. We already know that $FP(d, i)$ captures P and $FP(d, n)$ captures PSPACE. One would expect nondeterministic and alternating iteration to capture nondeterministic and alternating computation. This is indeed the case.

Theorem 2.2:

1. *$FP(n, i)$ captures NP.*
2. *$FP(a, i)$, and $FP(n, n)$ capture $PSPACE$.*
3. *$FP(a, n)$ captures $EXPTIME$.*

Theorem 2.2 says that the connection between fixpoint logics and the classes P and PSPACE discovered in [Imm86, Var82] is just one half of the picture; the complete picture is a broad connection between fixpoint logics and complexity classes between P and EXPTIME. Theorem 2.2 is strong enough to yield a separation between $FP(d, i)$ and $FP(a, n)$; this follows from the fact that P is strictly contained in EXPTIME.

The above broad connection between fixpoint logic and computational complexity still suffers from the order mismatch. Thus, we cannot translate containments between complexity classes into containments between fixpoint logics. For example, from the fact that $FP(d, n)$, $FP(n, n)$, and $FP(n, i)$ all capture PSPACE we can only infer that they have the same expressive power over *ordered* structures. This does not tell us anything about the relationship between these logics in general. To get around this difficulty, we develop a theory of *relational* complexity, which mediates beween standard complexity and fixpoint logics.

3 Relational Machines

3.1 The Model

A relational machine is a Turing machine augmented with a *relational store*. The relational store consists of a set of relations of certain arities. Some of these relations are designated

as input relations and some are designated as output relations. The *type* of a relational machine is the sequence of arities of its relations. The *arity* of the machine is the maximal arity of the relations in its store. The tape of the machine is a work tape and is initially empty. In addition to changing its internal state, moving the head on the tape and writing on the tape, the machine can check whether a relation in the store is empty or not and apply relational algebraic operations (which includes Boolean operations as well as the projection operation, see [Ull89]) to the relations in the store. For example, the machine can have instructions such as: "If the machine is in state s_3, the head is reading the symbol 1, and relation R_1 is empty, then change the state to s_4, replace the symbol 1 by 0, move the head to the right, and replace R_2 by $R_2 \cap R_3$".

Relational machines model the embedding of a relational database query language [Cod72] in a general purpose programming language. They are equivalent to, but simpler than, the "loose generic machines" of [AV91].[5] Unlike Turing machines, which get their input spread on the input tape, relational machines get their input in a more natural way. For example, if the input is a graph, then the input to the machine is the set of nodes and the set of edges. (The type of the machine is $\langle 1, 2 \rangle$ and the arity of the machine is 2.) Thus, the order issue does not come up between relational machines and fixpoint logics. We will come back later to the connection between relational machines and fixpoint logics.

Relational machines give rise to three types of computational devices. First, we can think of a relational machine as an acceptor of a *relational language*, i.e., a set of structures. In that case we assume that there is a single 0-ary output relation; the machine accepts if the output relation consists of the 0-ary tuple and rejects if the output relation is empty. We can also think of a relational machine as a *relational function*, computing an output structure for a given input structure. Finally, we can think of relational machine as a *mixed function*, i.e., a function from structures to strings, where the output is written on the machine's tape.

3.2 Discerning Power and Relational Complexity

Relational machines are less "discerning" then Turing machines. We say that a relational machine M cannot *discern* between k-tuples $\mathbf{u}$ and $\mathbf{v}$ over an input $\mathbf{D}$, if for each k-ary relation R in its store and throughout the computation of M over $\mathbf{D}$, we have that $\mathbf{u}$ is in R precisely when $\mathbf{v}$ is in R. For example, if there is an automorphism on $\mathbf{D}$ that maps $\mathbf{u}$ to $\mathbf{v}$, then $\mathbf{u}$ and $\mathbf{v}$ cannot be discerned by M over $\mathbf{D}$. The discerning power of relational machines is best understood by viewing them as an effective fragment of the logic $L^{\omega}_{\infty\omega}$. This is an infinitary logic with a finite number of variables studied recently in [KV90a, KV90b]. The connection between relational machines and $L^{\omega}_{\infty\omega}$ will be developed in [AVV92]; it suffices here to say that it yields a characterization of the discerning power of relational machines in terms of certain infinite 2-player pebble games.

The *k-pebble game* between the Spoiler and the Duplicator on the l-tuples ($l \leq k$) $\mathbf{u}$ and $\mathbf{v}$ of a structure $\mathbf{D}$ has the following rules. Each of the players has k pebbles, say $p_1, \ldots, p_k$ and $q_1, \ldots, q_k$, respectively. The game starts with the Spoiler choosing one of the tuples $\mathbf{u}$ or $\mathbf{v}$, placing $p_1, \ldots, p_l$ on the elements of the chosen tuple (i.e., on $u_1, \ldots, u_l$ or $v_1, \ldots, v_l$), and placing the pebbles $p_{l+1}, \ldots, p_k$ on some elements of $\mathbf{D}$. The Duplicator responds by placing the pebbles $q_1, \ldots, q_l$ on the elements of the other tuple, and placing the pebbles $q_{l+1}, \ldots, q_k$ on some elements of $\mathbf{D}$. In each following round of the game, the Spoiler moves

[5] Unlike relational machines, loose generic machines can apply general 1st-order transformations to the relational store. Relational machines are closely related to the 2nd-order pointer machines of [Lei89a]. Pointer machine manipulate their store by means of tagging and untagging operations, while relational machines manipulate their store by means of relational operations. The definition of relational machines is motivated by our search for an intermediate ground between fixpoint logics and complexity classes.

some pebble p_i (or q_i) to another element of $\mathbf{D}$, and the Duplicator responds by moving the corresponding pebble q_i (or p_i).

Let a_i (resp., b_i), $1 \leq i \leq k$, be the elements of $\mathbf{D}$ under the pebbles p_i (resp., q_i), $1 \leq i \leq k$, at the end of a round of the game. If the mapping h with $h(a_i) = b_i,\ 1 \leq i \leq k$, is not an isomorphism between the substructures of $\mathbf{D}$ with universes $\{a_1, \ldots, a_k\}$ and $\{b_1, \ldots, b_k\}$ respectively, then the Spoiler wins the game. The Duplicator wins the game if he can continue playing "forever", i.e. if the Spoiler can never win the game.

If the Duplicator wins the k-pebble game on the tuple $\mathbf{u}$ and $\mathbf{v}$ of a structure $\mathbf{D}$, then we say that $\mathbf{u}$ and $\mathbf{v}$ are k-equivalent, denoted $\mathbf{u} \equiv_k \mathbf{v}$, over $\mathbf{D}$. The relation $\equiv_k$ characterizes the discerning power of k-ary relational machines.

Proposition 3.1: *The tuples* $\mathbf{u}$ *and* $\mathbf{v}$ *are* k*-equivalent over a structure* $\mathbf{D}$ *if and only if no* k*-ary relational machine* M *can discern between* $\mathbf{u}$ *and* $\mathbf{v}$ *over* $\mathbf{D}$*.*

Since k-ary relational machines cannot discern among k-equivalent tuples, they cannot measure the size of their input. Thus, it seems natural to measure their input size with respect to $\equiv_k$. Let the *k-size* of a structure $\mathbf{D}$, denoted $size_k(D)$, be the number of $\equiv_k$-classes of k-tuples over $\mathbf{D}$. In the spirit of [AV91], we propose to measure the time or space complexity of k-ary relational machines as a function of the k-size of their input. This measure, however, can reasonably serve as a basis for measuring complexity only if it can be calculated by relational machines. The techniques of [AV91] can now be used to show that relational machines *can* measure k-size.

Proposition 3.2: *For each* $k > 0$ *and input type* σ*, there is a relational machine that outputs on its tape, for an input structure* $\mathbf{D}$ *of type* σ*, a string of length* $size_k(\mathbf{D})$ *in time polynomial in* $size_k(\mathbf{D})$

From now on we measure the complexity of k-ary relational machines in terms of the k-size of their input. We can now define relational complexity classes in terms of deterministic (resp., nondeterministic, alternating) relational time or relational space, e.g., $DTIME_r(f(n))$, $NSPACE_r(f(n))$, and so on. Thus, we can define the relational complexity class P_r (relational polynomial time), $\mathrm{NPSPACE}_r$, etc. Proposition 3.2 guarantees the effective enumerability of these classes. In contrast, it is not clear how to get effective enumerability if we define relational complexity in terms of the actual size of the input. Note that known relationships between deterministic, nondeterministic, and alternating complexity classes, such as PSPACE=NPSPACE=APTIME, need not hold for relational complexity classes, since these relationships are typically the result of simulations that use order.

To simplify references to complexity classes, we use the notation *Class*(*Resource*, *Control*, *Bound*) and $Class_r$(*Resource*, *Control*, *Bound*), where *Resource* can be time or space, *Control* can be deterministic, nondeterministic, or alternating, and *Bound* is the bounding function or family of functions. Thus, $Class(time, nondeterministic, poly)$ is NP, and $Class_r(space, deterministic, poly)$ is PSPACE_r. We will always assume that our bounds are at least linear. Typically, *Bound* will be a *polynomially closed* set of functions, i.e., a set of functions that contains the linear functions and contains $p(f(n))$ whenever it contains $f(n)$, for all polynomials $p(x)$.

4 Relational Complexity and Fixpoint Logics

4.1 Relational Complexity and Standard Complexity

What is the relationship between relational and standard complexity? It is easy to see that the k-size of a structure $\mathbf{D}$ is always bounded above by a polynomial ($O(n^k)$) in the size of

$\mathbf{D}$, for all $k > 0$. Furthermore, relational machines are strictly weaker than Turing machines because they cannot check that the cardinality of their input is even. The latter follows from the fact that the logic $L^{\omega}_{\infty\omega}$ has a 0-1 law, so it cannot express the evenness property [KV90a]. We thus obtain the following:

Proposition 4.1: *Let Φ be a polynomially closed set of functions. Then*

$$Class_r(resource, control, \Phi) \subset Class(resource, control, \Phi),$$

for any resource and control.

While relational machines are in some sense weaker than standard machines, the weakness is only due to the lack of order. This weakness disappears in the presence of order. Let w be string, say over the alphabet $\{0, 1\}$, of length n. We can encode w by a structure $rel(w)$ consisting of a total order on the elements $\{0, \ldots, n-1\}$ and a unary predicate on these elements. Note than the k-size of $rel(w)$ is bounded by n^k. For a language L, let $rel(L) = \{rel(w) \mid w \in L\}$. Since $rel(L)$ is ordered, a relational machine can simulate a machine for L.

Proposition 4.2: *Let L be a language in $Class(resource, control, bound)$, for some resource, control, and bound. Then $rel(L)$ is in $Class_r(resource, control, bound)$.*

Combining Propositions 4.1 and 4.2, we get:

Corollary 4.3: *Let Φ_1 and Φ_2 be polynomially closed sets of functions and let $resource_1$, $resource_2$, $control_1$, $control_2$ be resources and controls, respectively. Then we have that*

$$Class(resource_1, control_1, \Phi_1) \subseteq Class(resource_2, control_2, \Phi_2)$$

if

$$Class_r(resource_1, control_1, \Phi_1) \subseteq Class_r(resource_2, control_2, \Phi_2).$$

It follows from Corollary 4.3 that separation results among standard complexity classes translate into separation results among relational complexity classes. For example, it follows that P_r is strictly contained in $\mathrm{EXPTIME}_r$.

To further understand the relationship between standard and relational complexity, we introduce the notion of *reduction* from a relational language to a standard language. We say that a relational language L of type σ is *relationally reducible in polynomial time* to a standard language L' if there exists a deterministic polynomial-time relational machine M acting as a mixed function from σ-structures to strings, such that for each structure $\mathbf{D}$ of type σ we have that $\mathbf{D} \in L$ if and only if $M(\mathbf{D}) \in L'$.

One of the major technical results of this paper is that every relational language can be reduced to a standard language.

Theorem 4.4: *Let Φ be a polynomially closed set of functions and let L be a relational language in $Class_r(resource, control, \Phi)$ for resource and control. Then L is relationally reducible in polynomial time to a language L' in $Class(resource, control, \Phi)$. Furthermore, the reduction depends only on the type of L.*

According to the theorem, the reduction depends only on the type of L. We denote the relational machine that does the reduction for structures of type σ by M_σ, and call it the *relational reduction* machine of input type σ.

Combining Propositions 4.1 and 4.2, Corollary 4.3, and Theorem 4.4, we get that the relationships among relational complexity classes are analogous to the relationships among standard complexity classes.

Corollary 4.5: *Let Φ_1 and Φ_2 be polynomially closed sets of functions and let $resource_1$, $resource_2$, $control_1$, $control_2$ be kinds of resources and controls, respectively. Then*

$$Class(resource_1, control_1, \Phi_1) \subseteq Class(resource_2, control_2, \Phi_2)$$

if and only if

$$Class_r(resource_1, control_1, \Phi_1) \subseteq Class_r(resource_2, control_2, \Phi_2).$$

In particular, it follows from Corollary 4.5, that the known relationships between deterministic, nondeterministic, and alternating complexity classes, such as PSPACE=NPSPACE=APTIME, *do* hold for relational complexity classes, i.e. PSPACE_r=NPSPACE_r=APTIME_r. Also, the open questions about standard complexity classes translate to questions about relational complexity classes, e.g., P=NP if and only if P_r=NP_r.

Further insight into the relationship between standard and relational complexity can be obtained from a closer look at the relationship between size and k-size. Since k-size can be much smaller than size, relational complexity can be much higher than standard complexity. As a result, relational complexity is in some sense orthogonal to standard complexity.

To understand this better, let us re-examine the translation between strings and structures. On one hand, the mapping rel maps strings to structures. It is easy to see that if w is a string, then the k-size of $rel(w)$ is polynomially related to the size of w. On the other hand, according to Theorem 4.4, we have mappings from structures to strings. Let M be the machine mapping structures of the type of $rel(w)$ to strings. It can be checked that for a given string w, the size of $M(rel(w))$ is polynomially related to the size of w. It turns out, however, that using a result of Lindel about the k-size of trees [Lin91], we can encode strings by structures in way that blows up the size without blowing up the k-size.

Proposition 4.6: *For any function $f : N \to N$ there is a mapping rel_f from strings to structures such that:*

1. *if w is a string of length n, then $rel_i(w)$ is of size $f(n)^n$, and*
2. *the k-size of $rel(w)$ is polynomial in n.*

For example, by taking $f(n) = 2$, we can blow the k-size exponentially

Proposition 4.6 implies that we can find languages of different relational complexity but the same standard complexity. In particular, since $P_r \cap P = P_r$, we get the following consequences:[6]

Corollary 4.7:

1. $\text{P}_r = \text{NP}_r \cap \text{P}$ *iff P=NP.*
2. $\text{P}_r = \text{PSPACE}_r \cap \text{P}$ *iff P=PSPACE.*
3. $\text{P}_r \subset \text{EXPTIME}_r \cap \text{P}$.

Corolloary 4.7 tells us that questions about separation among standard complexity classes can be expressed as questions about relational complexity classes of the *same* standard complexity. The first clause says that P and NP separate if and only iff P_r and NP_r separate *inside* P. The second clause in the corollary says that EXPTIME_r is stronger than P_r not only because EXPTIME_r contains problems that cannot be solved in polynomial time, but also because it contains problems that cannot be solved in relational polynomial time even though they can be solved in standard polynomial time. See [DLW91] for related results.

[6] The equivalence of $\text{P}_r = \text{PSPACE}_r \cap \text{P}$ and P=PSPACE. was shown in [AV91].

4.2 Relational Machines and Fixpoint Logics

Fixpoint logics involve iterations of 1st-order formulas. Since relational algebra has the expressive power of 1st-order logic, it is clear that relational machines with the appropriate control can simulate all fixpoint logics. For example, $FP(d,i) \subseteq \mathrm{P}_r$ and $FP(a,n) \subseteq \mathrm{APSPACE}_r$.

To find the precise relationship between fixpoint logics and relational machines consider again Theorem 4.4. According to that theorem, every relational language of a certain relational complexity can be reduced in relational polynomial time to a standard language of the analogous standard complexity. For example, a relational language in NP_r can be reduced in relational polynomial time to a language in NP. According to Theorem 2.2, the fixpoint logic $FP(n,i)$ captures NP. Thus, the only gap now between NP_r and $FP(n,i)$ is the relational polynomial time reduction. This gap is now bridged by the following theorem, which uses the normal-form theorem of [AV91].

Theorem 4.8: *Let M_σ be the relational reduction machine of input type σ. There is a fixpoint formula φ_σ in $FP(d,i)$ such that $\varphi_\sigma(\mathbf{D}) = rel(M_\sigma(\mathbf{D}))$.*

Theorem 4.8 supplies the missing link between relational machines and fixpoint logics. To simulate a relational machine by a fixpoint logic, one first uses φ_σ to obtain $rel(M_\sigma(\mathbf{D}))$, and then uses the logic to simulate a standard machine.

Combining this with Theorem 4.4 we get:[7]

Theorem 4.9:

1. $FP(d,i) = \mathrm{P}_r$
2. $FP(n,i) = \mathrm{NP}_r$
3. $FP(a,i) = FP(d,n) = FP(n,n) = \mathrm{PSPACE}_r$
4. $FP(a,n) = \mathrm{EXPTIME}_r$

Theorem 4.9 should be contrasted with Theorem 2.2. While Theorem 2.2 talks about fixpoint logic *capturing* complexity classes, Theorem 4.9 provides a precise characterization of the expressive power of fixpoint logics in terms of relational complexity classes.

An immediate consequence of Theorem 4.9 is that $FP(a,i)$, $FP(d,n)$ $FP(n,n)$ all have the same expressive power. It follows that noninflationary fixpoint logic can express the problems of nonuniversality for finite automata and truth of quantified Boolean formula, albeit in a very non-straightforward fashion. Recall that Theorem 2.2 yielded the separation of $FP(d,i)$ and $FP(a,n)$. Both of these results demonstrate for the first time how complexity theory can yield *unconditional* results about expressive power of logics.[8]

We introduced relational complexity in order to mediate between fixpoint logic and standard complexity. Corollary 4.5 relates standard complexity to relational complexity, while Theorem 4.9 relates relational complexity to fixpoint logic. Together, they bridge the gap between standard complexity and fixpoint logic – the open questions about the complexity classes between P and EXPTIME can be translated to questions about fixpoint logic.[9]

Corollary 4.10:

[7] The equalities $FP(d,i) = \mathrm{P}_r$ and $FP(d,n) = \mathrm{PSPACE}_r$ were shown in [AV91].

[8] An example where complexity theory yields *conditional* results about expressive power of logics appeared in [Fag74]. Fagin showed that existential and universal second-order logics coincide iff Hamiltonicity is expressible in universal second-order logic.

[9] The equivalence of P=PSPACE and $FP(d,i) = FP(d,n)$ was shown in [AV91].

1. *P=NP iff* $FP(d,i) = FP(n,i)$
2. *P=PSPACE iff* $FP(d,i) = FP(d,n)$
3. *NP=PSPACE iff* $FP(n,i) = FP(n,n)$
4. *PSPACE = EXPTIME iff* $FP(a,i) = FP(a,n)$.

Thus, by the last three equivalences, some of the most tantalizing questions in complexity theory boil down to one fundamental issue: the relative power of inflationary vs. noninflationary 1st-order operators.

5 Concluding remarks

We established a general connection between fixpoint logic and complexity classes from P to EXPTIME. On one side, we have fixpoint logic, parameterized by the choices of 1st-order operators and iteration constructs. On the other side, we have the complexity classes between P and EXPTIME. While this connection is intimate, it is hampered by the order mismatch. Our parameterized fixpoint logics do capture all the complexity classes between P and EXPTIME, but equality is achieved only over ordered structures

To bridge this mismatch, we developed a theory of relational complexity, which bridges the gap between standard complexity and fixpoint logic. On one hand, we show that questions about containments among standard complexity classes can be translated to questions about containments among relational complexity classes. On the other hand, the expressive power of fixpoint logic can be precisely characterized in terms of relational complexity classes. This tight three-way relationship among fixpoint logics, relational complexity and standard complexity yields in a uniform way logical analogs to all containments among the complexity classes P, NP, PSPACE, and EXPTIME. The logical formulation shows that some of the most tantalizing questions in complexity theory boil down to a single question: the relative power of inflationary vs. noninflationary 1st-order operators.

References

[AU79] A. V. Aho and J. D. Ullman. Universality of data retrieval languages. In *Proc. 6th ACM Symp. on Principles of Programming Languages*, pages 110–117, 1979.

[AV89] S. Abiteboul and V. Vianu. Fixpoint extensions of first-order logic and Datalog-like languages. In *Proc. 4th IEEE Symp. on Logic in Computer Science*, pages 71–79, 1989.

[AV90] S. Abiteboul and V. Vianu. Non-deterministic languages to express deterministic transformations. In *Proc. 9th ACM Symp. on Principles of Database Systems*, pages 218–229, 1990.

[AV91] S. Abiteboul and V. Vianu. Generic computation and its complexity. In *Proc. 23rd ACM Symp. on Theory of Computing*, pages 209–219, 1991.

[AVV92] S. Abiteboul, Moshe Y. Vardi, and V. Vianu. Computing with infinitary logic. Unpublished manuscript, 1992.

[Bus86] S. R. Buss. *Bounded Arithmetics*. Bibliopolis, 1986.

[CH80] A. Chandra and D. Harel. Computable queries for relational databases. *Journal of Computer and System Sciences*, 21:156–178, 1980.

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

[CKS81] A. Chandra, D. Kozen, and L. Stockmeyer. Alternation. *Journal of the ACM*, 28:114–133, 1981.

[Cod72] E. F. Codd. Relational completeness of data base sublanguages. In R. Rustin, editor, *Database Systems*, pages 33–64. Prentice-Hall, 1972.

[Com88] K. J. Compton. An algebra and a logic for NC^1. In *Proc. 3rd IEEE Symp. on Logic in Computer Science*, pages 12–21, 1988.

[DLW91] A. Dawar, S. Lindell, and S. Weinstein. Infinitary logic and inductive definability over finite structures. Research report, Univ. of Pennsylvania, 1991.

[dR84] M. de Rougemont. Uniform definability on finite structures with successor. In *Proc. 16th ACM Symp. on Theory of Computing*, pages 409–417, 1984.

[Fag74] 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.

[Fag75] R. Fagin. Monadic generalized spectra. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:89–96, 1975.

[Fed81] T. Feder. *Stable networks and product graphs.* PhD thesis, Stanford University, 1981.

[Fri71] H. Friedman. Axiomatic recursive function theory. In R. O. Gandy and C. E. M. Yates, editors, *Logic Colloquium '69*, pages 113–137. North Holland, 1971.

[GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability - A Guide to the Theory of NP-Completeness.* W. H. Freeman and Co., 1979.

[Goe89] A. Goerdt. Characterizing complexity classesby higher-type primitive-recursive definitions. In *Proc. 4th IEEE Symp. on Logic in Computer Science*, pages 364–374, 1989.

[Gra84] E. Grandjean. The spectra of first-order sentences and computational complexity. *SIAM Journal on Computing*, 13:356–373, 1984.

[Gra85] E. Grandjean. Universal quantifiers and time complexity of random access machines. *Mathematical System Theory*, 13:171–187, 1985.

[GS86] Y. Gurevich and S. Shelah. Fixed-point extensions of first-order logic. *Annals of Pure and Applied Logic*, 32:265–280, 1986.

[GSS90] J-Y. Girard, A. Scedrov, and P. SCott. Bounded linear logic: a modular approach to polynomial time computability. In S. R. Buss and P. Scott, editors, *Feasible Mathematics*, pages 195–207. Birkhauser, 1990.

[Gur83] Y. Gurevich. Algebras of feasible functions. In *Proc. 24th IEEE Symp. on Foundations of Computer Science*, pages 210–214, 1983.

[Gur84] Y. Gurevich. Toward logic tailored for computational complexity. In M. M. Richter et al., editor, *Computation and Proof Theory, Lecture Notes in Mathematics 1104*, pages 175–216. Springer-Verlag, 1984.

[Gur88] 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.

[HP84] D. Harel and D. Peleg. Static logics, dynamic logics, and complexity classes. *Information and Control*, pages 86–102, 1984.

[IL90] N. Immerman and E. S. Lander. Describing graphs: a first-order approach to graph canonization. In A. Selman, editor, *Complexity Theory Retrospective*, pages 59–81. Springer-Verlag, 1990.

[Imm86] N. Immerman. Relational queries computable in polynomial time. *Information and Control*, 68:86–104, 1986.

[Imm87a] N. Immerman. Expressibility as a complexity measure: results and directions. In *Second Structure in Complexity Conference*, pages 194–202, 1987.

[Imm87b] N. Immerman. Languages that capture complexity classes. *SIAM Journal of Computing*, 16:760–778, 1987.

[Imm89] N. Immerman. Descriptive and computational complexity. In J. Hartmanis, editor, *Computational Complexity Theory, Proc. Symp. Applied Math., Vol. 38*, pages 75–91. American Mathematical Society, 1989.

[JS74] N. G. Jones and A.L. Selman. Turing machines and the spectra of first-order formulas. *Journal of Symbolic Logic*, 39:139–150, 1974.

[KV90a] Ph. G. Kolaitis and M. Y. Vardi. 0-1 laws for infinitary logics. In *Proc. 5th IEEE Symp. on Logic in Computer Science*, pages 156–167, 1990.

[KV90b] Ph. G. Kolaitis and M. Y. Vardi. On the expressive power of Datalog: tools and a case study. In *Proc. 9th ACM Symp. on Principles of Database Systems*, pages 61–71, 1990. Full version appeared in IBM Research Report RJ8010, March 1991.

[Lei89a] D. Leivant. Descriptive characterization of computational complexity. *Journal of Computer and System Sciences*, 39:51–83, 1989.

[Lei89b] D. Leivant. Monotonic use of space and computational complexity over abstract structures. Unpublished manuscript, 1989.

[Lei90] D. Leivant. Inductive definitions over finite structures. *Information and Computation*, 89:95–108, 1990.

[Lei91] D. Leivant. A foundational delineation of computational feasibility. In *Proc. 6th IEEE Symp. on Logic in Computer Science*, pages 2–11, 1991.

[Lin91] S. Lindell. An analysis of fixed point queries on binary trees. *Theoretical Computer Science*, 85:75–95, 1991.

[Liv82] A. B. Livchak. The relational model for systems of automatic testing. *Automatic Documentation and Mathematical Linguistics*, 4:17–19, 1982.

[Liv83] A. B. Livchak. The relational model for process control. *Automatic Documentation and Mathematical Linguistics*, 4:27–29, 1983.

[Lyn82] J. F. Lynch. Complexity classes and theories of finite models. *Mathematical System Theory*, 15:127–144, 1982.

[Mos74] Y. N. Moschovakis. *Elementary Induction on Abstract Structures*. North Holland, 1974.

[Sav80] W. J. Savitch. Relational between nondeterministic and deterministic tape complexity. *Journal of Computer and System Sciences*, 4:177–192, 1980.

[Saz80a] V. Yu. Sazonov. A logical approach to the problem of P=NP. In *Proc. 9th Conf. on Math. Found. of Computer Science*, pages 561–575. Springer-Verlag, Lecture Notes in Computer Science 88, 1980.

[Saz80b] V. Yu. Sazonov. Polynomial computability and recursivity in finite domains. *Elektronische Informationverarbeitung und Kybernetik*, 16:319–323, 1980.

[Sto77] L. J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3:1–22, 1977.

[TU88] J. Tiuryn and P. Urzyczyn. Some relationships between logic of programs and complexity theory. *Theoretical Computer Science*, 60:83–108, 1988.

[Ull89] J. D. Ullman. *Database and Knowledge-Base Systems, Volumes I and II.* Computer Science Press, 1989.

[Var82] M. Y. Vardi. The complexity of relational query languages. In *Proc. 14th ACM Symp. on Theory of Computing*, pages 137–146, 1982.