Logic in the Finite

1 Validity in the Finite

Is it simpler to reason about finite structures or about arbitrary structures? Some of the major results of logic in the twentieth century provide a clear and surprising answer to one precise version of this question. Suppose first that we restrict our reasonings to propositions which are expressible in first order logic. We may then understand the question as asking for a comparison between the complexity of

1. determining whether a first order sentence is valid, that is, true under every interpretation whatsoever, and

2. determining whether a first order sentence is valid in the finite, that is, true under every interpretation with a finite universe of discourse.

This question can be formulated more concisely and concretely in terms of Val, the set of valid sentences of L, the first order language with identity and a single binary relation symbol E, and Fval, the set of sentences of L which are valid in the finite, namely: is the decision problem for Fval simpler than the decision problem for Val?

Let's begin by analyzing the complexity of the decision problem for Fval. It is easy to see that we can make an effective list $A_1, A_2, \ldots$ of finite structures for L which contains every such structure up to isomorphism. We may now subject a sentence $\varphi \in L$ to the following effective procedure: successively test whether A_1 satisfies φ, A_2 satisfies $\varphi, \ldots$; at the first stage where the outcome is negative, halt the procedure and return the answer "no." Clearly, this procedure yields the correct answer to the query "is φ valid in the finite," if the answer is negative, and yields no answer, otherwise. That is, the complement of Fval is recursively enumerable, or in other words, Fval is co-r.e.

If we attempt such a direct approach to analyzing the complexity of Val, we are stymied at the outset. There is no possibility of effectively generating a list of all structures up to isomorphism, since there are structures of every infinite cardinality; moreover, there is, in general, no effective way to test whether a given infinite structure A satisfies a sentence $\varphi \in L$. Reflection on the apparent complexity of the notion of validity provides the proper context in which to appreciate the extraordinary depth of Gödel's Completeness Theorem for first order logic: there is a sound and complete effective proof procedure for first order validity. In other words, Val is recursively enumerable – in order to discover that a first order sentence is valid, if it is, we need

only look through an effectively generated list of finite objects and check that one is its proof.

So far so good: Val is r.e.; Fval is co-r.e. To complete the picture we need to invoke two more fundamental results of twentieth century logic. Church's Theorem tells us that Val is undecidable, from which it follows that Val is not co-r.e. On the other hand, Trakhtenbrot's Theorem (see Trakhtenbrot (1950)) tells us that Fval is undecidable, from which it follows that Fval is not r.e., that is, there is no sound and complete proof procedure for the first order sentences which are valid in the finite. This suggests one answer to the question with which we began: reasoning about finite structures is no simpler than reasoning about arbitrary structures – there is an effective proof procedure for validity, but no effective proof procedure for validity in the finite. Indeed, there is a good sense in which we can say that the complexity of the decision problems for Val and Fval are identical, namely, Val and Fval are Turing reducible to one another. That is, there is a Turing machine which will decide membership in Val given an oracle for Fval and there is a Turing machine which will decide membership in Fval given an oracle for Val. Remarkably, Val and Fval turn out to have effectively the same information content.

2 Model Theory in the Finite?

The last section suggests that, in a sense, there can be no proof theory for first order logic in the finite, since there can be no effective proof procedure for validity in the finite. How about model theory? At the outset, there are disappointments. One of the central results in the model theory of first order logic, the Compactness Theorem, does not extend to the finite case. Recall the Compactness Theorem: if every finite subset of a set of first order sentences Γ is satisfiable, then Γ itself is satisfiable. Call a set of sentences Γ *satisfiable in the finite*, if and only if, there is a finite structure A which satisfies every sentence in Γ. It is easy to construct a set of first order sentences Γ such that every finite subset of Γ is satisfiable in the finite, whereas Γ itself is not satisfiable in the finite. For example, let $\Gamma = \{\lambda_n \mid n > 0\}$, where λ_n is a first order sentence in the pure language of identity which is true in a structure A, if and only if, the size of A is at least n. Virtually all the finite analogs of well-known consequences of the Compactness Theorem fail as well, for example, the Beth Definability Theorem, the Craig Interpolation Theorem, most all "preservation theorems," *etc.* (See Gurevich (1984) for a compendium of such results; a notable exception is van Benthem's preservation theorem for the modal fragment of first order logic, see Rosen (1997).)

Further contrasts between the finite model theory of first order logic and classical model theory abound. A central phenomenon of first order model theory is that no infinite structure can be characterized up to isomorphism by a set of first order sentences. Recall that structures A and B are elementarily equivalent, if and only if, they satisfy the same first order sentences. It is a corollary of the Compactness Theorem that for every infinite structure A, there is a structure B (indeed, a proper class of pairwise nonisomorphic structures B) such that A is elementarily equivalent to B, but A is not isomorphic to B. In contrast, it is easy to show that for all structures A and B, if A is finite and B is elementarily equivalent to A, then B is isomorphic to A. Indeed, for every finite structure A whose signature is finite, there is a single first order sentence φ such that for every structure B, B satisfies φ, if and only if, B is isomorphic to A.

3 Definability and Complexity

In light of all these contrasts, one might legitimately wonder what finite model theory could be. The following sections attempt to answer this question by giving a feeling for some of the techniques, results, and open problems of the subject. For the most part, we will pursue questions in definability theory, that is, we will inquire into the expressive power of various logical languages in the context of finite structures. We will see that this study has close connections with the theory of computational complexity.

We collect together here some notions and notations that will ease our progress. A structure A, for us, consists of a universe of discourse $|A|$ and interpretations for a finite set of relation symbols and constant symbols; this set of symbols is called the signature of A. Whenever we mention two structures in the same breath, they are of the same signature; whenever we speak of a collection of structures, they are of the same signature. Let $\mathcal{K}$ be a class of structures. A collection of structures $\mathcal{Q} \subseteq \mathcal{K}$ is a *query relative to* $\mathcal{K}$, if and only if, $\mathcal{Q}$ is isomorphism closed in $\mathcal{K}$, that is,

$$\forall A, B \in \mathcal{K}((A \in \mathcal{Q} \wedge A \cong B) \rightarrow B \in Q).$$

We will drop the qualification "relative to $\mathcal{K}$" when the background class is clear from the context. Queries are the proper object of study in our

investigation of definability and complexity, since logical languages do not distinguish between isomorphic structures.

We think of a logical language L as consisting of a set of sentences S_L and a satisfaction relation $\models_L$. We will suppress the subscript to $\models$ as it will generally be clear from the context. Given a class of structures $\mathcal{K}$ and a sentence $\varphi \in S_L$, we write $\varphi(\mathcal{K})$ for the *query defined by φ relative to $\mathcal{K}$*, that is,

$$\varphi(\mathcal{K}) = \{A \in \mathcal{K} \mid A \models \varphi\}.$$

We write $L(\mathcal{K})$ for $\{\varphi(\mathcal{K}) \mid \varphi \in S_L\}$, the set of queries which are L-definable relative to $\mathcal{K}$.

In what follows, we will analyze and compare the logical and computational complexity of queries relative to classes of finite structures. It will be convenient to introduce, for each signature σ, a canonical countable set of finite structures $\mathcal{F}_\sigma$ which contains, up to isomorphism, every finite structure of signature σ. We let $\mathcal{F}_\sigma$ be the set of structures of signature σ with universe of discourse $[n](= \{1, \ldots, n\})$ for some $n \geq 1$. Unless otherwise indicated, all collections of finite structures we mention are understood to be subsets of $\mathcal{F}_\sigma$ for some σ. We write $\mathcal{D}$ for $\mathcal{F}_{\{E\}}$ where E is a binary relation symbol; $\mathcal{D}$ is, for us, the class of finite directed graphs. For simplicity and concreteness, our discussion will often focus on queries relative to $\mathcal{D}$.

In the following sections, we will address questions concerning the logical resources that are required to define a given query $\mathcal{Q} \subseteq \mathcal{D}$. For example, we will consider whether $\mathcal{Q}$ is definable in second order, but not in first order, logic; or whether $\mathcal{Q}$ is definable by an existential second order sentence, but not by the negation of such a sentence, *etc.* We can think of this study as yielding information about the complexity of $\mathcal{Q}$ – for example, if $\mathcal{Q}$ is not first order definable, while $\mathcal{Q}'$ is, we might want to say that the definitional, or descriptive, complexity of $\mathcal{Q}'$ is no greater than that of $\mathcal{Q}$. In this way, we can think of the classes of queries $L(\mathcal{D})$, for various languages L, as descriptive complexity classes, in analogy with the resource complexity classes studied in the theory of computation (see Papadimitriou (1994) for background on the theory of computational complexity). Let us pursue this analogy.

Consider a query $\mathcal{Q} \subseteq \mathcal{D}$. We have been thinking of $\mathcal{Q}$ under the guise of definability. We can, on the other hand, think of $\mathcal{Q}$ as a decision problem: given an $A \in \mathcal{D}$ answer the question whether or not A is a member of $\mathcal{Q}$. Rather than asking what logical resources are required to specify $\mathcal{Q}$, we can ask instead, what computational resources are required to decide membership in $\mathcal{Q}$. To make this precise, we can easily encode each $A \in \mathcal{D}$ as a bit string, thereby making it a suitable input to a Turing machine. If A is of size n, the adjacency matrix of A is the $n \times n$ matrix whose i, j-entry is

a 1, if $\langle i, j \rangle \in E^A$, and is a 0, otherwise. We encode A as the bit string $c(A)$ which consists of the concatenation of the rows of the adjacency matrix of A, and for $\mathcal{Q} \subseteq \mathcal{D}$, we let $c(\mathcal{Q}) = \{c(A) \mid A \in \mathcal{Q}\}$. If Y is a resource complexity class, then we write $\mathsf{Y}(\mathcal{D})$ for the collection of queries $\mathcal{Q} \subseteq \mathcal{D}$ such that $c(\mathcal{Q}) \in \mathsf{Y}$. (In a similar fashion, we may define $\mathsf{Y}(\mathcal{F}_\sigma)$ for any signature σ.) We are now in a position to make direct comparisons between resource and descriptive complexity classes. In the following sections, we will see that many important resource complexity classes, for example, P and NP, have natural logical characterizations relative to various sets of finite structures.

4 First Order Definability

One of the main tools for establishing limits on the expressive power of first order logic over arbitrary structures is the Compactness Theorem. As noted in Section 2, we are deprived of the use of this tool in the context of finite structures, so we will need to rely on other techniques. We begin with an exemplary application of the Compactness Theorem, so we can appreciate what we are missing; the example will reappear throughout the following sections.

Let $\mathcal{D}^*$ be the collection of arbitrary structures A of signature $\{E\}$; each $A \in \mathcal{D}^*$ is a, perhaps infinite, directed graph. We call such a graph A *simple*, if and only if, E^A is irreflexive and symmetric, and we let $\mathcal{G}^*$ be the collection of arbitrary simple graphs. A simple graph may be visualized as a loop-free, undirected graph. Note that $\mathcal{G}^*$ is first order definable relative to $\mathcal{D}^*$. Now let $\mathcal{D}^*_{st}$ (resp., $\mathcal{G}^*_{st}$) be the collection of expansions of structures in $\mathcal{D}^*$ (resp., $\mathcal{G}^*$) to the signature with two additional constant symbols s and t – this is the collection of directed (resp., simple) source–target graphs. A graph $A \in \mathcal{D}^*_{st}$ is *reachable*, if and only if, there is a path from s^A to t^A in A, that is, a sequence $a_1, \ldots, a_n$ of nodes of A such that $a_1 = s^A, a_n = t^A$, and for every $1 \leq i < n, \langle a_i, a_{i+1} \rangle \in E^A$. Let $\mathcal{S}^*$ be the collection of $A \in \mathcal{G}^*_{st}$ such that A is reachable. Is $\mathcal{S}^*$ first order definable relative to $\mathcal{G}^*_{st}$? An

application of the Compactness Theorem provides a negative answer. For suppose that there is a first order sentence φ with $\varphi(\mathcal{G}^*_{st}) = \mathcal{S}^*$. Let Γ be the set consisting of the following sentences:

$$\begin{array}{ll} \psi_0 & \neg s = t \\ \psi_1 & \neg Est \\ \psi_2 & \neg\exists x(Esx \wedge Ext) \\ \vdots & \vdots \end{array}$$

Notice that a graph A satisfies the conjunction $\psi_0 \wedge \ldots \wedge \psi_n$, if and only if, there is no path in A of length $\leq n$ from s^A to t^A. Therefore, the simple chain of length $n+1$ with end nodes labelled s and t satisfies $\psi_0 \wedge \ldots \wedge \psi_n$, from which it follows that every finite subset of $\Gamma \cup \{\varphi\}$ is satisfiable. Therefore, by the Compactness Theorem, $\Gamma \cup \{\varphi\}$ is satisfiable. On the other hand, it is clear that if a graph A satisfies Γ, then A is not reachable. But, this contradicts the hypothesis that φ defines $\mathcal{S}^*$.

Now, let $\mathcal{S} \subset \mathcal{S}^*$ be the set of finite reachable simple source-target graphs. The question whether $\mathcal{S}$ is first order definable is no longer immediately accessible to an application of the Compactness Theorem of the sort sketched above. The Compactness Theorem can be pressed into service to answer the question by exploiting "pseudofinite" structures, that is, infinite structures which satisfy every first order sentence which is valid in the finite (see Gaif-

man & Vardi (1985) for details); but, we will follow a different approach, due to Gurevich (1984), which proceeds via Ehrenfeucht games and yields additional information. The approach involves a reduction from a query on linear orders.

Let $\mathcal{L}_{st} \subseteq \mathcal{F}_{\{<,s,t\}}$ be the set of finite linear orders with minimal element s and maximal element t. The conjunction of the following first order conditions defines $\mathcal{L}_{st}$.

$\forall x \neg(x < x)$ (irreflexive)

$\forall x \forall y \forall z((x < y \wedge y < z) \rightarrow x < z)$ (transitive)

$\forall x \forall y(x < y \vee y < x \vee x = y)$ (total)

$\forall x(\neg(x < s) \wedge \neg(t < x))$ (endpoints)

Let $\mathcal{M} \subseteq \mathcal{L}_{st}$ be the set of odd linear orders, that is, linear orders with universe $[2n + 1]$, for some n. Is $\mathcal{M}$ first order definable relative to $\mathcal{L}_{st}$?

Here is one strategy for attempting to show that $\mathcal{M}$ is not first order definable. For each first order sentence φ, show that there are $A, B \in \mathcal{L}_{st}$ such that A and B agree about φ (either they both satisfy φ or they both fail to do so), $A \in \mathcal{M}$, and $B \notin \mathcal{M}$. It is clear that if we succeed in doing this, we have shown that $\mathcal{M}$ is not first order definable. (Indeed, the converse holds as well – the strategy is nothing more than a restatement of what's required.)

What makes the strategy worth pursuing is that there is a powerful, and entertaining, technique, the Ehrenfeucht game, for showing that pairs of structures agree about first order sentences. This technique applies to both finite and infinite structures and, to some extent, fills the void left by the failure of compactness in finite model theory.

The Ehrenfeucht game is played between two players, conventionally called the Spoiler and the Duplicator. The equipment for the game consists of two boards, one representing the graph A and the other representing the graph B, and an unlimited supply of pairs of pebbles $\langle \alpha_1, \beta_1 \rangle, \langle \alpha_2, \beta_2 \rangle, \ldots$. The game is played through a sequence of rounds as follows. At the i-th round of the game, the Spoiler chooses one of the pebbles from the pair $\langle \alpha_i, \beta_i \rangle$ and places it on a node of the corresponding board A or B, the α pebbles are played onto A and the β pebbles onto B. The Duplicator then places the remaining pebble on the other board, completing the round of play. Suppose the game has proceeded through n-rounds of play. Let a_i be the node in A covered by α_i and let b_i be the node in B covered by β_i. Let f be the mapping which sends a_i to b_i for all $1 \leq i \leq n$ and sends s^A to s^B and t^A to t^B. If f is a partial isomorphism from A to B (that is, a one to one, edge preserving map) we say the Duplicator wins the game through n-rounds of play. Thus, the Spoiler's goal is to reveal structural

distinctions between A and B, the Duplicator's goal is to hide them. We say that A is n-similar to B, if and only if, the Duplicator has a strategy to win every play of the Ehrenfeucht game on A and B through n-rounds. We say structures A and B are n-equivalent, if and only if, A and B satisfy exactly the same first order sentences of quantifier rank $\leq n$ (recall that the quantifier rank of a formula is the maximum depth of nesting of quantifiers in the formula). The Ehrenfeucht-Fraïssé Theorem tells us that n-similarity and n-equivalence coincide, that is, for all structures A and B and for every n, A is n-similar to B, if and only if, A is n-equivalent to B (see Ehrenfeucht (1961) and Fraïssé(1954)).

Armed with the Ehrenfeucht-Fraïssé Theorem, we can now implement our strategy for showing that $\mathcal{M}$ is not first order definable. For each n, it suffices to construct a pair of finite linear orders A and B such that $A \in \mathcal{M}$, $B \notin \mathcal{M}$, and A is n-similar to B. We accomplish this by overkill – for each n, if A and B are finite linear orders of length $> 2^n$, then A is n-similar to B. To see this, consider the following strategy for the Duplicator in the n-round game played on two such linear orders. At round m, the Duplicator plays as follows. Suppose, without loss of generality, that the Spoiler has played into A. This play falls into one of m intervals into which A has been divided by the play of pebbles at earlier rounds of the game and it

determines distances d_1 and d_2 between the newly pebbled point and the left and right endpoints of that interval, respectively. The Duplicator plays into the corresponding interval in B so as to achieve the following approximation between these distances and the corresponding distances d_1' and d_2' between the point she pebbles and the endpoints of her interval. Namely, for $i = 1, 2$ if $d_i \leq 2^{(n-m)}$, then $d_i = d_i'$, and if $d_i > 2^{n-m}$, then $d_i' > 2^{n-m}$. The initial condition on the lengths of A and B insures that the Duplicator can maintain these approximations through n-rounds of play. Thus, $\mathcal{M}$ is not first order definable. Indeed, any first order definable collection of finite linear orders is a finite or cofinite subset of $\mathcal{L}_{st}$.

Now, we reduce the problem of defining odd length linear orders ($\mathcal{M}$) to the problem of defining reachability ($\mathcal{S}$). Let $\rho(x, y)$ be a first order formula which is true of a pair of elements of a linear order, if and only if, the second is the successor of the successor of the first, and let $\chi(x, y)$ be the formula $\rho(x, y) \vee \rho(y, x)$. Suppose $A \in \mathcal{L}_{st}$. We may use the formula χ to define a simple source-target graph B from A. We let $|B| = |A|$, $s^B = s^A$, $t^B = t^A$, and $E^B = \{\langle u, v\rangle \mid A \models \chi[u, v]\}$. Now, observe that the graph B thus defined is reachable, if and only if, $A \in \mathcal{M}$. Suppose that there is a first order sentence θ which defines $\mathcal{S}$. Let θ' be the result of replacing each subformula of the form Exy in θ with $\chi(x, y)$. Then, θ' defines $\mathcal{M}$. We

have exhibited a "first order reduction" of $\mathcal{M}$ to $\mathcal{S}$; it follows at once that $\mathcal{S}$ is not first order definable, since $\mathcal{M}$ is not. Such first order reductions are an important descriptive analog of the resource bounded reductions of computational complexity theory.

The foregoing examples show that some simple properties of finite graphs are not first order definable. These examples can be easily multiplied – acyclicity, regularity, 2-colorability, *etc.* all fail to be first order definable. Lest the reader be left with the impression that no interesting classes of finite graphs are first order definable, note that the collection $\mathcal{FR}$ of finite nonempty ranks of the cumulative hierarchy of sets equipped with the membership relation as their edge relation is first order definable (see Dawar, Doets, Lindell, & Weinstein (1998)). In Section 6, we will see that questions concerning the expressive power of first order logic relative to $\mathcal{FR}$ are directly related to open problems in the theory of computational complexity.

5 Second Order Definability

What logical resources are required to define reachability over finite graphs? As we've just seen, first order logic doesn't suffice. There are several routes to the definability of reachability. Let's begin with Frege's (see Frege (1884)). The transitive closure (sometimes called the ancestral) of a binary relation R is the smallest relation (in the sense of inclusion) which is transitive and includes R. For example, the relation "ancestor of" is the transitive closure of the relation "parent of." If R is a binary relation, we write $\mathrm{tc}(R)$ for the transitive closure of R.

Frege observed that the relational operator tc is uniformly definable by a formula $\tau(x, y)$ of second order logic; that is, for every structure $A \in \mathcal{D}^*, \mathrm{tc}(E^A) = \{\langle u, v\rangle \mid A \models \tau[u, v]\}$. The formula $\tau(x, y)$ may be chosen to be:

$$\forall P((\forall z(Exz \to Pz) \wedge \forall v \forall w((Pv \wedge Evw) \to Pw)) \to Py).$$

This formula has a couple of noteworthy features. First, it is a universal second order formula, that is, it is of the form

$$\forall P_1 \ldots \forall P_n \theta$$

with θ first order. Second, it is monadic universal, that is, each of the universal quantifiers binds a monadic second order variable. We call the

fragment of second order logic consisting of all such formulas mon-Π^1_1. Now, let $\mathcal{R}^* \subseteq \mathcal{D}^*_{st}$ be the collection of reachable directed source-target graphs. It is clear that $\tau(s,t)$ defines $\mathcal{R}^*$ relative to $\mathcal{D}^*_{st}$; directed reachability is mon-Π^1_1 definable.

Is $\mathcal{R}^*$ also definable by a monadic existential second order sentence? Since the full existential fragment of second order logic is compact, the argument we gave at the beginning of Section 3 to show that $\mathcal{S}^*$ is not first order definable, also shows that $\mathcal{S}^*$ (and hence $\mathcal{R}^*$ as well) is not definable by an existential second order sentence, monadic or otherwise. In the finite case, the situation is subtler. Paris Kanellakis observed (see Immerman (1999)) that $\mathcal{S}$ is definable by a monadic existential second order sentence $\exists P\theta$, where θ is the conjunction of the following first order conditions.

$$Ps \wedge \exists! x(Px \wedge Esx) \quad (s \text{ has degree 1 in } P)$$

$$Pt \wedge \exists! x(Px \wedge Etx) \quad (t \text{ has degree 1 in } P)$$

$$\forall x((Px \wedge x \neq s \wedge x \neq t) \rightarrow \exists y \exists z(Py \wedge Pz \wedge y \neq z \wedge \forall w(Pw \rightarrow (Exw \leftrightarrow (w = y \vee w = z))))) \quad (\text{all other nodes have degree 2 in } P)$$

If a finite simple graph A satisfies θ with respect to an assignment of a set of nodes X to P, then the nodes in X form a simple chain with end nodes s^A and t^A. (The reader should construct an infinite simple graph which is

not reachable, but satisfies $\exists P\theta$.)

Let $\mathcal{R} \subset \mathcal{R}^*$ be the collection of finite reachable source-target graphs; this class differs from $\mathcal{S}$ in omitting the requirement of simplicity. Ajtai & Fagin (1990) established that $\mathcal{R}$ is not definable by a monadic existential second order sentence. Their argument blends an extension of the Ehrenfeucht game to monadic existential second order logic with probabilistic techniques (see Section 8 for a discussion of such techniques). This result establishes a difference in the descriptive complexity of $\mathcal{S}$ and $\mathcal{R}$, the former is definable in both mon-Π^1_1 and mon-Σ^1_1 (the monadic existential fragment of second order logic), the latter only in mon-Π^1_1. From an intuitive point of view, the problem of determining whether a finite directed graph is reachable is more complex than the same problem restricted to simple graphs. It appears that descriptive complexity provides a more convincing acount of this intuitive distinction than analysis of the computational complexity of these problems has yet been able to offer (see Ajtai & Fagin (1990) for further discussion).

The foregoing considerations leave open the question whether $\mathcal{R}$ is definable by an existential second order sentence not subject to the monadic restriction. Rather than exhibiting such a sentence directly, which is straightforward, we will see that a positive answer to this question is a corollary of a celebrated result of Fagin (1974), namely: for all σ, $\mathsf{NP}(\mathcal{F}_\sigma) = \Sigma^1_1(\mathcal{F}_\sigma)$ (Σ^1_1

is the set of existential second order sentences). Fagin's Theorem has been dubbed the first theorem of descriptive complexity theory. It equates the important computational complextiy class of queries whose decision problems are solvable by nondeterministic Turing machines in polynomial time with the descriptive complexity class of queries which are definable by existential second order sentences. Fagin's Theorem provides a machine independent characterization of NP – in order to verify that a query is in NP, one needn't tinker with machines and time bounds, just produce a Σ_1^1 sentence which defines it. In a sense, Fagin's Theorem shows that existential second order logic is an alternative, what might be called, "higher-level," programming language for specifying exactly the NP queries: the proof of the theorem yields an effective procedure F for "compiling" an arbitrary existential second order sentence φ into a polynomially clocked nondeterministic Turing machine $F(\varphi)$ which accepts the query defined by φ and establishes that every query in NP is accepted by one of the machines $F(\varphi)$. Thus, existential second order logic yields an effective enumeration of the NP queries, with the relation of satisfaction as the enumerating relation.

To return to our story of reachability, $\mathcal{R}$ is in NP – indeed it is in NL, the class of problems solvable by nondeterministic Turing machines using only logarithmic work space, and this class is included in P the class of problems

solvable by deterministic Turing machines in polynomial time. It is generally believed that both the inclusions $\mathsf{NL} \subseteq \mathsf{P}$ and $\mathsf{P} \subseteq \mathsf{NP}$ are strict, but three decades of intense investigation have failed to produce a proof for the strictness of either. Fagin's Theorem opened up the possibility of attacking such outstanding problems in the theory of computational complexity by means of logical techniques. For example, in order to show that $\mathsf{P} \neq \mathsf{NP}$, it would suffice to show that there is a query $\mathcal{Q}$ such that $\mathcal{Q} \notin \Sigma^1_1(\mathcal{D})$ and $\mathcal{Q} \in \Pi^1_1(\mathcal{D})$, for, by Fagin's Theorem, this would establish that NP is not closed under complementation. The results mentioned earlier on the monadic fragments of Π^1_1 and Σ^1_1 are of some interest in this connection. We saw that $\mathcal{R} \in \text{mon}{-}\Pi^1_1(\mathcal{D})$ whereas $\mathcal{R} \notin \text{mon}{-}\Sigma^1_1(\mathcal{D})$. This does not resolve any outstanding problem concerning computational complexity since mon-Σ^1_1 does not correspond to any natural level of computational complexity. On the one hand, as we've just noted, $\mathcal{R}$ is in NL but not in mon-Σ^1_1. On the other hand, mon-Σ^1_1 contains NP-complete problems, that is, problems which are of maximal complexity among problems in NP with respect to polynomial time reduction. For example, the NP-complete query graph 3-colorability is easily seen to be in mon-Σ^1_1. Thus, though the result of Ajtai & Fagin (1990) does not lead to a separation of computational complexity classes, it does indicate how logic can contribute to a richer understanding

of complexity by focussing attention on complexity classes which are orthogonal to the standard computational complexity measures, yet natural from a descriptive point of view.

6 Inductive Definability

In this section, we will pursue a more constructive approach to the definability of the set of reachable graphs. We will see that there are interesting connections between constructivity and complexity in this context.

One of the outstanding open problems of descriptive complexity theory concerns the existence of logics which characterize computational complexity classes below NP. An important result, due independently to Immerman (1986) and Vardi (1982), is that P is characterized by FO+LFP relative to ordered finite structures. FO+LFP is the extension of first order logic by a least fixed point operator for defining relations by induction. Least fixed point operators have played a major role in studies of definability on fixed infinite structures (see Moshovakis (1974)). Let $\varphi(R, x_1, \ldots, x_k)$ be a first-order formula with a distinguished k-ary relation symbol R. On a structure, A, we can use φ to define the relational operator, $\Phi_A(X) = \{\langle a_1, \ldots a_k\rangle | A \models \varphi[X, a_1, \ldots, a_k]\}$ (here, X is a k-ary relation on A and the notation stands for the assignment of X to R). If φ is an R–positive formula, Φ_A is monotone in the sense that for all $X \subseteq Y \subseteq |A|^k, \Phi_A(X) \subseteq \Phi_A(Y)$. We may view φ as determining an induction on A the stages of which are defined as follows: $\varphi_A^0 = \emptyset$; $\varphi_A^{m+1} = \Phi_A(\varphi_A^m)$. Since Φ_A is monotone and A is finite, it follows immediately that for some m, $\varphi_A^m = \varphi_A^{m+1}$. The least such

m is called the *closure ordinal* of φ on A and is denoted $||\varphi||_A$. It is easy to see that $||\varphi||_A \leq l^k$, for a finite structure A of size l (in the case of an infinite structure A, the closure ordinal of an induction may be a transfinite ordinal α whose cardinality is equal to the cardinality of $|A|$). Moreover, one can readily verify that for $m = ||\varphi||_A, \varphi_A^m$ is the *least fixed point* of the relational operator Φ_A, that is, $\Phi_A(\varphi_A^m) = \varphi_A^m$ and for all $X \subseteq |A|^k$, if $\Phi_A(X) = X$, then $\varphi_A^m \subseteq X$. We use φ_A^∞, to denote the least fixed point of the operator Φ_A. For example, if $\chi(R, x, y)$ is the formula

$$Exy \vee \exists z(Exz \wedge Rzy)$$

then for every structure $A \in D$, χ_A^∞ is the transitive closure of E^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, $\text{lfp}(R, x_1, \ldots, x_k)\varphi$ is a formula of FO+LFP and $A \models \text{lfp}(R, x_1, \ldots, x_k)\varphi[\overline{a}]$, if and only if, $\overline{a} \in \varphi_A^\infty$.

Let us attend once again to reachability. For $\chi(R, x, y)$ as above, the sentence $\text{lfp}(R, x, y)\chi(s, t)$ defines $\mathcal{R}$ relative to $\mathcal{D}$. This approach to the definability of $\mathcal{R}$ has been regarded as more constructive than the Fregean approach described in the preceding section: many find the general notion of iteration to be more transparent than universal second order quantification. Since, as we will see in the next section, FO+LFP($\mathcal{D}$) is *properly* included in

$\mathsf{P}(\mathcal{D})$, the "more constructive" approach actually yields a stronger bound on the descriptive complexity of $\mathcal{R}$. It is interesting to observe, as a corollary of Fagin's Theorem and the Immerman-Vardi Theorem, that in the case of finite ordered structures, the relative power of first order positive induction versus universal second order quantification amounts exactly to the question whether $\mathsf{P} = \mathsf{NP}$.

Let us look a bit more carefully at the case of ordered structures. For simplicity, let's focus on the set $\mathcal{O} \subseteq \mathcal{F}_{\{E,<\}}$ of ordered graphs – a structure A is a member of $\mathcal{O}$, if and only if, the reduct of A to $\{E\}$ is in $\mathcal{D}$ and the reduct of A to $\{<\}$ is in $\mathcal{L}$, the set of finite linear orders. The Immerman-Vardi Theorem tells us that $\text{FO+LFP}(\mathcal{O}) = \mathsf{P}(\mathcal{O})$. It follows from the results of Section 4 that the set of ordered graphs of odd size, a query in $\mathsf{P}(\mathcal{O})$, is not first order definable relative to $\mathcal{O}$. We may conclude that that $\text{FO}(\mathcal{O})$ is properly included in $\text{FO+LFP}(\mathcal{O})$. In fact, there is no known example of an infinite query $\mathcal{Q} \subseteq \mathcal{O}$ such that $\text{FO}(\mathcal{Q}) = \text{FO+LFP}(\mathcal{Q})$. Kolaitis & Vardi (1992a) conjectured that for every infinite query $\mathcal{Q} \subseteq \mathcal{O}$, $\text{FO}(\mathcal{Q})$ is properly included in $\text{FO+LFP}(\mathcal{Q})$. This Ordered Conjecture is an important open problem in finite model theory which turns out to have connections to a number of open problems in the theory of computational complexity. Even the special case of this conjecture concerning the power of first order

versus fixed point definability relative to the set $\mathcal{FR}$ of finite ranks of the cumulative hierarchy of sets is open, and its resolution would have significant complexity theoretic consequences. (see Dawar, Lindell, & Weinstein (1995) and Gurevich, Immerman, & Shelah (1994)). (This counts as a special case, since a linear order is uniformly first order definable on the structures in $\mathcal{FR}$, see Dawar, Doets, Lindell, & Weinstein (1998).)

The Ordered Conjecture asks whether there is an infinite set of finite ordered structures relative to which first order logic characterizes polynomial time computability. If we turn our attention away from ordered structures, we can formulate what has been regarded as the central open problem of descriptive complexity theory, namely: Is there a logical characterization of polynomial time computability over structures without a built-in order? Gurevich (1988) has given a rigorous formulation of this question. In connection with Fagin's Theorem, we noted that existential second order logic characterizes NP in a strong sense – not only is $\mathsf{NP}(\mathcal{F}_\sigma) = \Sigma^1_1(\mathcal{F}_\sigma)$, for all σ; there is an effective procedure for transforming sentences of existential second order logic into polynomially clocked nondeterministic Turing machines that witness the membership of the queries they define in NP. Likewise, in the case of P, we can ask if there is a logic $L = \langle S_L, \models_L \rangle$ such that both S_L and $\models_L$ are recursive and

1. $L(\mathcal{F}_\sigma) = \mathsf{P}(\mathcal{F}_\sigma)$;

2. there is an effective procedure F such that for every $\varphi \in S_L$, $F(\varphi)$ is a polynomially clocked deterministic Turing machine which accepts $c(\varphi(\mathcal{F}_\sigma))$.

We call a logic meeting these requirements a logic for P. A logic for P amounts to an effective list of polynomially clocked deterministic Turing machines, each of which decides a query, and which lists at least one machine deciding each query in P. The difficulty in constructing such an effective list lies in the requirement that the machines must decide queries, that is, isomorphism invariant sets of structures. The set of machines meeting this requirement is not recursively enumerable. This is not fatal to the enterprise of constructing a logic for P, since we do not need to enumerate all the polynomially clocked, isomorphism invariant machines, just a rich enough subset of them. An obvious way to proceed would be as follows. A function $C : \mathcal{D} \mapsto \mathcal{D}$ is called a graph canon, if and only if,

1. $\forall G \in \mathcal{D}(G \cong C(G))$, and

2. $\forall G, H \in \mathcal{D}(G \cong H \rightarrow C(G) = C(H))$.

A graph canon extracts a unique representative from each equivalence class of $\mathcal{D}$ under the equivalence relation of isomorphism. If there is a graph canon

C that is computable in polynomial time, then there is a logic for P. This is easily seen by composing C with an effective list of polynomially clocked deterministic Turing machines which, for each set of strings $X \in \mathsf{P}$, includes a machine which decides X – such an effective list can be constructed absent the requirement that the machines decide queries. It is well-known that if $\mathsf{P} = \mathsf{NP}$, then there is a polynomial time computable graph canon, which yields the conclusion that if there is no logic for P, then $\mathsf{P} \neq \mathsf{NP}$. There is no evidence that the converse holds, and the quest for a logic for P remains an active area of research in descriptive complexity theory.

7 Infinitary Logics

In this section, we investigate a measure of logical complexity that has played a prominent role in recent research in finite model theory. The measure is the total number of variables, both free and bound, which occur in a formula of first order logic, or its infinitary extension, $L_{\infty\omega}$. First order sentences which involve the reuse of bound variables within the scopes of quantifiers already binding those same variables are generally frowned on from a pedagogical and stylistic point of view. Thus, the study of finite variable fragments of first order logic and infinitary logic, whose point is to exploit the possibility of such reuse, typically seems a bit unusual, if not perverse, to most logicians.

Consider the following sequence of first order sentences, each of which contains occurrences of only the two variables x_1 and x_2 :

$$\begin{array}{ll} \varphi_0 & Est \\ \varphi_1 & \exists x_1(Esx_1 \wedge Ex_1t) \\ \varphi_2 & \exists x_1\exists x_2(Esx_1 \wedge Ex_1x_2 \wedge Ex_2t) \\ \varphi_3 & \exists x_1\exists x_2(Esx_1 \wedge Ex_1x_2 \wedge \exists x_1(Ex_2x_1 \wedge Ex_1t)) \\ \varphi_4 & \exists x_1\exists x_2(Esx_1 \wedge Ex_1x_2 \wedge \exists x_1(Ex_2x_1 \wedge \exists x_2(Ex_1x_2 \wedge Ex_2t))) \\ \vdots & \vdots \end{array}$$

Clearly, the sentences φ_i are pairwise inequivalent (consider the structures A_n for $n > 1$ which interpret E as the successor relation on $[n]$ and assign 1

to s and and n to t; $A_n \models \varphi_i$, if and only if, $i+2=n$). Note that although the sentences involve only two variables, their quantifier rank is unbounded. Needless to say, these sentences cannot be brought to prenex normal form without increasing the number of variables.

The logic $L_{\infty\omega}$ is the infinitary extension of first order logic which is closed under the formation of arbitrary conjunctions and disjunctions of sets of formulas. In Section 2, we observed that every finite structure is characterized up to isomorphism by a single first order sentence, from which it follows that for every signature σ, every query $\mathcal{Q} \subseteq \mathcal{F}_\sigma$ is $L_{\infty\omega}$ definable. Thus, $L_{\infty\omega}$ is too strong to be of interest from the point of view of finite model theory. Let us consider the weaker finite variable fragments of $L_{\infty\omega}$. We define $L^k_{\infty\omega}$ to be the k-variable fragment of $L_{\infty\omega}$, that is, $L^k_{\infty\omega}$ consists of all formulas of $L_{\infty\omega}$ all of whose individual variables, either free or bound, are among $x_1, \ldots, x_k$. We let $L^\omega_{\infty\omega} = \bigcup_{k<\omega} L^k_{\infty\omega}$. For example, let θ, a sentence of $L^2_{\infty\omega}$, be the infinite disjunction of the sentences $\varphi_0, \varphi_1, \ldots$, exhibited above. Observe that θ defines $\mathcal{R}$ (directed reachability) relative to $\mathcal{D}$ (the set of finite directed graphs). This is no accident: Kolaitis & Vardi (1992b) established that for every σ, $\mathrm{FO} + \mathrm{LFP}(\mathcal{F}_\sigma) \subseteq L^\omega_{\infty\omega}(\mathcal{F}_\sigma)$. Thus, the finite variable fragment of infinitary logic provides a tool for analyzing inductive definability over finite structures.

One of the main techniques for studying $L^{\omega}_{\infty\omega}$ definability is the k-pebble game, a variant of the Ehrenfeucht game, essentially due to Barwise (1977). In the k-pebble game, instead of an unlimited supply of pebble pairs, the equipment contains only the pebble pairs $\langle\alpha_1, \beta_1\rangle, \ldots \langle\alpha_k, \beta_k\rangle$. At each round of play, the Spoiler may now either play a pebble from a pair that has not yet been played and place it on the associated board, or move a pebble that has already been played to a new position. As before, the Duplicator must follow by moving the matched pebble on the other board. The winning condition for the n-round game remains the same as before. There is also an infinite version of the k-pebble game which we call the eternal k-pebble game. In this version, play continues through a sequence of rounds of order type ω. The Spoiler wins a play of the eternal game, if and only if, he wins at some finite round; otherwise, the Duplicator wins. We say that structures A and B are indistinguishable by sentences of $L^{k}_{\infty\omega}$ ($A \equiv^{k}_{\infty\omega} B$), if and only if, for every sentence $\varphi \in L^{k}_{\infty\omega}$,

$$A \models \varphi \iff B \models \varphi.$$

Barwise proved that the Duplicator has a winning strategy for the eternal k-pebble game played on A and B, if and only if, $A \equiv^{k}_{\infty\omega} B$. Thus, we can show that a query $\mathcal{Q} \subseteq \mathcal{D}$ is not $L^{k}_{\infty\omega}$ definable by exhibiting structures $A, B \in \mathcal{D}$, such that $A \in \mathcal{Q}, B \notin \mathcal{Q}$, and the Duplicator has a winning strategy for the

eternal k-pebble game played on A and B.

As an illustration of this technique, we show that $\mathsf{P}(\mathcal{D}) \not\subseteq L^{\omega}_{\infty\omega}(\mathcal{D})$. We say that $A \in \mathcal{D}$ is an *empty graph*, if and only if, $E^A = \emptyset$. It is easy to see, by playing the k-pebble game, that for all empty graphs A and B, if A and B both have at least k nodes, then $A \equiv^k_{\infty\omega} B$. It follows at once that the set of graphs which have an odd number of nodes, a query in P, is not definable in $L^{\omega}_{\infty\omega}$. It also follows that the languages $L^k_{\infty\omega}$ form a strict hierarchy in terms of expressive power relative to $\mathcal{D}$. We will meet $L^{\omega}_{\infty\omega}$ again in the next section.

8 Random graphs and 0-1 laws

In this section, we will take up some connections between finite model theory and combinatorics. We focus attention on the study of random graphs, an active area of research in contemporary combinatorics.

Random graphs Consider the following procedure for determining a directed graph with node set $[n]$. For each of the n^2 ordered pairs of nodes flip a fair coin to determine whether or not there is a directed edge from the first to the second; we assume the outcomes of the tosses are mutually independent. For each n, this procedure gives rise to the uniform probability distribution over $\mathcal{D}_n$, the collection of directed graphs with node set $[n]$. We may use this probability distribution to answer questions about how many graphs there are with certain properties. We write $\Pr_n(\theta)$ for the probability (with respect to this distribution) that a graph with node set $[n]$ satisfies θ. Note that,

$$\Pr_n(\theta) = \frac{\text{card}\{G \in \mathcal{D}_n \mid G \models \theta\}}{\text{card}\mathcal{D}_n}.$$

We will be interested in the behavior of $\Pr_n(\theta)$ as a function of n for various choices of θ. We write $\Pr(\theta) = \lim_{n\to\infty} \Pr_n(\theta)$. In general, $\Pr(\theta)$ may not be defined. For example, when $\theta \in \Sigma^1_1$ expresses the condition that there are an even number of nodes, $\Pr_n(\theta)$ endlessly oscillates between the

values 0 and 1 and thus has no well defined limit. On the other hand, many interesting graph theoretic properties do possess a "limit probability" with respect to the uniform distribution. We will see how logic provides some explanation of this fact.

Let us begin with the example of connectivity: a directed graph A is connected, if and only if, for each pair i, j of distinct nodes of A, there is a path from i to j. Let θ be the sentence of FO+LFP that defines the set of connected graphs relative to $\mathcal{D}$. We wish to discover whether $\Pr(\theta)$ is well defined, and if it is, whether we can determine its value. In order to do so, we will attempt to approximate the value of $\Pr_n(\theta)$ for large values of n.

Rather than dealing directly with θ, let us consider the following first order condition which implies θ. Let φ be the following sentence:

$$(\forall x)(\forall y)(x \neq y \rightarrow (\exists z)(x \neq z \wedge y \neq z \wedge (Exz \wedge Ezy).$$

The sentence φ expresses the "two degrees of separation" property – we can proceed from any node to any other by a path of length two. Clearly, φ implies θ. Hence, for all n,

$$\Pr_n(\varphi) \leq \Pr_n(\theta).$$

Therefore, if we can show that $\Pr_n(\varphi)$ becomes large, as a function of n, the same will be true of $\Pr_n(\theta)$.

Let's perform the calculation. Fix a pair of distinct nodes $i, j \in [n]$. We say that a node k links i to j, if and only if, there is an edge from i to k and and edge from k to j. Clearly, for any fixed node k, distinct from i and j, the probability that k does not link i to j is .75. So the probability that no node distinct from i and j links i to j is $(.75)^{n-2}$. Now, there are $n(n-1)$ ordered pairs of distinct nodes in $[n]$. Therefore, the probability that some pair of distinct nodes in $[n]$ fail to be linked is bounded from above by $n(n-1) \cdot (.75)^{n-2}$. That is,

$$\Pr_n(\neg\varphi) \leq n(n-1) \cdot (.75)^{n-2}.$$

It is easy to show that

$$\lim_{n\to\infty} n(n-1) \cdot (.75)^{n-2} = 0.$$

It follows at once that

$$\Pr(\theta) = \Pr(\varphi) = 1.$$

So we have succeeded in analyzing the limiting behavior of graph connectivity by reducing the problem to a simple calculation of the limiting behavior of a first order condition; and the limit probability of that condition is 1. To what extent can we generalize this example?

0-1 Laws In this section we will consider a sweeping generalization of the preceding example of connectivity. We say that a logical language L satisfies

the 0-1 law with respect to the uniform distribution over directed graphs, if and only if, for every sentence φ of L,

$$\Pr(\varphi) = 0 \quad \text{or} \quad \Pr(\varphi) = 1.$$

A bold generalization of the example of connectivity would be the following: FO+LFP satisfies the 0-1 law for the uniform distribution over directed graphs. Indeed, this generalization is true, as was established by Blass, Gurevich, and Kozen (1985). This result itself generalized the 0-1 law for first order logic due to Fagin (1976) and Glebskij (1969). A striking generalization of these (and additional) results, which provides a beautiful explanation for the limiting behavior of a variety of graph theoretic properties, is the following 0-1 law for $L^{\omega}_{\infty\omega}$ due to Kolaitis & Vardi (1992b): $L^{\omega}_{\infty\omega}$ satisfies the 0-1 law for the uniform distribution over directed graphs. Not only does this result generalize the example of connectivity given above; it's proof also follows the lines of the argument given for the example. In particular, the theorem is a corollary of the following fascinating result, also due to Kolaitis & Vardi (1992b): For every $k \geq 2$, there is a k-variable first order sentence γ_k such that

1. $\Pr(\gamma_k) = 1$, and

2. for every sentence $\theta \in L^{k}_{\infty\omega}$, either $\gamma_k \models \theta$ or $\gamma_k \models \neg\theta$.

In other words, for each k, there is a single first order sentence which has limit probability 1 with respect to the uniform distribution on directed graphs and axiomatizes a complete $L^k_{\infty\omega}$ theory.

The sentence γ_k may be constructed as follows. A k-literal is a formula of the form Ex_ix_j or its negation with $1 \leq i, j \leq k$. A basic k-type is a maximal consistent conjunction of k-literals. A k-extension condition is a sentence of the form:

$$\forall x_1 \ldots \forall x_{k-1}((\bigwedge_{i \neq j} x_i \neq x_j \wedge \varphi) \rightarrow \exists x_k(\bigwedge_{i<k} x_i \neq x_k \wedge \psi)),$$

where φ is a $(k-1)$-type, ψ is a k-type, and ψ extends φ. A graph satisfies such a k-extension condition, if and only if, each of its size $k-1$ subgraphs of type φ can be extended to a size k subgraph of type ψ. We let γ_k be the conjunction of all the l-extension conditions for $2 \leq l \leq k$. The sentence γ_k expresses a "bounded principle of plenitude:" every subgraph of size $l < k$ can be extended in every possible way to a subgraph of size $l+1$ (compare the two degrees of separation principle above). For $k \geq 3$, it is not at first sight obvious that there are finite structures which satisfy γ_k. However, an easy computation, of just the sort sketched for the two degrees of separation principle, reveals that $\Pr(\gamma_k) = 1$ for all $k \geq 2$. That is, for every $\varepsilon > 0$, for large enough n, all but an ε fraction of the directed graphs of size n satisfy γ_k.

In order to verify that γ_k axiomatizes a complete $L^k_{\infty\omega}$ theory, it suffices to show that for all directed graphs A, B, if $A \models \gamma_k$ and $B \models \gamma_k$, then $A \equiv^k_{\infty\omega} B$. But this follows directly from Barwise's characterization of $\equiv^k_{\infty\omega}$ given in Section 7, since it is easy to see that the Duplicator has a winning strategy for the eternal k-pebble game played on A and B, if both A and B satisfy γ_k. (Play the game! The description of γ_k as a bounded principle of plenitude is exactly what's required for the Duplicator's strategy.)

Let us call a sentence φ of first order logic *stochastically valid*, if and only if, $\Pr(\varphi) = 1$, and let Sval be the set of stochastically valid sentences of first order logic. It clear from the preceding discussion that $\Gamma = \{\gamma_k \mid k \geq 2\}$ axiomatizes a complete first order theory, a result due to Gaifman (1964). In particular, Γ axiomatizes Sval. It follows at once that Sval is decidable. This provides an interesting contrast to the results described in Section 1.

9 Further Reading

Two excellent texts are available which cover the topics presented here in depth. They are Ebbinghaus & Flum (1999) and Immerman (1999). An invaluable introduction to the theory of computational complexity is Papadimitriou (1994). For readers wishing further background on finite variable logics there are valuable survey articles by Dawar (1999) and Grohe (1998) and an excellent monograph by Otto (1997). An excellent introduction to the theory of random graphs is Alon & Spencer (1992).

Acknowledgements The preparation of this work was supported in part by NSF CCR-9820899. I would like to thank the Graduate Program in Logic and Algorithms at the University of Athens for support while on leave from the University of Pennsylvania and for providing the stimulating research environment in which my work on this paper was completed. I am especially grateful to Steven Lindell for a decade of valuable discussions on the topics of this paper and to Mary-Angela Papalaskari for valuable comments on earlier drafts of this paper.

SCOTT WEINSTEIN

References

- Ajtai, M. and Fagin, R. (1990). Reachability is harder for directed than for undirected finite graphs. *Journal of Symbolic Logic*, 55:113–150.

- Alon, N. and Spencer, J. (1992). *The Probabilistic Method*. New York: John Wiley and Sons.

- Barwise, J. (1977). On Moschovakis closure ordinals. *Journal of Symbolic Logic*, 42:292–296.

- Blass, A., Gurevich, Y., and Kozen, D. (1985). A zero-one law for logic with a fixed point operator. *Information and Control*, 67:70–90.

- Dawar, A. (1999). Finite models and finitely many variables. In Niwinski, D. and Maron, R., editors, *Logic, Algebra and Computer Science*, volume 46 of *Banach Center Publications*, pages 93–117. Polish Academy of Sciences.

- Dawar, A., Doets, D., Lindell, S., and Weinstein, S. (1998). Elementary properties of the finite ranks. *Mathematical Logic Quarterly*, 44:349–353.

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

- Ebbinghaus, H.-D. and Flum, J. (1999). *Finite Model Theory.* Berlin: Springer-Verlag.

- Ehrenfeucht, A. (1961). An application of games to the completeness problem for formalized theories. *Fund. Math.*, 49:129–141.

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

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

- Fraïssé, R. (1954). Sur quelques classifications des systèmes de relations. *Publications Scientifiques de l'Université d'Algerie, Séries A*, 1:35–182.

- Frege, G. (1884). *Die Grundlagen der Arithmetik.* Breslau: Wilhelm Koebner.

- Gaifman, H. (1964). Concerning measures in first-order calculi. *Israel Journal of Mathematics*, 2:1–18.

- Gaifman, H. and Vardi, M. (1985). A simple proof that connectivity of finite graphs is not first order. *Bulletin of the EATCS*, pages 43–45.

- Glebskij, Y., Kogan, D., Liogon'kij, M., and Talanov, V. (1969). Range and degree of realizability of formulas in the restricted predicate calculus. *Cybernetics*, 5:142–154.

- Grohe, M. (1998). Finite variable logics in descriptive complexity theory. *Bulletin of Symbolic Logic*, 4:345–398.

- Gurevich, Y. (1984). Toward logic tailored for computational complexity. In Richter, M. et al., editors, *Computation and Proof Theory*, pages 175–216. Heidelberg: Springer-Verlag.

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

- Gurevich, Y., Immerman, N., and Shelah, S. (1994). McColm's conjecture. In *Proceedings of the 9th IEEE Symposium on Logic in Computer Science*, pages 10–19.

- Immerman, N. (1986). Relational queries computable in polynomial time. *Information and Control*, 68:86–104.

- Immerman, N. (1999). *Descriptive Complexity*. New York: Springer-Verlag.

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

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

- Moschovakis, Y. N. (1974). *Elementary Induction on Abstract Structures.* Amsterdam: North Holland.

- Otto, M. (1997). *Bounded Variable Logics and Counting.* Berlin: Springer-Verlag.

- Papadimitriou, C. (1994). *Computational Complexity.* Reading: Addison-Wesley.

- Rosen, E. (1997). Modal logic over finite structures. *Journal of Logic, Language, and Information*, 6:427–439.

- Trakhtenbrot, B. A. (1950). Impossibility of an algorithm for the decision problem in finite classes. *Dokdaly Akademii Nauk SSSR*, 70:569–572.

- Vardi, M. Y. (1982). The complexity of relational query languages. In *Proceedings of the 14th ACM Symposium on the Theory of Computing*, pages 137–146.