

available at www.sciencedirect.comjournal homepage: www.elsevier.com/locate/cosrev

Book review

E. Grädel, P. Kolaitis, L. Libkin, M. Marx, I. Spencer, M. Vardi, Y. Venema, S. Weinstein. *Finite Model Theory and Its Applications*. Springer-Verlag, (2007)

1. Introduction

It is tempting to start by defining the subject the book is about, in this case finite model theory, in my own words. I succumb to the temptation. Finite model theory is concerned with studying, classifying, and making use of formal means for expressing properties of finite mathematical structures.

This was risky. Defining a broad subarea of mathematics in a short sentence is inevitably oversimplifying the matter, which has the danger of rendering the definition useless, misleading, or both. And in fact, did I need to succumb to the temptation? It might be more appropriate to discuss what the preface of the book itself says about its subject. I look for it and I find a wise reference to the definition of model theory from the classic on the topic by Chang and Keisler [3]: “the branch of mathematical logic which deals with the relation between a formal language and its interpretations”. If the finiteness is to be understood on the interpretations, also known as models or structures, this is quite close to what I had in mind.

But the preface, which is unusually short for a good reason to be understood later, actually starts not by defining what finite model theory is about, but by asserting that it is “an area of mathematical logic that has developed in close connection with applications to computer science”. This explains the conjunctive second component in the book title, “Finite Model Theory and Its Applications”. And it is in fact this emphasis on “applications”, whatever it means for the moment, that gives this book a distinguishing feature over previous books on finite model theory. I will come back to it.

1.1. About the authors and book organization

Before discussing the actual contents and goals of the book, and the degree to which it realizes them, I want to comment briefly on the authors and on the book organization.

This is an eight-authored monograph: Erich Grädel, Phokion Kolaitis, Leonid Libkin, Maarten Marx, Joel Spencer,

Moshe Vardi, Yde Venema, and Scott Weinstein. There is no need to investigate the prominence of these authors in their respective areas of expertise. What requires comment is how a text with such a long list of authors is actually organized.

All eight authors sign the monograph, but each chapter appears signed by one or two of them, and has its own list of references and particular style of writing. Thus, while the book does not seem to be conceived as an edited collection of independent articles on a common subject, it is certainly not a cohesive single text with a rigid and uniform style throughout either. In this respect, the book organization is somewhat non-standard and innovative. If it were not for its shorter length, this format and concept of text does have some resemblance with the notably successful genre commonly known as *handbook* [1,8,5,2].

The criterion that divides the text into chapters is what the authors call “applications”, which I modestly understand by the less glamorous word “opportunities”. What I mean by this is that the applications the authors build on are nothing but legitimate (and peaceful) incursions of finite model theory into diverse subareas of mathematics and computer science that very often look separate from logic itself.

1.2. About the declared goals

From the table of contents, it is quite clear the volume does not aim to be a textbook on finite model theory. In fact, the preface refers to the three textbooks in finite model theory published to date [4,7,6] to claim that the goal of this volume is different. As the preface puts it, it “aims at highlighting applications” of finite model theory. To be precise, the theme of the book is built around some of the successful insight provided by finite model theory to five areas: computational complexity, probabilistic combinatorics, databases, constraint solving and (specification through) modal logics. This sort of structure traces a significant fraction of the development of finite model theory to its motivations, and is coherent with the thesis raised in the first sentence of the preface referred to above.

2. Summary of contents

The book is organized in seven chapters. The first chapter is in some sense an overview. It could well serve as a long preface,

and it could also play the role of a descriptive chapter-by-chapter review. This explains the short length of the preface referred to before. It also gives me an excellent excuse not to accurately describe the contents of each chapter here. Instead, the rest of this review will focus on commenting on some selected passages. The leading criterion that I followed to select a passage is that I have found it either particularly original, or particularly instructive, or both.

2.1. Weinstein on unifying themes

Under the title “Unifying Themes in Finite Model Theory”, Scott Weinstein offers the pieces that link the different topics that the rest of the chapters will develop. Written beautifully, in the literary style allowed by the expository nature of the chapter, Weinstein aims at captivating the reader through the aesthetics of analogy. For those who, like this reviewer, appreciate this style of exposition, the early example illustrating the benefits of definability theory through a celebrated result in descriptive set theory, will be nice to read.

It may sound from this that reading even the first chapter will require some solid prior knowledge in mathematical logic. It turns out that not much is needed to understand what is going on, but it would not be accurate to say that no background is needed to appreciate all the parts. This also applies to other chapters of the book. At any rate, this initial chapter is a good sample where readers can test if the assumed background meets their own.

2.2. Kolaitis on expressive power

The second chapter of the book, titled “On the Expressive Power of Logics on Finite Models” and signed by Phokion Kolaitis, serves the goal of introducing the main concepts and tools of finite model theory as a field. The chapter starts explaining the fundamental concepts of definability and uniform definability at an appropriate level of abstraction. This is followed by a careful exposition of Ehrenfeucht-Fraïssé games for first-order logic, including a complete proof that the games provide a purely combinatorial characterization of the expressive power of first-order logic, and some illustrative examples.

Within the section on logics with fixed-point operators, the author offers an exposition of one of the early results of finite model theory that gives its distinguished character as a field. This is Immerman’s normal form result for least fixed-point logic on finite structures. I will describe the result trying not to spoil the pleasure of reading the passage from the book itself.

The syntax of least fixed-point logic (LFP) extends that of first-order logic with a formula formation rule to denote least fixed-points of positive formulas. More formally, the class of LFP-formulas is the smallest class of formulas that contains all atomic formulas and is closed under negation, conjunction, disjunction, existential and universal quantification, and such that if $\varphi(x, X)$ is a previously formed formula where x is a tuple of k first-order variables and X is a k -ary second-order variable that occurs within an even number of negations in φ , then $\text{lfp}X.\varphi$ is also a formula. The

semantics given to $\text{lfp}X.\varphi$ on a structure M is the smallest k -ary relation $A \subseteq M^k$ that satisfies the equation $A(x) \leftrightarrow \varphi(x, A)$; that is, the least fixed-point of φ . Such a relation always exist by the celebrated Knaster–Tarski Theorem, which is actually proved in an earlier section of Kolaitis’ chapter.

We are now ready to state Immerman’s Theorem as stated in Kolaitis’ chapter. Let LFP_1 be the fragment of LFP that consists of all formulas that contain at most one application of lfp which moreover appears positively only (within an even number of negations). Immerman’s Theorem states that on every class of finite structures \mathcal{C} , every LFP-formula is equivalent on \mathcal{C} to an LFP_1 -formula. In particular, LFP_1 -definable subclasses of \mathcal{C} are closed under complementation. Quite significantly, this important normal form result is true only of classes of finite structures as it is known to fail otherwise.

Besides giving a nearly complete proof of the normal form result, Kolaitis provides very useful historical perspective. The result refuted a conjecture by Chandra and Harel that arose in the context of database theory and which was inspired by the celebrated Kleene-Spector Theorem of recursion theory. I strongly recommend reading this passage in detail and, for those interested, tracing the classical references. Last but not least, the proof of Immerman’s Theorem is a nice application of the beautiful Stage Comparison Theorem, also discussed in the chapter, which is really worth knowing about. Finding out what is special about finite structures that makes this application possible is a pleasure to read.

2.3. Grädel on descriptive complexity

The third chapter of the book is about “Finite Model Theory and Descriptive Complexity” and is signed by Erich Grädel. The main topic to be studied here is the relationship between logical definability and computational complexity. The well-known example in this class is Fagin’s Theorem, which states that the problems in NP are precisely those that can be expressed in the existential fragment of second-order logic. A full proof of this result is given in detail. This is followed by the result due to Grädel himself that the problems in P are precisely those expressible in the Horn fragment of existential second-order logic provided the structures come equipped with a linear order.

There is a relatively subtle detail in this result that might pass unnoticed to some and on which I want to comment. In the original proof of Grädel’s Theorem, the author had to assume that the structures were equipped, not only with a linear order, but also with the successor relation of this linear order and the constants for the maximum and the minimum of the order. Quite pleasantly, the need for the successor and the constants is removed in the current version at the expense of a little (but natural) twist in the definition of the logic that I will not disclose in this review.

One aspect that makes this chapter particularly original is the discussion about model checking games. Actually, Grädel starts the chapter with a discussion of model checking problems, which he describes as occupying a “central place in finite model theory”. This is the problem of deciding, for a given finite structure and a given sentence of a logic, whether the sentence holds in the structure. This

leads him to discuss *model checking games* for first-order logic quite early in the chapter, and by abstraction, the strategy problem for abstract two-player games. Such games reappear in the section on fixed-point logics to discuss the complexity of the model checking problem for a fixed-point logic called μ -calculus, a modal logic of great importance in application areas as diverse as hardware verification and knowledge representation. A complete proof is given for the intriguing result that the model checking problem for the μ -calculus is in $NP \cap co-NP$. This is proved by reduction to the strategy problem for *parity games*, one of the most fascinating problems at the interface between P and NP. The main and fundamental property of parity games, which places its strategy problem in $NP \cap co-NP$, is the so-called Positional Determinacy Theorem, for which a detailed clean proof is provided.

The chapter concludes with a section titled “Algorithmic Model Theory” where the author argues that the general approach and research programme that has led finite model theory should be extended “to interesting domains of infinite structures”, where the word “interesting” is understood with respect to applications. The section surveys some such domains of structures, including a brief reference to the recently developed and promising theory of automatic structures.

2.4. Spencer on logic and random structures

The third chapter “Logic and Random Structures” is written by Joel Spencer. The application area here is the theory of random graphs and structures. The incursion of finite model theory into this area started when Fagin and Glebskii, Kogan, Liagonskii, and Talanov proved that every first-order definable property of graphs holds either in almost all finite graphs, or in almost none. This result is known as the 0-1 law for first-order logic and could not miss this chapter. The given proof is actually quite short and nice as it builds on the powerful machinery developed earlier in the same chapter.

The main theme of the chapter is however not the uniform distribution on finite graphs, but the more general model of random graphs $G(n, p(n))$ introduced by Erdős and Renyi. Here, a random graph on n vertices is generated by placing each possible edge independently with probability $p(n)$, which of course is a real number in the unit interval $(0, 1)$. Random graph theorists study the properties that hold in the graph $G(n, p(n))$ almost surely, which means with probability approaching 1 as n goes to infinity. A typical example, also mentioned in the text, is that if $p(n) = \omega(n^{-2/3})$ then $G(n, p(n))$ contains a K_4 subgraph almost surely, and if $p(n) = o(n^{-2/3})$ then $G(n, p(n))$ does not contain any K_4 subgraph almost surely. What happens at $p(n) = \Theta(n^{-2/3})$? The answer is that, in this case, $G(n, p(n))$ contains a K_4 subgraph with a probability that is bounded away from 0 and 1 and that depends on the hidden constants in the Θ notation.

The chapter starts by developing the necessary machinery to prove the deep theorem of Shelah and Spencer: for every irrational number $\alpha \in (0, 1)$ and every first-order definable property P , the probability that the graph $G(n, n^{-\alpha})$ satisfies P is either asymptotically 1 or asymptotically 0. Thus, we

say that $G(n, n^{-\alpha})$ has the 0-1 law for first-order logic when $\alpha \in (0, 1)$ is irrational. For the curious, if $\alpha \in (0, 1)$ is rational, then $G(n, n^{-\alpha})$ does not have the 0-1 law for first-order logic as it can be easily shown with the classical results on the existence of subgraphs in the Erdős-Renyi model. The proof of the Shelah-Spencer Theorem given in the text is almost complete except for a key and rather difficult probability lemma which is recalled from the references. The required intuitions are beautifully illustrated with concrete examples. Judging from those in the text and those in the previous expositions of the same material by the author, it seems that the author’s favorite example of an irrational number in the unit interval is $\pi/7$.

The chapter continues with a rather advanced material, including a discussion of the models of the set of sentences that hold almost surely in $G(n, n^{-\alpha})$, and a construction of a first-order sentence having an *infinite spectrum*. In this context, the *spectrum* of a first-order sentence is the set of all $\alpha \in (0, 1)$ at which the function $f_P(\alpha) = \lim_{n \rightarrow \infty} \Pr[G(n, n^{-\alpha}) \models P]$ is discontinuous. Of course, there is no way this could be confused with the spectrum of a first-order sentence understood, classically, as the set of cardinalities of its finite models. I thought I would mention this clash of names in this review because the term *spectrum* (actually *spectra*) appears in the table of contents.

2.5. Libkin on embedded and constraint databases

Leonid Libkin’s chapter is about “Embedded Finite Models and Constraint Databases”. Arguably, database theory is the application area where finite model theory has had the greatest impact even at practical levels. Codd’s model of relational databases, together with all the succeeding work in the area, constitutes one of the biggest success-stories of theoretical computer science, logic in computer science, and finite model theory in particular, with respect to real-world computer systems. The current chapter is about a relatively recent development in finite model theory that is relevant for database theory and uses methods of infinite model theory.

An embedded finite model is a finite structure living within a larger, usually infinite, background structure. The typical example, which is actually given at the beginning of the chapter, is a finite database holding relationships between real numbers on which we can also specify common arithmetic relationships between the points such as sums and products. The expressive power of logical languages on such structures is again the main topic of study. How do we prove that certain queries on databases living within some background structure are not expressible in a logic? The chapter is primarily dedicated to the development of a method, which is interesting in its own right, that can be used to answer such questions. The method consists in showing how certain first-order definable queries on embedded databases are also definable without using the background structure (perhaps with the help of a linear order, or some limited access to the background). These methods come under the name of “collapse results”. The chapter studies conditions on the background structure that allow such collapses to be proved.

After discussing the collapse results that follow from the classical arguments in Ramsey theory, the author provides full proofs for the collapse results that follow from the assumption that the background structure satisfies a well-studied condition in infinite model theory known as ω -minimality. Many interesting structures are ω -minimal, such as the field of real numbers even extended with exponential-like functions. These advanced methods give as a result rather impressive consequences that the interested reader will surely enjoy. For those eager to see further applications of infinite model theory to “databases”, the author offers collapse results for structures beyond ω -minimality using notions of infinite model theory such as saturation.

The deep collapse results due to Baldwin and Benedikt on structures with finite VC-dimension are only stated, with no proofs. However, the author provides the proof of an interesting and surprising result which says that a certain type of collapse known as “restricted quantifier collapse” implies finite VC-dimension in the background structure. The proof of this result is as well surprising, as it works by making use of the deep result in Boolean circuit complexity that non-uniform bounded-depth circuits of polynomial-size are not able to compute NP-complete problems. The most amazing part of this application is that the non-uniformity in the circuit lower-bound seems to be essential.

2.6. Kolaitis and Vardi on constraint satisfaction problems

The chapter titled “A Logical Approach to Constraint Satisfaction” is signed by two authors: Phokion Kolaitis and Moshe Vardi. In brief, a constraint satisfaction problem is given by a set of variables, a domain of values that the variables can take, and a set of constraints each specifying the set of allowed tuples of values for a given tuple of variables. A typical example is graph k -colorability: each vertex is seen as a variable taking one of k possible values, and each edge specifies the constraint that the end-points should take different values. Another example is the satisfiability of systems of equations over a ring.

Following the influential work by Feder and Vardi, the authors start noting how the problem may be rephrased as the homomorphism problem between relational structures: given two relational structures A and B , is there a mapping from the universe of A to the universe of B that preserves the tuples of the relations? Here, the universe of A stands for the variables of the problem, the universe of B stands for the values these variables can take, and the relations of A and B encode the constraint-scopes and the constraint-relations, respectively. The restriction of this problem when A and B are taken from classes of structures \mathcal{A} and \mathcal{B} respectively is denoted by $\text{CSP}(\mathcal{A}, \mathcal{B})$. The authors then phrase the grand classification challenge: “identify or characterize classes \mathcal{A} and \mathcal{B} of relational structures such that $\text{CSP}(\mathcal{A}, \mathcal{B})$ is solvable in polynomial time”.

The reformulation of the constraint satisfaction problem as the homomorphism problem for structures puts us in the world of finite model theory. We can ask for model-theoretic properties of the classes \mathcal{A} and \mathcal{B} that lead to tractability. The authors discuss some such conditions, and in particular one that has led a satisfactory solution of the grand challenge for

problems of the form $\text{CSP}(\mathcal{A}, -)$, where $-$ denotes the class of all finite structures. We discuss this next.

An early and easy observation in the area of constraint satisfaction problems is that in restriction to tree-shaped instances, that is, when the graph underlying the instance is a tree, the problem is tractable. The authors discuss how this has been generalized significantly through the notion of treewidth from the work on graph minors by Robertson and Seymour. The important result here is that on instances of treewidth k , the constraint satisfaction problem can be solved in polynomial time. This result is discussed in the book and re-interpreted in logical terms through finite-variable logics. The necessary background can be found in the earlier chapter by Kolaitis, which comes in handy. The flow of the chapter moves us to analyzing this connection further to find out that the logical re-interpretation provides, in a natural way, a slightly larger class of instances for which the CSP is tractable. These are the instances that are homomorphically equivalent to an instance of treewidth k . Mysteries of science, it turns out that this slightly larger class is optimal by an important result of Grohe also discussed in the text. But perhaps this is not that mysterious; after all, the authors’ thesis all throughout the chapter is that finite-variable logics are indeed at the heart of the deepest questions in this area. As a final remark on this chapter, let us point out that the emphasis is judiciously put in making the logic of the text flow, which explains why it contains no proofs at all. The generous list of 69 references, many of which quite recent, will guide the reader who needs to go deeper and see the proofs.

2.7. Marx and Venema on modal logic

Maarten Marx and Yde Venema close the book with the chapter titled “Local Variations on a Loose Theme: Modal Logic and Decidability”.

Modal logic is understood classically as the extension of propositional logic with modal operators that express the possibility or the necessity of a formula being true in the *reachable worlds*. If we understand *worlds* as *states* of a system, we are facing a powerful specification language with applications as diverse as program and hardware verification, game theory, or even natural language processing. With this broadness in mind, the authors introduce modal logic abstractly, as a triple formed by a set of formulas, a set of models, and a satisfaction relation, where the notion of bisimulation formalizes indistinguishability.

The chapter pivots around two fundamental principles that the authors call *looseness* and *locality*. These state that the reasons for the class of modal logics to be so well-behaved computationally is to be found, on the one hand, on the fact that bisimulation guarantees that the relevant models are “loose”, in the sense that they are tree-shaped, and on the other hand, on the fact that the language is typically “local”, which means that it is unable to express global properties of the models. While the principle of looseness was widely recognized a long time ago thanks to the influential work of Vardi, the principle of locality for modal logics is one of the hot topics in the area. At its heart we find again finite-variable

logics, nicely gluing the topics of this chapter with those that precede it.

The theorems in the chapter are beautifully linked together with useful discussions on the intuitions that underlie each new idea and definition. As an example, I want to mention the interesting discussion on modal logics with grid-like models. These are at the extreme of non-looseness, as grids are hardly tree-shaped. With the aim of identifying the fundamental reasons for computational tameness, the authors adventurously explore its border showing how small variations on such logics can cause the satisfiability problem for the logic to switch from being decidable to being undecidable.

It is remarkable that this chapter contains complete proofs of all its claims, except for a few hardness results, which is not the main focus anyway. As a final remark, I want to point out how the inclusion of this chapter makes this book significantly different from any other book in finite model theory. It might be for this reason that I felt I had learned a good deal of methods, techniques, and intuitions by the time I had finished reading it.

3. Concluding remarks

Closing a book by its back cover is always pleasant. But it is more so if you have the feeling that you learned new things as you went along. I think it is safe to say that even the experts in finite model theory will have this feeling when closing their back cover.

The analogies by Weinstein, the historical perspective by Kolaitis, the originality through model-checking games by Grädel, the advanced topics in probability theory by Spencer, the fascinating connections to classical model theory by Libkin, the logical flow that shapes a field by Kolaitis and Vardi, or the intuitions and full proofs by Marx and Venema: these are only selected aspects that I personally liked from each chapter. No need to say that, as a reader, you will find your own. On the other hand, it could well be that my selection made this review technical. In this case, keep in

mind my suggestion to have a look at the first chapter to get a representative sample. And even if you feel that the assumed background might not fit your own, keep in mind that this is not a book conceived as a reference textbook that you need to read linearly from cover to cover. There is sufficient overlap between the chapters.

Since its initial steps, finite model theory has followed an interesting route of development. Its focus has shifted back and forth between fundamental questions, understood as questions of purely mathematical and intellectual interest, and applications, understood in the sense discussed in the beginning of this review. The choice in this monograph is to present finite model theory to its readers from the perspective of its applications. This is a great honorable goal, complementing the previous literature on the topic, that also serves the good purpose of educating a sense of historical perspective and broadness.

REFERENCES

- [1] J. Barwise (Ed.), *Handbook of Mathematical Logic*, North-Holland, 1977.
- [2] S.R. Buss (Ed.), *Handbook of Proof Theory*, Elsevier Science Publishers, 1998.
- [3] C.C. Chang, H.J. Keisler, *Model Theory*, North Holland, 1973.
- [4] H.-D. Ebbinghaus, J. Flum, *Finite Model Theory*, second edition, Springer-Verlag, 1995.
- [5] R.L. Graham, M. Grötschel, L. Lovász (Eds.), *Handbook of Combinatorics*, vol. I, II, Elsevier Science Publishers, 1995.
- [6] N. Immerman, *Descriptive Complexity*, Springer-Verlag, 1998.
- [7] Leonid Libkin, *Elements of Finite Model Theory*, Springer-Verlag, 2004.
- [8] J. van Leeuwen (Ed.), *Handbook of Theoretical Computer Science*, vol. A, B, Elsevier Science Publishers, 1990.

Albert Atserias
 Universitat Politècnica de Catalunya,
 Barcelona, Spain
 E-mail address: atserias@lsi.upc.edu.