

Descriptive Complexity

Neil Immerman

February 25, 2015

first printing: Jan., 1999, Springer Graduate Texts in Computer Science

This book is dedicated to Daniel and Ellie.

Preface

This book should be of interest to anyone who would like to understand computation from the point of view of logic. The book is designed for graduate students or advanced undergraduates in computer science or mathematics and is suitable as a textbook or for self study in the area of descriptive complexity. It is of particular interest to students of computational complexity, database theory, and computer aided verification. Numerous examples and exercises are included in the text, as well as a section at the end of each chapter with references and suggestions for further reading.

The book provides plenty of material for a one semester course. The core of the book is contained in Chapters 1 through 7, although even here some sections can be omitted according to the taste and interests of the instructor. The remaining chapters are more independent of each other. I would strongly recommend including at least parts of Chapters 9, 10, and 12. Chapters 8 and 13 on lower bounds include some of the nicest combinatorial arguments. Chapter 11 includes a wealth of information on uniformity; to me, the low-level nature of translations between problems that suffice to maintain completeness is amazing and provides powerful descriptive tools for understanding complexity. I assume that most readers will want to study the applications of descriptive complexity that are introduced in Chapter 14.

Map of the Book

Chapters 1 and 2 provide introductions to logic and complexity theory, respectively. These introductions are fast-moving and specialized. (Alternative sources are suggested at the end of these chapters for students who would prefer more background.) This background material is presented with an eye toward the descriptive point of view. In particular, Chapter 1 introduces the notion of queries. In Chapter 2, all complexity classes are defined as sets of boolean queries. (A boolean query

is a query whose answer is a single bit: yes or no. Since traditional complexity classes are defined as sets of yes/no questions, they are exactly sets of boolean queries.)

Chapter 3 begins the study of the relationship between descriptive and computational complexity. All first-order queries are shown to be computable in the low complexity class deterministic logspace (L). Next, the notion of *first-order reduction* — a first-order expressible translation from one problem to another — is introduced. Problems complete via first-order reductions for the complexity classes L, nondeterministic logspace (NL), and P are presented.

Chapter 4 introduces the least-fixed-point operator, which formalizes the power of making inductive definitions. P is proved equal to the set of boolean queries expressible in first-order logic plus the power to define new relations by induction. It is striking that such a significant descriptive class is equal to P. We thus have a natural, machine-independent view of feasible computation. This means we can understand the P versus NP question entirely from a logical point of view: P is equal to NP iff every second-order expressible query is already expressible in first-order logic plus inductive definitions (Corollary 7.23).

Chapter 5 introduces the notion of parallel computation and ties it to descriptive complexity. In parallel computation, we can take advantage of many different processors or computers working simultaneously. The notion of quantification is inherently parallel. I show that the parallel time needed to compute a query corresponds exactly to its quantifier depth. The number of distinct variables occurring in a first-order inductive query corresponds closely with the amount of hardware — processors and memory — needed to compute this query. The most important tradeoff in complexity theory — between parallel time and hardware — is thus identical to the tradeoff between inductive depth and number of variables.

Chapter 6 introduces a combinatorial game that serves as an important tool for ascertaining what can and cannot be expressed in logical languages. Ehrenfeucht-Fraïssé games offer a semantics for first-order logic that is equivalent to, but more directly applicable than, the standard definitions. These games provide powerful tools for descriptive complexity. Using them, we can often decide whether a given query is or is not expressible in a given language.

Chapter 7 introduces second-order logic. This is much more expressive than first-order logic because we may quantify over an exponentially larger space of objects. I prove Fagin's theorem as well as Stockmeyer's characterization of the polynomial-time hierarchy as the set of second-order describable boolean queries. It follows from previous results that the polynomial-time hierarchy is the set of boolean queries computable in constant time, but using exponentially much hard-

ware (Corollary 7.28). This insight exposes the strange character of the polynomial-time hierarchy and of the class NP.

Chapter 8 uses Ehrenfeucht-Fraïssé games to prove that certain queries are not expressible in some restrictions of second-order logic. Since second-order logic is so expressive, it is surprising that we can prove results about non-expressibility. However, the restrictions needed on the second-order languages — in particular, that they quantify only monadic relations — are crucial.

Chapter 9 studies the transitive-closure operator, a restriction of the least-fixed-point operator. I show that transitive closure characterizes the power of the class NL. When infinite structures are allowed, both the least-fixed-point and transitive-closure operators are not closed under negation. In this case there is a strict expressive hierarchy as we alternate applications of these operators with negation. However, for finite structures I show that these operators are closed under negation. A corollary is that nondeterministic space classes are closed under complementation. This was a very unexpected result when it was proved. It constitutes a significant contribution of descriptive complexity to computer science.

Chapter 10 studies the complexity class polynomial space, PSPACE, which is the set of all boolean queries that can be computed using a polynomial amount of hardware, but with no restriction on time. Thus PSPACE is beyond the realm of what is feasibly computable. It is obvious that NP is contained in PSPACE, but it is not known whether PSPACE is larger than NP. Indeed, it is not even known that PSPACE is larger than P. PSPACE is a very robust complexity class. It has several interesting descriptive characterizations, which expose more information about the tradeoff between inductive depth and number of variables.

Chapter 11 studies precomputation — the work that may go into designing the program, formula, or circuit before any input is seen. Precomputation — even less well understood than time and hardware — has an especially crisp formulation in descriptive complexity.

In order for a structure such as a graph to be input to a real or idealized machine, it must be encoded as a character string. Such an encoding imposes an ordering on the universe of the structure, e.g., on the vertices of the graph. All first-order, descriptive characterizations of complexity classes assume that a total ordering relation on the universe is available in the languages. Without such an ordering, simple lower bounds from Chapter 6 show that certain trivial properties — such as computing the PARITY of the cardinality of the universe — are not expressible. However, the ordering relation allows us to distinguish isomorphic structures which all plausible queries should treat the same. In addition, an ordering relation spoils the power of Ehrenfeucht-Fraïssé games for most languages.

The mathematically rich search for a suitable alternative to ordering is described in Chapter 12.

Chapter 13 describes some interesting combinatorial arguments that provide lower bounds on descriptive complexity. The first is the optimal lower bound due to Håstad on the quantifier depth needed to express PARITY. One of many corollaries is that the set of first-order boolean queries is a strict subset of L. The second two lower bounds are weaker: they use Ehrenfeucht-Fraïssé games without ordering and thus, while quite interesting, do not separate complexity classes.

Chapter 14 describes applications of descriptive complexity to databases and computer-aided verification. Relational databases are exactly finite logical structures, and commercial query languages such as SQL are simple extensions of first-order logic. The complexity of query evaluation, the expressive power of query languages, and the optimization of queries are all important practical issues here, and the tools that have been developed previously can be brought to bear on these issues.

Model checking is a burgeoning subfield of computer-aided verification. The idea is that the design of a circuit, protocol, or program can be automatically translated into a transition system, i.e., a graph whose vertices represent global states and whose edges represent possible atomic transitions. Model checking means deciding whether such a design satisfies a simple correctness condition such as, “Doors are not opened between stations”, or, “Division is always performed correctly”. In descriptive complexity, we can see on the face of such a query what the complexity of checking it will be.

Finally, Chapter 15 sketches some directions for future research in and applications of descriptive complexity.

Acknowledgements:

I have been intending to write this book for more years than I would like to admit. In the mean time, many researchers have changed and extended the field so quickly that it is not possible for me to really keep up. I have tried to give pointers to some of the many topics not covered.

I am grateful to everyone who has found errors or made suggestions or encouraged me to write this book. All the errors remaining are mine alone. There will be a page on the world wide web with corrections, recent developments, etc., concerning this book and descriptive complexity in general. Just search for “Neil Immerman” on the web and you will find it. All comments, corrections, etc., will be greatly appreciated.

Some of the people who have already provided help and helpful comments are: Natasha Alechina, Jose Balcázar, Dave Mix Barrington, Jonathan Buss, Russ Ellsworth, Miklos Erdelyi-Szabo, Ron Fagin, Erich Grädel, Jens Gramm, Martin Grohe, Brian Hanechak, Lauri Hella, Janos Makowsky, Yiannis Moschovakis, John Ridgway, Jose Antonio Medina, Gleb Naumovich, Sushant Patnaik, Nate Segerlind, Richard Shore, Wolfgang Thomas, and especially Kousha Etessami.

I am grateful to David Gries for taking his role as an editor of this series so seriously that he read this book in detail, making numerous helpful comments and corrections.

Thanks to the following institutions for financial support during the long process of writing this book: NSF Grant CCR-9505446, Cornell University Computer Science Department, and the DIMACS special year in logic and algorithms.

I want to acknowledge my debt to the many inspiring teachers that I have had over the years. I have a vivid memory of Larry Carter on crutches because he had broken a blood vessel in his leg during the previous afternoon's soccer game hopping back and forth in front of the room as he built a Turing machine to multiply two numbers. This was at an NSF sponsored summer program at the University of New Hampshire in 1969. David Kelly was the codirector of that program and has been making magic ever since, creating summer programs where mathematics as a creative and cooperative endeavor is taught and shared. These programs more than anything else taught me the value and pleasure of teaching and research. Larry later roomed with Ron Fagin as a graduate student at Berkeley, and it was because of this connection that I learned of Ron's research connecting logic and complexity.

A few other of my teachers that I would like to thank by name are Shizuo Kakutani, Angus Macintyre, John Hopcroft, and Juris Hartmanis.

I thank my wife, Susan Landau, who, in part to let me pursue my own obscure interests, has given up more than anyone should have to. I am delighted that you finished your book [DL98] first and to such acclaim. Thanks for your love, for your unswerving integrity, and for your amazing ability to keep moving forward.

Contents

0	Introduction	1
1	Background in Logic	5
1.1	Introduction and Preliminary Definitions	5
1.2	Ordering and Arithmetic	15
1.2.1	FO(BIT) = FO(PLUS, TIMES)	17
1.3	Isomorphism	20
1.4	First-Order Queries	21
2	Background in Complexity	29
2.1	Introduction	29
2.2	Preliminary Definitions	30
2.3	Reductions and Complete Problems	34
2.4	Alternation	43
2.5	Simultaneous Resource Classes	51
2.6	Summary	52
3	First-Order Reductions	57
3.1	$\text{FO} \subseteq \text{L}$	57
3.2	Dual of a First-Order Query	59
3.3	Complete problems for L and NL	64
3.4	Complete Problems for P	67

4	Inductive Definitions	75
4.1	Least Fixed Point	75
4.2	The Depth of Inductive Definitions	81
4.3	Iterating First-Order Formulas	83
5	Parallelism	89
5.1	Concurrent Random Access Machines	90
5.2	Inductive Depth Equals Parallel Time	93
5.3	Number of Variables versus Number of Processors	98
5.4	Circuit Complexity	102
5.5	Alternating Complexity	112
5.5.1	Alternation as Parallelism	114
6	Ehrenfeucht-Fraïssé Games	119
6.1	Definition of the Games	119
6.2	Methodology for First-Order Expressibility	129
6.3	First-Order Properties are Local	134
6.4	Bounded Variable Languages	135
6.5	Zero-One Laws	140
6.6	Ehrenfeucht-Fraïssé Games with Ordering	143
7	Second-Order Logic and Fagin's Theorem	147
7.1	Second-Order Logic	147
7.2	Proof of Fagin's Theorem	150
7.3	NP-Complete Problems	154
7.4	The Polynomial-Time Hierarchy	157
8	Second-Order Lower Bounds	161
8.1	Second-Order Games	161
8.2	$SO\exists$ (monadic) Lower Bound on Reachability	167
8.3	Lower Bounds Including Ordering	172

9	Complementation and Transitive Closure	177
9.1	Normal Form Theorem for FO(LFP)	177
9.2	Transitive Closure Operators	182
9.3	Normal Form for FO(TC)	184
9.4	Logspace is Primitive Recursive	189
9.5	$\text{NSPACE}[s(n)] = \text{co-NSPACE}[s(n)]$	191
9.6	Restrictions of SO	194
10	Polynomial Space	199
10.1	Complete Problems for PSPACE	199
10.2	Partial Fixed Points	203
10.3	$\text{DSPACE}[n^k] = \text{VAR}[k + 1]$	206
10.4	Using Second-Order Logic to Capture PSPACE	210
11	Uniformity and Precomputation	215
11.1	An Unbounded Number of Variables	216
11.1.1	Tradeoffs Between Variables and Quantifier Depth	217
11.2	First-Order Projections	218
11.3	Help Bits	224
11.4	Generalized Quantifiers	225
12	The Role of Ordering	229
12.1	Using Logic to Characterize Graphs	230
12.2	Characterizing Graphs Using \mathcal{L}^k	232
12.3	Adding Counting to First-Order Logic	234
12.4	Pebble Games for \mathcal{C}^k	237
12.5	Vertex Refinement Corresponds to \mathcal{C}^2	239
12.6	Abiteboul-Vianu and Otto Theorems	243
12.7	Toward a Language for Order-Independent P	252

13 Lower Bounds	257
13.1 Håstad's Switching Lemma	257
13.2 A Lower Bound for REACH_a	263
13.3 Lower Bound for Fixed Point and Counting	271
14 Applications	281
14.1 Databases	281
14.1.1 SQL	282
14.1.2 Datalog	285
14.2 Dynamic Complexity	287
14.2.1 Dynamic Complexity Classes	289
14.3 Model Checking	298
14.3.1 Temporal Logic	299
14.4 Summary	305
15 Conclusions and Future Directions	307
15.1 Languages That Capture Complexity Classes	307
15.1.1 Complexity on the Face of a Query	310
15.1.2 Stepwise Refinement	310
15.2 Why Is <i>Finite</i> Model Theory Appropriate?	311
15.3 Deep Mathematical Problems: P versus NP	312
15.4 Toward Proving Lower Bounds	313
15.4.1 Role of Ordering	314
15.4.2 Approximation and Approximability	314
15.5 Applications of Descriptive Complexity	315
15.5.1 Dynamic Complexity	315
15.5.2 Model Checking	316
15.5.3 Abstract State Machines	316
15.6 Software Crisis and Opportunity	317
15.6.1 How can Finite Model Theory Help?	318

Chapter 0

Introduction

In the beginning, there were two measures of computational complexity: time and space. From an engineering standpoint, these were very natural measures, quantifying the amount of physical resources needed to perform a computation. From a mathematical viewpoint, time and space were somewhat less satisfying, since neither appeared to be tied to the inherent mathematical complexity of the computational problem.

In 1974, Ron Fagin changed this. He showed that the complexity class NP — those problems computable in nondeterministic polynomial time — is exactly the set of problems describable in second-order existential logic. This was a remarkable insight, for it demonstrated that the computational complexity of a problem can be understood as the richness of a language needed to specify the problem. Time and space are not model-dependent engineering concepts, they are more fundamental.

Although few programmers consider their work in this way, a computer program is a completely precise description of a mapping from inputs to outputs. In this book we follow database terminology and call such a map a *query* from input structures to output structures. Typically a program describes a precise sequence of steps that compute a given query. However, we may choose to describe the query in some other precise way. For example, we may describe queries in variants of first- and second-order mathematical logic.

Fagin's Theorem gave the first such connection. Using first-order languages, this approach, commonly called descriptive complexity, demonstrated that virtually all measures of complexity can be mirrored in logic. Furthermore, as we will see, the most important classes have especially elegant and clean descriptive characterizations.

Descriptive complexity provided the insight behind a proof of the Immerman-Szelepcsényi Theorem, which states that nondeterministic space classes are closed under complementation. This settled a question that had been open for twenty-five years; indeed, almost everyone had conjectured the negation of this theorem.

Descriptive complexity has long had applications to database theory. A relational database is a finite logical structure, and commonly used query languages are small extensions of first-order logic. Thus, descriptive complexity provides a natural foundation for database theory, and many questions concerning the expressibility of query languages and the efficiency of their evaluation have been settled using the methods of descriptive complexity. Another prime application area of descriptive complexity is to the problems of Computer Aided Verification.

Since the inception of complexity theory, a fundamental question that has bedeviled theorists is the P versus NP question. Despite almost three decades of work, the problem of proving P different from NP remains. As we will see, P versus NP is just a famous and dramatic example of the many open problems that remain. Our inability to ascertain relationships between complexity classes is pervasive. We can prove that more of a given resource, e.g., time, space, nondeterministic time, etc., allows us to compute strictly more queries. However, the relationship between different resources remains virtually unknown.

We believe that descriptive complexity will be useful in these and many related problems of computational complexity. Descriptive complexity is a rich edifice from which to attack the tantalizing problems of complexity. It gives a mathematical structure with which to view and set to work on what had previously been engineering questions. It establishes a strong connection between mathematics and computer science, thus enabling researchers of both backgrounds to use their various skills to set upon the open questions. It has already led to significant successes.

The Case for Finite Models

A fundamental philosophical decision taken by the practitioners of descriptive complexity is that computation is inherently finite. The relevant objects — inputs, databases, programs, specifications — are all finite objects that can be conveniently modeled as finite logical structures. Most mathematical theories study infinite objects. These are considered more relevant, general, and important to the typical mathematician. Furthermore, infinite objects are often simpler and better behaved than their finite cousins. A typical example is the set of natural numbers, $\mathbf{N} = \{0, 1, 2, \dots\}$. Clearly this has a simpler and more elegant theory than the set of natural numbers representable in 64-bit computer words. However, there is a

significant danger in taking the infinite approach. Namely, the models are often wrong! Properties that we can prove about \mathbf{N} are often false or irrelevant if we try to apply them to the objects that computers have and hold. We find that the subject of finite models is quite different in many respects. Different theorems hold and different techniques apply.

Living in the world of finite structures may seem odd at first. Descriptive complexity requires a new way of thinking for those readers who have been brought up on infinite fare. Finite model theory is different and more combinatorial than general model theory. In Descriptive complexity, we use finite model theory to understand computation. We expect that the reader, after some initial effort and doubt, will agree that the theory of computation that we develop has significant advantages. We believe that it is more accurate and more relevant in the study of computation.

I hope the reader has as much pleasure in discovering and using the tools of Descriptive complexity as I have had. I look forward to new contributions in the modeling and understanding of computation to be made by some of the readers of this book.

Chapter 1

Background in Logic

Mathematics enables us to model many things abstractly. Group theory, for example, abstracts features of such diverse activities as English change ringing and quantum mechanics. Mathematical logic carries the abstraction one level higher: it is a mathematical model of mathematics. This book shows that the computational complexity of all problems in computer science can be understood via the complexity of their logical descriptions. We begin with a high-level introduction to logic. Although much of the material is well-known, we urge readers to at least skim this background chapter as the concentration on finite and ordered structures, i.e., relational databases, is not standard in most treatments of logic.

1.1 Introduction and Preliminary Definitions

All logic books begin with definitions. We have to introduce the language before we start to speak. Thus, a *vocabulary*

$$\tau = \langle R_1^{a_1}, \dots, R_r^{a_r}, c_1, \dots, c_s, f_1^{r_1}, \dots, f_t^{r_t} \rangle$$

is a tuple of relation symbols, constant symbols, and function symbols. R_i is a relation symbol of arity a_i and f_j is a function symbol of arity r_j . Two important examples are $\tau_g = \langle E^2, s, t \rangle$, the vocabulary of graphs with specified source and terminal nodes, and $\tau_s = \langle \leq^2, S^1 \rangle$, the vocabulary of binary strings.

A *structure* with vocabulary τ is a tuple,

$$\mathcal{A} = \langle |\mathcal{A}|, R_1^{\mathcal{A}}, \dots, R_r^{\mathcal{A}}, c_1^{\mathcal{A}}, \dots, c_s^{\mathcal{A}}, f_1^{\mathcal{A}}, \dots, f_t^{\mathcal{A}} \rangle$$




Figure 1.1: Graphs G and H

whose universe is the nonempty set $|\mathcal{A}|$. For each relation symbol R_i of arity a_i in τ , \mathcal{A} has a relation $R_i^{\mathcal{A}}$ of arity a_i defined on $|\mathcal{A}|$, i.e., $R_i^{\mathcal{A}} \subseteq |\mathcal{A}|^{a_i}$. For each constant symbol $c_j \in \tau$, \mathcal{A} has a specified element of its universe $c_j^{\mathcal{A}} \in |\mathcal{A}|$. For each function symbol $f_i \in \tau$, $f_i^{\mathcal{A}}$ is a total function from $|\mathcal{A}|^{r_i}$ to $|\mathcal{A}|$. A vocabulary without function symbols is called a *relational vocabulary*. In this book, unless stated otherwise, all vocabularies are relational. The notation $\|\mathcal{A}\|$ denotes the cardinality of the universe of \mathcal{A} .

In the history of mathematical logic most interest has concentrated on infinite structures. Indeed, many mathematicians consider the study of finite structures trivial. Yet, the objects computers have and hold are always finite. To study computation we need a theory of finite structures.

Logic restricted to finite structures is rather different from the theory of infinite structures. We mention infinite structures from time to time, most often when we comment on whether a given theorem also holds in the infinite case. However, we concentrate on finite structures. We define $\text{STRUC}[\tau]$ to be the set of finite structures of vocabulary τ .

As an example, the graph $G = \langle V^G, E^G, 1, 3 \rangle$ defined by,

$$V^G = \{0, 1, 2, 3, 4\}, \quad E^G = \{(1, 2), (3, 0), (3, 1), (3, 2), (3, 4), (4, 0)\}$$

is a structure of vocabulary τ_g consisting of a directed graph with two specified vertices s and t . G has five vertices and six edges. (See Figure 1.1, which shows G as well as another graph H which is isomorphic but not equal to G .)

For another example, consider the binary string $w = \text{“01101”}$. We can code w as the structure $\mathcal{A}_w = \langle \{0, 1, \dots, 4\}, \leq, \{1, 2, 4\} \rangle$ of vocabulary τ_s . Here \leq represents the usual ordering on $0, 1, \dots, 4$. Relation $S^w = \{1, 2, 4\}$ represents

the positions where w is one. (Relation symbols of arity one, such as S^w , are sometimes called *monadic*.)

A relational database is exactly a finite relational structure. The following begins a running example of a genealogical database.

Example 1.2 Consider a genealogical database $\mathcal{B}_0 = \langle U_0, F_0, P_0, S_0 \rangle$; where U_0 is a finite set of people,

$$U_0 = \{\text{Abraham, Isaac, Rebekah, Sarah, } \dots\}$$

F_0 is a monadic relation that is true of the female elements of U_0 ,

$$F_0 = \{\text{Sarah, Rebekah, } \dots\}$$

P_0 and S_0 are the binary relations for parent and spouse, respectively, e.g.,

$$\begin{aligned} P_0 &= \{\langle \text{Abraham, Isaac} \rangle, \langle \text{Sarah, Isaac} \rangle, \dots\} \\ S_0 &= \{\langle \text{Abraham, Sarah} \rangle, \langle \text{Isaac, Rebekah} \rangle, \dots\} \end{aligned}$$

Thus, \mathcal{B}_0 is a structure of vocabulary $\langle F^1, P^2, S^2 \rangle$. □

For any vocabulary τ , define the *first-order language* $\mathcal{L}(\tau)$ to be the set of formulas built up from the relation and constant symbols of τ ; the logical relation symbol $=$; the boolean connectives \wedge, \neg ; variables: $\text{VAR} = \{x, y, z, \dots\}$; and quantifier \exists .

We say that an occurrence of a variable v in φ is *bound* if it lies within the scope of a quantifier ($\exists v$) or ($\forall v$), otherwise it is *free*. Variable v is *free* in φ iff it has a free occurrence in φ . For example, the free variables in the following formula are x and y . We use the symbol “ \equiv ” to define or denote equivalence of formulas.

$$\alpha \equiv [(\exists y)(y + 1 = x)] \wedge x < y$$

In a similar way we sometimes use “ \Leftrightarrow ” to indicate that two previously defined formulas or conditions are equivalent.

Bound variables are “dummy” variables and may be renamed to avoid confusion. For example, α is equivalent to the following α' which also has free variables x and y ,

$$\alpha' \equiv [(\exists z)(z + 1 = x)] \wedge x < y$$

We write $\mathcal{A} \models \varphi$ to mean that \mathcal{A} satisfies φ , i.e., that φ is true in \mathcal{A} . Since φ may contain some free variables, we will let an *interpretation* into \mathcal{A} be a map $i : V \rightarrow |\mathcal{A}|$ where V is some finite subset of VAR. For convenience, for every constant symbol $c \in \tau$ and any interpretation i for \mathcal{A} , we let $i(c) = c^{\mathcal{A}}$. If τ has function symbols, then the definition of i extends to all terms via the recurrence,

$$i(f_j(t_1, \dots, t_{r_j})) = f_j^{\mathcal{A}}(i(t_1), \dots, i(t_{r_j})).$$

We can be completely precise about the semantics of mathematical logic. In particular, we can definitively define what it means for a sentence φ to be true in a structure \mathcal{A} .

Definition 1.3 (Definition of Truth) Let $\mathcal{A} \in \text{STRUC}[\tau]$ be a structure, and let i be an interpretation into \mathcal{A} whose domain includes all the relevant free variables. We inductively define whether a formula $\varphi \in \mathcal{L}(\tau)$ is true in (\mathcal{A}, i) :

$$\begin{aligned} (\mathcal{A}, i) \models t_1 = t_2 &\Leftrightarrow i(t_1) = i(t_2) \\ (\mathcal{A}, i) \models R_j(t_1, \dots, t_{a_j}) &\Leftrightarrow \langle i(t_1), \dots, i(t_{a_j}) \rangle \in R_j^{\mathcal{A}} \\ (\mathcal{A}, i) \models \neg\varphi &\Leftrightarrow \text{it is not the case that } (\mathcal{A}, i) \models \varphi \\ (\mathcal{A}, i) \models \varphi \wedge \psi &\Leftrightarrow (\mathcal{A}, i) \models \varphi \text{ and } (\mathcal{A}, i) \models \psi \\ (\mathcal{A}, i) \models (\exists x)\varphi &\Leftrightarrow (\text{there exists } a \in |\mathcal{A}|)(\mathcal{A}, i, a/x) \models \varphi \end{aligned}$$

where $(i, a/x)(y) = \begin{cases} i(y) & \text{if } y \neq x \\ a & \text{if } y = x \end{cases}$

Write $\mathcal{A} \models \varphi$ to mean that $(\mathcal{A}, \emptyset) \models \varphi$. □

Definition 1.3 is our first example of an inductive definition, a device that is often used by logicians. It deserves a few comments. Note that the equality symbol ($=$) is not treated as an ordinary binary relation symbol — the definition insists that this symbol be interpreted as equality. Many students, on first seeing this definition, feel that it is circular. It is not. We are defining the meaning of the symbol “ $=$ ” in terms of the intuitively well-understood standard equality. In the same way, we define the meaning of “ \neg ”, “ \wedge ”, and “ \exists ” in terms of their intuitive counterparts.

We define the “for all” quantifier as the dual of \exists and the boolean “or” as the dual of \wedge ,

$$(\forall x)\varphi \equiv \neg(\exists x)\neg\varphi; \quad \alpha \vee \beta \equiv \neg(\neg\alpha \wedge \neg\beta)$$

It is convenient to introduce other abbreviations into our formulas. For example, “ $y \neq z$ ” is an abbreviation for “ $\neg y = z$ ”. Similarly “ $\alpha \rightarrow \beta$ ” is an abbreviation for “ $\neg\alpha \vee \beta$ ”, and “ $\alpha \leftrightarrow \beta$ ” is an abbreviation for “ $\alpha \rightarrow \beta \wedge \beta \rightarrow \alpha$ ”. In some sense, the symbols we introduce formally into our language are part of our low-level “machine language”, and abbreviations are analogous to what computer scientists call macros. Abbreviations are directly translatable into the real language, and they make formulas more readable. Without abbreviations and the breaking of formulas into modular descriptions, it would be impossible to communicate complicated ideas in first-order logic.

We use spacing and parentheses to make the order of operations clear. Our convention for operator precedence is that “ \neg ”, “ \forall ”, and “ \exists ” have highest precedence, then “ \wedge ” and “ \vee ”, and finally, “ \rightarrow ” and “ \leftrightarrow ”. The operators “ \wedge ” and “ \vee ” are evaluated left to right, but “ \rightarrow ” and “ \leftrightarrow ” are evaluated right to left. For example, the following two formulas are equivalent,

$$\begin{aligned} \neg R(a) \rightarrow R(b) \wedge R(c) \vee R(d) \leftrightarrow R(e) \\ (\neg R(a)) \rightarrow (((R(b) \wedge R(c)) \vee R(d)) \leftrightarrow R(e)) \end{aligned}$$

A *sentence* is a formula with no free variables. Every sentence $\varphi \in \mathcal{L}(\tau)$ is either true or false in any structure $\mathcal{A} \in \text{STRUC}[\tau]$.

Example 1.4 We give a few examples of first-order formulas in the language of graphs:

$$\varphi_{undir} \equiv (\forall x)(\forall y)(\neg E(x, x) \wedge (E(x, y) \rightarrow E(y, x)))$$

Formula φ_{undir} says that the graph in question is undirected and has no loops.

$$\begin{aligned} \varphi_{out2} \equiv (\forall x)(\exists yz)(y \neq z \wedge E(x, y) \wedge E(x, z) \wedge \\ (\forall w)(E(x, w) \rightarrow (w = y \vee w = z))) \end{aligned}$$

$$\varphi_{deg2} \equiv \varphi_{undir} \wedge \varphi_{out2}$$

Formula φ_{out2} says that every vertex has exactly two edges leaving it. Thus, φ_{deg2} says that the graph in question is undirected, has no loops, and is regular of degree two, i.e., every vertex has exactly two neighbors.

$$\begin{aligned}
\varphi_{dist1} &\equiv x = y \vee E(x, y) \\
\varphi_{dist2} &\equiv (\exists z)(\varphi_{dist1}(x, z) \wedge \varphi_{dist1}(z, y)) \\
\varphi_{dist4} &\equiv (\exists z)(\varphi_{dist2}(x, z) \wedge \varphi_{dist2}(z, y)) \\
\varphi_{dist8} &\equiv (\exists z)(\varphi_{dist4}(x, z) \wedge \varphi_{dist4}(z, y))
\end{aligned}$$

Formulas φ_{dist1} , φ_{dist2} , and so on say that there is a path from x to y of length at most 1, 2, 4, and 8, respectively. Note that these formulas have free variables x and y .

Formulas express properties about their free variables. For example, for a pair of vertices a, b from the universe of a graph G , the meaning of

$$(G, a/x, b/y) \models \varphi_{dist8}$$

is that the distance from a to b in G is at most 8.

Sometimes we will make the free variables in a formula explicit, e.g., writing $\varphi_{dist8}(x, y)$ instead of just φ_{dist8} . This offers the advantage of making substitutions more readable: we can write $\varphi_{dist8}(a, b)$ instead of $\varphi_{dist8}(a/x, b/y)$. \square

Exercise 1.5 For $n \in \mathbf{N}$, consider the logical structures

$$\mathcal{A}_n = \langle \{0, 1, \dots, n-1\}, \text{PLUS}^{\mathcal{A}_n}, \text{TIMES}^{\mathcal{A}_n}, 0, 1, n-1 \rangle$$

of vocabulary $\tau_a = \langle \text{PLUS}^3, \text{TIMES}^3, 0, 1, \text{max} \rangle$, where PLUS and TIMES are the arithmetic relations, i.e., for $i, j, k < n$,

$$\begin{aligned}
\mathcal{A}_n \models \text{PLUS}(i, j, k) &\Leftrightarrow i + j = k \\
\mathcal{A}_n \models \text{TIMES}(i, j, k) &\Leftrightarrow i \cdot j = k
\end{aligned}$$

Write formulas in $\mathcal{L}(\tau)$ that represent the following arithmetic relations,

1. DIVIDES(x, y), meaning that y is a multiple of x .
2. PRIME(x), meaning that x is a prime number.
3. $p_2(x)$, meaning that x is a power of 2.

[Hint for (3): x is a power of 2 iff 2 is the only prime divisor of x .] \square

Example 1.6 Here are a few formulas in the language of strings. The first describes the set of strings that have no consecutive “1”s. It uses the abbreviation “ $x < y$ ”, meaning “ $x \leq y \wedge x \neq y$ ”.

$$\varphi_{no11} \equiv (\forall x)(\forall y)(\exists z)((S(x) \wedge S(y) \wedge x < y) \rightarrow (x < z < y \wedge \neg S(z)))$$

Formula φ_{five1} below says that the given string contains at least five “1”s. To do so, it uses the abbreviation “distinct”:

$$\text{distinct}(x_1, \dots, x_k) \equiv (x_1 \neq x_2 \wedge \dots \wedge x_1 \neq x_k \wedge \dots \wedge x_{k-1} \neq x_k)$$

$$\varphi_{five1} \equiv (\exists uvwxy)(\text{distinct}(u, v, w, x, y) \wedge S(u) \wedge S(v) \wedge S(w) \wedge S(x) \wedge S(y))$$

Note that φ_{five1} uses five variables to say that there are five “1”s. Using the ordering relation, we can reduce the number of variables. The following formula is equivalent to φ_{five1} but uses only two variables:

$$(\exists x) \left(S(x) \wedge (\exists y) \left(x < y \wedge S(y) \wedge (\exists x) (y < x \wedge S(x) \wedge (\exists y) (x < y \wedge S(y) \wedge (\exists x) y < x \wedge S(x))) \right) \right)$$

Read the above sentence carefully. A good way to think of it is that we have two fingers and are trying to count the number of “1”s in a string. We put finger x down on the first “1”. Then we put finger y down on the next “1” to the right. Now we don’t need x anymore so we can move it to the next “1” to the right of y , and so on.

We will see later that the number of variables is an important descriptive resource. Note that the standard semantics of first-order logic (Definition 1.3) allows us to requantify variables. Each quantifier $(\exists x)$ or $(\forall x)$ bounds only the free occurrences of x within its scope. We will see in Theorem 6.31 that every first-order sentence in $\mathcal{L}(\tau_s)$ — i.e., every sentence about strings — is equivalent to a sentence with only three distinct variables. \square

Exercise 1.7 Prove that if interpretations i and i' agree on all the free variables in φ then

$$(\mathcal{A}, i) \models \varphi \iff (\mathcal{A}, i') \models \varphi$$

[Hint: by induction on φ using Definition 1.3.] □

Exercise 1.8 Let $(\exists!x)\alpha(x)$ mean that there exists a unique x such that α . Show how to write $(\exists!x)\alpha(x)$ using the usual quantifiers \forall, \exists . □

As another example, let $\tau_{ab} = \langle \leq^2, A^1, B^1 \rangle$ consist of an ordering relation and two monadic relation symbols A and B , each serving the same role as the symbol S in τ_s . Let $\mathcal{A} \in \text{STRUC}[\tau_{ab}]$, and let $n = \|\mathcal{A}\|$. Then \mathcal{A} is a pair of binary strings A, B , each of length n . These binary strings represent natural numbers, where we think of the bit zero as most significant and bit $n - 1$ as least significant. Here $A(i)$ is true iff bit i of A is “1”.

The following sentence expresses the ordering relation on such natural numbers represented in binary.

$$\text{LESS}(A, B) \equiv (\exists x)(B(x) \wedge \neg A(x) \wedge (\forall y. y < x)(A(y) \rightarrow B(y)))$$

The above sentence uses a very useful abbreviation, that of restricted quantifiers,

$$(\forall x.\alpha)\varphi \equiv (\forall x)(\alpha \rightarrow \varphi); \quad (\exists x.\alpha)\varphi \equiv (\exists x)(\alpha \wedge \varphi)$$

In the next proposition we show that addition is first-order expressible. Addition of natural numbers represented in binary is one of the most basic computations. We will see in Theorem 5.2 that the first-order queries characterize the problems computable in constant parallel time. Thus the following may be thought of as an addition algorithm that runs in constant parallel time.

Proposition 1.9 *Addition of natural numbers, represented in binary, is first-order expressible.*

Proof We use the well-known “carry-look-ahead” algorithm. In order to express addition, we first express the carry bit,

$$\varphi_{\text{carry}}(x) \equiv (\exists y. x < y)[A(y) \wedge B(y) \wedge (\forall z. x < z < y)A(z) \vee B(z)]$$

The formula $\varphi_{carry}(x)$ holds if there is a position y to the right of x where $A(y)$ and $B(y)$ are both one (i.e. the carry is generated) and for all intervening positions z , at least one of $A(z)$ and $B(z)$ holds (that is, the carry is propagated). Let \oplus be an abbreviation for the commutative and associative “exclusive or” operation.

We can express φ_{add} as follows,

$$\begin{aligned}\alpha \oplus \beta &\equiv \alpha \leftrightarrow \neg\beta \\ \varphi_{add}(x) &\equiv A(x) \oplus B(x) \oplus \varphi_{carry}(x)\end{aligned}$$

Note that the formula $\varphi_{add}(x)$ has the free variable x . Thus, φ_{add} is a description of n bits: one for each possible value of x . \square

An important relation between two structures of the same type is that one may be a substructure of the other. \mathcal{A} is a substructure of \mathcal{B} if the universe of \mathcal{A} is a subset of the universe of \mathcal{B} and the relations and constants on \mathcal{A} are inherited from \mathcal{B} .

Definition 1.10 (Substructure) Let \mathcal{A} and \mathcal{B} be structures of the same vocabulary $\tau = \langle R_1^{a_1}, \dots, R_r^{a_r}, c_1, \dots, c_s \rangle$. We say that \mathcal{A} is a *substructure* of \mathcal{B} , written $\mathcal{A} \leq \mathcal{B}$, iff the following conditions hold,

1. $|\mathcal{A}| \subseteq |\mathcal{B}|$
2. For $i = 1, 2, \dots, r$, $R_i^{\mathcal{A}} = R_i^{\mathcal{B}} \cap |\mathcal{A}|^{a_i}$
3. For $j = 1, 2, \dots, s$, $c_j^{\mathcal{A}} = c_j^{\mathcal{B}}$. \square

See Figure 1.11 where A and B are substructures of G . Note that C is not a substructure of G for two reasons: it doesn't contain the constant t and the induced edge from vertex 1 to vertex 2 is missing.

Exercise 1.12 Let $\mathcal{A} \in \text{STRUC}[\tau]$ be a structure and let $\alpha(x)$ be a formula such that $\mathcal{A} \models (\exists x)\alpha(x)$. Assume also that for every constant symbol c in τ , $\mathcal{A} \models \alpha(c)$. Let \mathcal{B} be the substructure of \mathcal{A} with universe

$$|\mathcal{B}| = \{a \in |\mathcal{A}| \mid \mathcal{A} \models \alpha(a)\}$$

Let φ be a sentence in $\mathcal{L}(\tau)$. Define the *restriction* of φ to α to be the sentence φ^α , the result of changing every quantifier $(\forall y)$ or $(\exists y)$ in φ to the restricted quantifier $(\forall y.\alpha(y))$ or $(\exists y.\alpha(y))$ respectively. Prove the following,




Figure 1.11: A and B but not C are substructures of G .

$$\mathcal{A} \models \varphi^\alpha \quad \Leftrightarrow \quad \mathcal{B} \models \varphi \quad \square$$

We say that φ is *universal* iff it can be written in prenex form — i.e. with all quantifiers at the beginning — using only universal quantifiers. Similarly, we say that φ is *existential* iff it can be written in prenex form with only existential quantifiers.

The following “preservation theorems” provide a good way of proving that a formula is existential or universal.

Exercise 1.13 Prove the following preservation theorems. Let $\mathcal{A} \leq \mathcal{B}$ be structures and φ a first-order sentence.

1. Suppose φ is existential. If $\mathcal{A} \models \varphi$ then $\mathcal{B} \models \varphi$.
2. Suppose φ is universal. If $\mathcal{B} \models \varphi$ then $\mathcal{A} \models \varphi$.

[Hint: by induction on φ using Definition 1.3.] □

1.2 Ordering and Arithmetic

A logical structure such as a graph does not need to have an ordering on its vertices. However, if we use a computer to store or manipulate this graph, it must be encoded in some way that imposes an ordering on the vertices. In order to discuss computation in general, it is necessary to assume that the universes of our structures are ordered. This section introduces the issue of ordering and explains what we will assume about the ordering of structures in the remainder of this book.

When we code an input to a computer, we do so as a string of characters. There is always an ordering here: the first character, the second character, and so on. Indeed the concept of ordering is deeply embedded in the concepts of string and of computation.

For this reason, the binary relation symbol “ \leq ” plays a special role in descriptive complexity. When “ \leq ” is an element of τ , and $\mathcal{A} \in \text{STRUC}[\tau]$, then \mathcal{A} must interpret \leq as a total ordering on its universe. In this case, we also place constant symbols $0, 1, \max$ in τ and insist that these be interpreted as the minimum, second, and maximum elements under the \leq ordering. In order for formulas in first-order logic to express general computation, they need access to a total ordering of the universe. The requirement of a total ordering in descriptive complexity is analogous to the assumption of the Axiom of Choice in set theory.

Let $\mathcal{A} \in \text{STRUC}[\tau]$ be an ordered structure. Let $n = \|\mathcal{A}\|$. Let the elements of $|\mathcal{A}|$ in increasing order be a_0, a_1, \dots, a_{n-1} . Then there is a 1:1 correspondence $i \mapsto a_i, i = 0, 1, \dots, n-1$. We usually identify the elements of the universe with the set of natural numbers less than n . In a computer these would be represented as $\lceil \log n \rceil$ -bit words, and the operations plus, times, and even picking out bit j of such a word would all be wired in. The following numeric relations are useful:

1. PLUS(i, j, k) meaning $i + j = k$
2. TIMES(i, j, k) meaning $i \times j = k$
3. BIT(i, j) meaning bit j in the binary representation of i is 1

In the definition of BIT we will take bit 0 to be the low order bit, so BIT($i, 0$) holds iff i is odd. We will see in Chapter 11 that adding BIT (or equivalently PLUS and TIMES) to our vocabularies makes the set of first-order definable boolean queries a more robust complexity class.

When working with very weak reductions or proving normal form theorems, we will sometimes use the successor relation SUC in lieu of or in addition to \leq . Of course, SUC is first-order definable from \leq .

$$\text{SUC}(x, y) \equiv (x < y) \wedge (\forall z)(\neg(x < z \wedge z < y))$$

The symbols \leq , PLUS, TIMES, BIT, SUC, 0, 1, *max* are called *numeric* relation and constant symbols. They depend only on the size of the universe. We call the remainder of τ the *input* relation and constant symbols. We will see in Chapter 5 that the choice of numeric relations for weak languages such as FO corresponds to the definition of uniformity for complexity classes defined by uniform sequences of circuits. The numeric relations and constants are not explicitly given in the input since they are easily computable as functions of the size of the input. Whenever any of the numeric relation or constant symbols occur, they are required to have their standard meanings.

Proviso 1.14 (Ordering Proviso) *From now on, unless stated otherwise, we assume that the numeric relations and constants: \leq , PLUS, TIMES, BIT, SUC, 0, 1, *max* are present in all vocabularies. When we define vocabularies, we do not explicitly mention or show these symbols unless they are not present. In Chapter 6 we prove lower bounds on what can be expressed in some first-order language. We use the notation $\mathcal{L}(\text{wo}\leq)$ to indicate language \mathcal{L} without any of the numeric relations. We will write $\mathcal{L}(\text{wo BIT})$ to indicate language \mathcal{L} , including ordering, but not arithmetic, i.e., only the numeric relations \leq and SUC and the constants 0, 1, *max* are included.*

The following proviso is useful. It eliminates the trivial and sometimes annoying case of the structure with only one element which would thus satisfy the equation $0 = 1$. We assume this proviso unless otherwise noted. (The only time we do not assume the existence of boolean constants is in Section 6.5.)

Proviso 1.15 (Boolean Constants) *From now on, we assume that all structures have at least two elements. In particular, we will assume that we have two unequal constants denoted by 0 and 1.*

Next, we define what it means to have a boolean variable in a first-order formula. Boolean variables allow a more robust measure of the number of first-order variables needed to express a query. When we measure the number of first-order variables needed, we discount the (bounded) number of boolean variables.

Definition 1.16 *A boolean variable in a first-order formula is a variable that is restricted to being either 0 or 1. Here 0 is identified with **false** and 1 is identified*

with **true**. We typically use the letters b, c, d, e for boolean variables. We use the following abbreviations:

$$\begin{aligned} \text{bool}(b) &\equiv b \leq 1 \\ (\exists b) &\equiv (\exists b.\text{bool}(b)) \\ (\forall b) &\equiv (\forall b.\text{bool}(b)) \end{aligned}$$

□

1.2.1 FO(BIT) = FO(PLUS, TIMES)

In the remainder of this section we prove that adding BIT to first-order logic is equivalent to adding PLUS and TIMES. In order to prove this, we also need to prove the Bit Sum Lemma which is interesting in its own right. The proofs in this subsection are very technical and may safely be skipped at first reading.

Theorem 1.17 *Let τ be a vocabulary that includes ordering. Then*

1. *If $\text{BIT} \in \tau$ then PLUS and TIMES are first-order definable.*
2. *If PLUS, TIMES $\in \tau$ then BIT is first-order definable.*

Proof To prove (1), we have essentially seen in Proposition 1.9 that PLUS is expressible using BIT. To prove that TIMES is expressible we first need the following:

Lemma 1.18 (Bit Sum Lemma) *Let $\text{BSUM}(x, y)$ be true iff y is equal to the number of ones in the binary representation of x . BSUM is first-order expressible using ordering and BIT.*

Proof The bit-sum problem is to add a column of $\log n$ 0's and 1's. The idea is to keep a running sum. Since the sum of $\log n$ 1's requires at most $\log \log n$ bits to record, we maintain running sums of $\log \log n$ bits each. With one existentially quantified variable, we can guess $\log n / \log \log n$ of these. Thus, to express $\text{BSUM}(x, y)$ we existentially quantify s — the $\log \log n \cdot (\log n / \log \log n)$ bits of running sums. In the following example, $n = 2^{16}$, so x and y each

have 16 bits. To assert $\text{BSUM}(0110110110101101, 1010)$ we would guess $s = 0010010101111010$ as our partial sum bit string.

$$\begin{array}{r}
 0 \\
 1 \\
 1 \\
 0 \ 0010 \\
 1 \\
 1 \\
 0 \\
 1 \ 0101 \\
 1 \\
 0 \\
 1 \\
 0 \ 0111 \\
 1 \\
 1 \\
 0 \\
 1 \ 1010
 \end{array}
 \qquad
 \text{BSUM}(0110110110101101, 1010)$$

Next we say that for all i where $i \leq \log n / \log \log n$, running sum i , plus the number of 1's in segment $(i + 1)$ is equal to the running sum $(i + 1)$.

Thus, it suffices to express the bit sum of a segment of length $\log \log n$. This we can do by keeping a running sum at every position because this requires only $\log \log \log n \cdot \log \log n$, which is less than $\log n$ for sufficiently large n . \square

We next show that TIMES is first-order expressible using BIT . TIMES is equivalent to the addition of $\log n \log n$ -bit numbers,

$$A = A_1 + A_2 + \cdots + A_{\log n}$$

The first trick we employ is to split each A_i into a sum of two numbers, $A_i = B_i + C_i$, so that B_i and C_i have blocks of $\log \log n$ bits separated by $\log \log n$ 0's. We compute the sum of the B_i 's and of the C_i 's. In this way, we insure that no carries extend more than $\log \log n$ bits. Finally, we add the two sums with a single use of PLUS . In the following, let $\ell = \lceil \log \log n \rceil$.

$$\begin{array}{r}
 B_i = a_{i,1} \ \cdots \ a_{i,\ell} \ \ 0 \ \cdots \ 0 \ \cdots \ a_{i,\log n+1-\ell} \ \cdots \ a_{i,\log n} \\
 + C_i = \ 0 \ \cdots \ 0 \ a_{i,\ell+1} \ \cdots \ a_{i,2\ell} \ \cdots \ \ 0 \ \cdots \ 0 \\
 \hline
 A_i = a_{i,1} \ \cdots \ a_{i,\ell} \ a_{i,\ell+1} \ \cdots \ a_{i,2\ell} \ \cdots \ a_{i,\log n+1-\ell} \ \cdots \ a_{i,\log n}
 \end{array}$$

Position	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Y	0	1	0	0	0	0	0	0	0	1	0	0	0	1	0	1	0
Z	0	0	0	0	1	0	0	0	0	0	0	1	0	0	1	0	0
I	0	1	1	1	1	0	0	0	0	1	1	1	0	1	1	0	0

Table 1.19: Encoding of an arithmetic fact: $2^{15} = 32,768$.

In this way, we have reduced the problem of adding $\log n$ $\log n$ -bit numbers to that of adding $\log n$ $\log \log n$ -bit numbers. We can simultaneously guess the sums of each of the $\log \log n$ columns in a single variable, c . Using BSUM and a universal quantifier we can verify that each section of c is correct. Finally, we can add the $\log \log n$ numbers in c maintaining all the running sums as in the last paragraph of the proof of the Bit Sum Lemma.

2. In this direction, we want to show that BIT is first-order expressible using PLUS and TIMES. We do this with a series of definitions. First, let $p_2(y)$ mean that y is a power of 2. (See Exercise 1.5).

Next, define $\text{BIT}'(x, y)$ to mean that for some i , $y = 2^i$ and $\text{BIT}(x, i)$,

$$\text{BIT}'(x, y) \equiv p_2(y) \wedge (\exists uv)(x = 2uy + y + v \wedge v < y)$$

Using BIT' we can copy a sequence of bits. For example, the following formula says that if $y = 2^i$ and $z = 2^j$, then bits $i + j..i$ of x are the same as bits $j..0$ of c :

$$\text{COPY}(x, y, z, c) \equiv (\forall u.p_2(u) \wedge u \leq z)(\text{BIT}'(x, yu) \leftrightarrow \text{BIT}'(c, u))$$

Finally, to express BIT, we would like to express the relation $2^i = y$. We express this using the following recurrence,

$$2^i = y \Leftrightarrow (\exists j)(\exists z.2^j = z)(i = 2j + 1 \wedge y = 2z^2 \vee i = 2j \wedge y = z^2)(1.20)$$

We can guess three variables, Y, Z, I , that simultaneously include all but a bounded number of the $\log i$ computations indicated by Equation (1.20), namely all those such that $i > 2 \log i$. This is done as follows: Place a “1” in positions i, j , etc., of Y . Place the binary encoding of i starting at position i of I , the binary encoding of j starting at position j of I and so on. Finally, place a “1” in Z at the end of each of the binary encodings of exponents.

Using a universal quantifier we say that the variables Y, Z , and I encode all the relevant and sufficiently large computations of Equation (1.20). Table 1.19

shows the encodings Y, Z , and I for the proposition that $2^{15} = 32,768$. Note that I records the exponent 15, which is 1111 in binary, starting at position 15; 7 which is 111 in binary, starting at position 7; and 3 which is 11 in binary, starting at position 3. We leave the details of actually writing the relevant first-order formula as an exercise. \square

1.3 Isomorphism

When we impose an ordering on the universe of a structure, we have essentially labeled its elements 0, 1, and so on. It becomes interesting and important to know when we have used this ordering in an essential way. For this we need the concept of isomorphism. Two structures are isomorphic iff they are identical except perhaps for the names of the elements of their universes:

Definition 1.21 (Isomorphism of Unordered Structures) Let \mathcal{A} and \mathcal{B} be structures of vocabulary $\tau = \langle R_1^{a_1}, \dots, R_r^{a_r}, c_1, \dots, c_s \rangle$. We say that \mathcal{A} is *isomorphic* to \mathcal{B} , written, $\mathcal{A} \cong \mathcal{B}$, iff there is a map $f : |\mathcal{A}| \rightarrow |\mathcal{B}|$ with the following properties:

1. f is 1:1 and onto.
2. For every input relation symbol R_i and for every a_i -tuple of elements of $|\mathcal{A}|$, e_1, \dots, e_{a_i} ,

$$\langle e_1, \dots, e_{a_i} \rangle \in R_i^{\mathcal{A}} \quad \Leftrightarrow \quad \langle f(e_1), \dots, f(e_{a_i}) \rangle \in R_i^{\mathcal{B}}$$

3. For every input constant symbol c_i , $f(c_i^{\mathcal{A}}) = c_i^{\mathcal{B}}$

The map f is called an *isomorphism*. \square

As an example, see graphs G and H in Figure 1.1 which are isomorphic using the map that adds one mod five to the numbers of the vertices of G .

Note that we have defined isomorphisms so that they need only preserve the input symbols, not the ordering and other numeric relations. If we included the ordering relation then we would have $\mathcal{A} \cong \mathcal{B}$ iff $\mathcal{A} = \mathcal{B}$. To be completely precise, we should call the mapping f defined above an “isomorphism of unordered structures” and say that \mathcal{A} and \mathcal{B} are “isomorphic as unordered structures”. (Note also that, since “unordered string” does not make sense, neither does the concept of isomorphism for strings. By a strict interpretation of Definition 1.21, two strings

would be isomorphic as unordered structures iff they had the same number of each symbol.)

The following proposition is basic.

Proposition 1.22 *Suppose \mathcal{A} and \mathcal{B} are isomorphic. Then for all sentences $\varphi \in \mathcal{L}(\tau - \{\leq\})$, \mathcal{A} and \mathcal{B} agree on φ .*

Exercise 1.23 Prove Proposition 1.22. [Hint: do this by induction using Definition 1.3.] \square

1.4 First-Order Queries

As mentioned in the introduction, we use the concept of query as the fundamental paradigm of computation:

Definition 1.24 A *query* is any mapping $I : \text{STRUC}[\sigma] \rightarrow \text{STRUC}[\tau]$ from structures of one vocabulary to structures of another vocabulary, that is polynomially bounded. That is, there is a polynomial p such that for all $\mathcal{A} \in \text{STRUC}[\sigma]$, $\|I(\mathcal{A})\| \leq p(\|\mathcal{A}\|)$. A *boolean query* is a map $I_b : \text{STRUC}[\sigma] \rightarrow \{0, 1\}$. A boolean query may be thought of as a subset of $\text{STRUC}[\sigma]$ — the set of structures \mathcal{A} for which $I_b(\mathcal{A}) = 1$.

An important subclass of queries are the order-independent queries. (In database theory the term “generic” is often used instead of “order-independent”.) Let I be a query defined on $\text{STRUC}[\sigma]$. Then I is *order-independent* iff for all isomorphic structures $\mathcal{A}, \mathcal{B} \in \text{STRUC}[\sigma]$, $I(\mathcal{A}) \cong I(\mathcal{B})$. For boolean queries, this last condition translates to $I(\mathcal{A}) = I(\mathcal{B})$. \square

From our point of view, the simplest kind of query is a first-order query. As an example, any first-order sentence $\varphi \in \mathcal{L}(\tau)$ defines a boolean query I_φ on $\text{STRUC}[\tau]$ where $I_\varphi(\mathcal{A}) = 1$ iff $\mathcal{A} \models \varphi$.

For example, let $\text{DIAM}[8]$ be the query on graphs that is true of a graph iff its diameter is at most eight. This is a first-order query given by the formula,

$$\text{DIAM}[8] \equiv (\forall xy)\varphi_{dist8}$$

where φ_{dist8} , meaning that there is a path from x to y of length at most eight, was written in Example 1.4.

As another example, consider the query I_{add} , which, given a pair of natural numbers represented in binary, returns their sum. This query is defined by the first-order formula φ_{add} from Proposition 1.9. More explicitly, let $\mathcal{A} = \langle |\mathcal{A}|, \leq, A, B \rangle$ be any structure in $\text{STRUC}[\tau_{ab}]$. \mathcal{A} is a pair of natural numbers each of $n = \|\mathcal{A}\|$ bits. Their sum is given by $I_{add}(\mathcal{A}) = \langle |\mathcal{A}|, S \rangle$ where,

$$S = \{a \in |\mathcal{A}| \mid (\mathcal{A}, a/x) \models \varphi_{add}\} \quad (1.25)$$

The first-order query $I_{add} : \text{STRUC}[\tau_{ab}] \rightarrow \text{STRUC}[\tau_s]$ maps structure \mathcal{A} to another structure with the same universe, i.e., $|\mathcal{A}| = |I_{add}(\mathcal{A})|$. The following is a general definition of a k -ary first-order query. Such a query maps any structure \mathcal{A} to a structure whose universe is a first-order definable subset of all k -tuples from $|\mathcal{A}|$. Each relation R_i over $I(\mathcal{A})$ is a first-order definable subset of $|I(\mathcal{A})|^{a_i}$. The constants of $I(\mathcal{A})$ are first-order definable elements of $|\mathcal{A}|^k$.

Definition 1.26 (First-Order Queries) Let σ and τ be any two vocabularies where $\tau = \langle R_1^{a_1}, \dots, R_r^{a_r}, c_1, \dots, c_s \rangle$, and let k be a fixed natural number. We want to define the notion of a first-order query,

$$I : \text{STRUC}[\sigma] \rightarrow \text{STRUC}[\tau] .$$

I is given by an $r + s + 1$ -tuple of formulas, $\varphi_0, \varphi_1, \dots, \varphi_r, \psi_1, \dots, \psi_s$, from $\mathcal{L}(\sigma)$. For each structure $\mathcal{A} \in \text{STRUC}[\sigma]$, these formulas describe a structure $I(\mathcal{A}) \in \text{STRUC}[\tau]$,

$$I(\mathcal{A}) = \langle |I(\mathcal{A})|, R_1^{I(\mathcal{A})}, \dots, R_r^{I(\mathcal{A})}, c_1^{I(\mathcal{A})}, \dots, c_s^{I(\mathcal{A})} \rangle .$$

The universe of $I(\mathcal{A})$ is a first-order definable subset¹ of $|\mathcal{A}|^k$,

$$|I(\mathcal{A})| = \{ \langle b^1, \dots, b^k \rangle \mid \mathcal{A} \models \varphi_0(b^1, \dots, b^k) \}$$

Each relation $R_i^{I(\mathcal{A})}$ is a first-order definable subset of $|I(\mathcal{A})|^{a_i}$,

$$R_i^{I(\mathcal{A})} = \{ \langle \langle b_1^1, \dots, b_1^{a_i} \rangle, \dots, \langle b_{a_i}^1, \dots, b_{a_i}^{a_i} \rangle \rangle \in |I(\mathcal{A})|^{a_i} \mid \mathcal{A} \models \varphi_i(b_1^1, \dots, b_{a_i}^{a_i}) \} .$$

Each constant symbol $c_j^{I(\mathcal{A})}$ is a first-order definable element of $|I(\mathcal{A})|$,

$$c_j^{I(\mathcal{A})} = \text{the unique } \langle b^1, \dots, b^k \rangle \in |I(\mathcal{A})| \text{ such that } \mathcal{A} \models \psi_j(b^1, \dots, b^k) .$$

¹Usually we will take $\varphi_0 \equiv \mathbf{true}$, thus letting $|I(\mathcal{A})| = |\mathcal{A}|^k$, cf. Remark 1.32.

When we need to be formal, we let $a = \max\{a_i \mid 1 \leq i \leq r\}$ and let the free variables of φ_i be $x_1^1, \dots, x_1^k, \dots, x_{a_i}^1, \dots, x_{a_i}^k$. The free variables of φ_0 and the ψ_j 's are x_1^1, \dots, x_1^k .

If the formulas ψ_j have the property that for all $\mathcal{A} \in \text{STRUC}[\sigma]$,

$$|\{(b^1, \dots, b^k) \in |\mathcal{A}|^k \mid (\mathcal{A}, b^1/x_1^1, \dots, b^k/x_1^k) \models \varphi_0 \wedge \psi_j\}| = 1$$

then we write $I = \lambda_{x_1^1 \dots x_1^k} \langle \varphi_0, \dots, \psi_s \rangle$ and say that I is a k -ary first-order query from $\text{STRUC}[\sigma]$ to $\text{STRUC}[\tau]$.

It is often possible to name constant $c_j^{I(\mathcal{A})}$ explicitly as a k -tuple of constants, $\langle t^1, \dots, t^k \rangle$. In this case, we may simply write this tuple in place of its corresponding defining formula,

$$\psi_j \equiv x_1^1 = t^1 \wedge \dots \wedge x_1^k = t^k .$$

As another example, in a 3-ary query I , the numerical constants 0, 1, and \max will be mapped to the following:

$$0^{I(\mathcal{A})} = \langle 0, 0, 0 \rangle; \quad 1^{I(\mathcal{A})} = \langle 0, 0, 1 \rangle; \quad \max^{I(\mathcal{A})} = \langle \max, \max, \max \rangle$$

A first-order query is either boolean, and thus defined by a first-order sentence, or is a k -ary first-order query, for some k .

Let FO be the set of first-order boolean queries. Let $Q(\text{FO})$ be the set of all first-order queries. \square

Example 1.27 Consider the genealogical database from Example 1.2. The following pair of formulas define a unary query, $I_{sa} = \lambda_{xy} \langle \mathbf{true}, \varphi_{sibling}, \varphi_{aunt} \rangle$, from genealogical databases to structures of vocabulary $\langle \text{SIBLING}^2, \text{AUNT}^2 \rangle$:

$$\begin{aligned} \varphi_{sibling}(x, y) &\equiv (\exists fm)(x \neq y \wedge f \neq m \wedge P(f, x) \wedge P(f, y) \wedge P(m, x) \wedge P(m, y)) \\ \varphi_{aunt}(x, y) &\equiv (\exists ps)(P(p, y) \wedge \varphi_{sibling}(p, s) \\ &\quad \wedge (s = x \vee S(x, s))) \wedge F(x) \end{aligned}$$

Codd defined a database query language as “complete” if it could express all first-order queries. As we will see, many queries of interest are not first-order. One such example is the ancestor query on genealogical databases (Exercise 6.46). \square

As another example, the first-order query I_{add} (Equation (1.25)) is a unary query, i.e., $k = 1$, given by $I_{add} = \lambda_x \langle \mathbf{true}, \varphi_{add} \rangle$. In this case, $\varphi_0 = \mathbf{true}$ means that the universe of $I_{add}(\mathcal{A})$ is equal to the universe of \mathcal{A} .

We will see later that $Q(\text{FO})$ is a very robust class of queries. For now, the reader should check the following proposition, which says that first-order queries are closed under composition.

Proposition 1.28 *Let $I_1 : \text{STRUC}[\sigma] \rightarrow \text{STRUC}[\tau]$ be a k -ary first-order query and let $I_2 : \text{STRUC}[\tau] \rightarrow \text{STRUC}[v]$ be an m -ary first-order query. Then $I_2 \circ I_1 : \text{STRUC}[\sigma] \rightarrow \text{STRUC}[v]$ is an mk -ary first-order query.*

Exercise 1.29 Consider the following binary first-order query from graphs to graphs: $I = \lambda_{x,y,x',y'} \langle \mathbf{true}, \alpha, \langle 0, 0 \rangle, \langle \mathit{max}, \mathit{max} \rangle \rangle$, where

$$\alpha(x, y, x', y') \equiv (x = x' \wedge E(y, y')) \vee (\text{SUC}(x, y) \wedge x' = y' = y)$$

Recall that part of the meaning of this query is that given a structure $\mathcal{A} \in \text{STRUC}[\tau_g]$, with $n = \|\mathcal{A}\|$,

$$|I(\mathcal{A})| = \{ \langle i, j \rangle \mid i, j \in |\mathcal{A}| \}; \quad s^{I(\mathcal{A})} = \langle 0, 0 \rangle; \quad t^{I(\mathcal{A})} = \langle n-1, n-1 \rangle .$$

1. Show that I has the following interesting property: For all undirected graphs G ,

$$(G \text{ is connected}) \Leftrightarrow (t \text{ is reachable from } s \text{ in } I(G))$$

2. Recall that a graph is *strongly connected* iff for every pair of vertices g and h , there is a path in G from g to h . Modify I to be a 3-ary query I' such that for all directed graphs G ,

$$(G \text{ is strongly connected}) \Leftrightarrow (t \text{ is reachable from } s \text{ in } I'(G)) .$$

[Hint: I almost works, but we need to also make sure that there is a path in G from max to 0 .] \square

Exercise 1.30 Show that the set of first-order queries is closed under composition, i.e., prove Proposition 1.28. \square

Exercise 1.31 The first-order query I_{add} defined in Equation (1.25) has the defect that it ignores the possibility that the sum of two n -bit numbers might be $n+1$ bits. Show how to define a more robust first-order query that returns the always correct $n+1$ -bit sum. Going further, show how to define the first-order query that always returns the correct sum and has no superfluous leading 0's. \square

Remark 1.32 If I is a first-order query on ordered structures, then it must include first-order definitions of the numeric relations and constants. Unless we state otherwise, the ordering on $I(\mathcal{A})$ will be the lexicographic ordering of k -tuples \leq^k inherited from \mathcal{A} : $\leq^1 = \leq$ and inductively,

$$\langle x_1, \dots, x_k \rangle \leq^k \langle y_1, \dots, y_k \rangle \equiv x_1 < y_1 \vee (x_1 = y_1 \wedge \langle x_2, \dots, x_k \rangle \leq^{k-1} \langle y_2, \dots, y_k \rangle)$$

In the following exercise, you are asked to write the definition of the remaining numeric relations and constants, assuming that $\varphi_0 \equiv \mathbf{true}$. For the first-order queries in this book, we usually limit ourselves to the case that $\varphi_0 \equiv \mathbf{true}$. If not, we must express the new numeric relations explicitly.

Exercise 1.33 Let I be a first-order query on ordered structures. The successor and bit relations must be defined.

1. Give the formulas defining 0, 1, and max the minimum, second, and maximum elements of the new universe under the lexicographical ordering. Note that if $\varphi_0 \equiv \mathbf{true}$, then the resulting constants are just k -tuples of constants:

$$0^{I(\mathcal{A})} = \langle 0, \dots, 0 \rangle; 1^{I(\mathcal{A})} = \langle 0, \dots, 0, 1 \rangle; max^{I(\mathcal{A})} = \langle max, \dots, max \rangle$$

However, in the more general case you must use quantifiers to say that the given element is the minimum, second, maximum in the lexicographical ordering.

2. Assuming that $\varphi_0 \equiv \mathbf{true}$, write a quantifier-free formula defining the new SUC relation.
3. Assuming that $\varphi_0 \equiv \mathbf{true}$, write the formula defining the new BIT relation. [Hint: by Theorem 1.17 you may define addition and multiplication on k -tuples.]

□

Without the assumption that $\varphi_0 \equiv \mathbf{true}$, BIT need not be first-order definable in the image structures. For example, if $\sigma = \tau_s$ and $\varphi_0(x) \equiv S(x)$, then the parity of the universe of $I(\mathcal{A})$ is not first-order expressible in \mathcal{A} (Theorem 13.1). If BIT were definable in $I(\mathcal{A})$ then so would the parity of its universe. For this reason, when we define first-order reductions, we restrict our attention to very simple formulas, φ_0 , that define the universe of the image structure.

Historical Notes and Suggestions for Further Reading

There are many excellent introductions to logic. We especially recommend [End72] and [EFT94]. The recent books on finite model theory [EF95], [LR96] and [Ott97] complement this book. For history of logic, it is wonderful to go back to some of the original sources, carefully translated and annotated in [vH67].

The definition of semantics for first-order logic (Definition 1.3) is usually attributed to Tarski [Tar36].

Lemma 1.18 was originally proved by Barrington, Immerman, and Straubing in [BIS88]. We got part 2 of Theorem 1.17 from Lindell [L], who says that the result comes from page 299 of Hajek and Pudlak [HP93]. However, on page 406 of [HP93] the result is attributed to Bennett [Ben62]. See also [DDLW98], where Dawar, Lindell, and Weinstein prove that ordering is definable when the only given numeric predicate is BIT. It follows that Theorem 1.17 remains valid when ordering is not given because ordering is easily definable from PLUS.

Exercise 1.12 is from [vD94].

See [CH80a] for perhaps the first study of the set of order-independent, computable queries.

One very important topic in a standard course on first-order logic that is omitted here is the study of proofs. In particular, the following two theorems are basic and appear in every standard logic book. They were originally proved by Gödel in his Ph.D. thesis [Göd30].

Theorem 1.34 (Completeness Theorem for First-Order Logic) *There is a complete recursive axiomatization for the set of formulas valid in all — finite and infinite — structures.*

Theorem 1.35 (Compactness Theorem for First-Order Logic) *Let Γ be a set of first-order formulas with the property that every finite subset of Γ has a (perhaps infinite) model. Then Γ has a (perhaps infinite) model.*

These theorems fail when we restrict our attention to finite structures. From the Completeness Theorem it follows that the set of valid formulas for first-order logic is recursively enumerable (r.e.), and VALID is in fact r.e.-complete. Thus, the set of satisfiable formulas is not r.e. For finite structures, Trahtenbrot's Theorem says that the reverse is true: The set of formulas satisfiable in a finite structure is

r.e.-complete, so the set of formulas valid in all finite structures is not r.e., and thus not axiomatizable [Tra50].

For a finite alphabet, $\Sigma = \{\sigma_1, \dots, \sigma_r\}$, consider the vocabulary $\tau_\Sigma = \langle S_{\sigma_1}^1, \dots, S_{\sigma_r}^1 \rangle$. The set $\text{STRUC}[\tau_\Sigma]$ consists of the set of non-empty words of vocabulary Σ . Languages without BIT are well-studied for these vocabularies. The following two theorems are fundamental:

Theorem 1.36 *The set of boolean queries expressible in second-order, monadic logic, without BIT, over the vocabularies τ_Σ consist exactly of the regular languages. In symbols,*

$$\text{SO(monadic)(wo BIT)} = \text{Regular} .$$

Theorem 1.37 *The set of boolean queries expressible in first-order logic, without BIT, over the vocabularies τ_Σ consist exactly of the star-free regular languages. In symbols,*

$$\text{FO(wo BIT)} = \text{star-free Regular} .$$

Theorem 1.36 is due to Büchi [Büc60] and Theorem 1.37 is due to McNaughton and Papert, [MP71]. From the point of view of this book, the languages without BIT are slightly too weak to provide a robust view of computation. A good reference for these results and other relations between logic and automata theory is [Str94] by Straubing.

Chapter 2

Background in Complexity

Computational Complexity measures the amount of computational resources, such as time and space, that are needed, as a function of the size of the input, to compute a query. This chapter introduces the reader to complexity theory. We define the complexity measures and complexity classes that we study in the rest of the book. We also explain some of their basic properties, complete problems, and inter-relationships.

2.1 Introduction

In the 1930's many models of computation were invented, including Church's lambda calculus, Gödel's recursive functions, Markov algorithms and Turing machines. It is very striking that these interesting and apparently different models — all independent efforts to precisely define the intuitive notion of “mechanical procedure” — were proved equivalent. This leads to the universally accepted “Church's thesis”, which states that the intuitive concept of what can be “automatically computed” is appropriately captured by the Turing machine (and all its variants).

If one appropriately measures the complexity of computation in a Markov algorithm, a lambda expression, a recursive function, or a Turing machine, one obtains equivalent values. A consequence of this is that efficiency is not model-dependent, but is in fact a fundamental concept.

We get the same theory of complexity whether we approach it via Turing machines or any of these other models. As we will see in this book, descriptive complexity gives definitions of complexity that are equivalent to those of all the

above models. Different formalisms lend themselves to different ways of thinking and working. We find that the insights gained from the descriptive approach to complexity offer a different point of view from more traditional machine-based complexity. In particular, there are well understood methods in logic for ascertaining what can and cannot be expressed in a given language. We will introduce some of these methods in Chapter 6.

In Descriptive complexity, we measure the difficulty of describing queries. We will see that natural measures of descriptive complexity such as depth of nesting of quantifiers and number of variables correspond closely to natural notions of complexity in Turing machines.

2.2 Preliminary Definitions

We assume that the reader is familiar with the Turing machine. We start from there and present a survey of computational complexity theory.

We write $M(w)\downarrow$ to mean that Turing machine M accepts input w , and we write $L(M)$ to denote the language accepted by M ,

$$L(M) = \{w \in \{0, 1\}^* \mid M(w)\downarrow\}.$$

Instead of just accepting or rejecting, Turing machines may compute functions from binary strings to binary strings. We use $T(w)$ to denote the binary string that Turing machine T leaves on its write-only output tape when it is started with the binary string w on its input tape. If T does not halt on input w , then $T(w)$ is undefined.

Everything that a Turing machine does may be thought of as a query from binary strings to binary strings. In order to make Descriptive complexity rich and flexible it is useful to consider queries that use other vocabularies. To relate such queries to Turing machine complexity, we fix a scheme that encodes the structures of vocabulary τ as boolean strings. To do this, for each τ , we define an encoding query,

$$\text{bin}_\tau : \text{STRUC}[\tau] \rightarrow \text{STRUC}[\tau_s]$$

Recall that $\tau_s = \langle S^1 \rangle$ is the vocabulary of boolean strings. The details of the encoding are not important, but it is useful to know that for each τ , bin_τ and its inverse are first-order queries (Exercise 2.3).

Definition 2.1 (The binary encoding of structures: $\text{bin}(\mathcal{A})$) Let $\tau = \langle R_1^{a_1}, \dots, R_r^{a_r}, c_1, \dots, c_s \rangle$ be a vocabulary, and let $\mathcal{A} = \langle \{0, 1, \dots, n -$

$1\}, R_1^{\mathcal{A}} \dots R_r^{\mathcal{A}}, c_1^{\mathcal{A}} \dots c_s^{\mathcal{A}})$ be an ordered structure of vocabulary τ . The relation $R_i^{\mathcal{A}}$ is a subset of $|\mathcal{A}|^{a_i}$. We encode this relation as a binary string $\text{bin}^{\mathcal{A}}(R_i)$ of length n^{a_i} where “1” in a given position indicates that the corresponding tuple is in $R_i^{\mathcal{A}}$. Similarly, for each constant $c_j^{\mathcal{A}}$, its number is encoded as a binary string $\text{bin}^{\mathcal{A}}(c_j)$ of length $\lceil \log n \rceil$.

The binary encoding of the structure \mathcal{A} is then just the concatenation of the bit strings coding its relations and constants,

$$\text{bin}_{\tau}(\mathcal{A}) = \text{bin}^{\mathcal{A}}(R_1)\text{bin}^{\mathcal{A}}(R_2) \cdots \text{bin}^{\mathcal{A}}(R_r)\text{bin}^{\mathcal{A}}(c_1) \cdots \text{bin}^{\mathcal{A}}(c_s)$$

We do not need any separators between the various relations and constants because the vocabulary τ and the length of $\text{bin}_{\tau}(\mathcal{A})$ determines where each section belongs. Observe that the length of $\text{bin}_{\tau}(\mathcal{A})$ is given by

$$\hat{n}_{\tau} = \|\text{bin}_{\tau}(\mathcal{A})\| = n^{a_1} + \cdots + n^{a_r} + s \lceil \log n \rceil \quad (2.2)$$

Note: We do not bother to include any numeric predicates or constants in $\text{bin}_{\tau}(\mathcal{A})$ since they can be easily recomputed. However, the coding $\text{bin}_{\tau}(\mathcal{A})$ does presuppose an ordering on the universe. There is no way to code a structure as a string without an ordering. Since a structure determines its vocabulary, in the sequel we usually write $\text{bin}(\mathcal{A})$ instead of $\text{bin}_{\tau}(\mathcal{A})$ for the binary encoding of $\mathcal{A} \in \text{STRUC}[\tau]$. Here bin is the union of bin_{τ} over all vocabularies τ . In the special case where τ includes no input relations symbols, we pretend that there is a unary relation symbol that is always false. For example, if $\tau = \emptyset$, then $\text{bin}(\mathcal{A}) = 0^{|\mathcal{A}|}$. We do this to insure that the size of $\text{bin}(\mathcal{A})$ is at least as large as $\|\mathcal{A}\|$. \square

When $\tau = \tau_s$, the map bin_{τ_s} maps strings to strings. The reader should check from the above definition that in this case, bin_{τ_s} is the identity map and thus $\hat{n}_{\tau_s} = n$.

In complexity theory, n is usually reserved for the length of the input. However, in this book, n is used to denote the size of the input structure, $n = \|\mathcal{A}\|$. When the inputs are structures of vocabulary τ , the length of the input is \hat{n}_{τ} . For the case of binary strings, these two sizes coincide because $\hat{n}_{\tau_s} = n$. When τ is understood, we write \hat{n} for \hat{n}_{τ} . Observe that n and \hat{n} are always polynomially related.

There are two requirements of a coding function such as “bin”. First, it must be computationally very easy to encode and decode. Secondly, the coding must be fairly space efficient, e.g., coding in unary would not be acceptable. In the next

exercise, the reader is asked to show that both the encoding and decoding of bin are first-order expressible.

Exercise 2.3 Show that for any vocabulary τ , the queries $\text{bin}_\tau : \text{STRUC}[\tau] \rightarrow \text{STRUC}[\tau_s]$ and its inverse $\text{bin}_\tau^{-1} : \text{STRUC}[\tau_s] \rightarrow \text{STRUC}[\tau]$ are first-order queries. More explicitly, let $\tau = \langle R_1^{a_1}, \dots, R_r^{a_r}, c_1, \dots, c_s \rangle$.

1. Construct a first-order query β_τ that is equal to the mapping bin_τ .
2. Construct a first-order query $\delta_\sigma : \text{STRUC}[\tau_s] \rightarrow \text{STRUC}[\tau]$, such that for all $\mathcal{A} \in \text{STRUC}[\tau]$, $\delta_\sigma(\text{bin}_\tau(\mathcal{A})) = \mathcal{A}$. The query δ should be unary, that is, $k = 1$ in the definition of k -ary first-order query (Definition 1.26). \square

Using the encoding bin , we define what it means for a Turing machine to compute a query:

Definition 2.4 Let $I : \text{STRUC}[\sigma] \rightarrow \text{STRUC}[\tau]$ be a query. Let T be a Turing machine. Suppose that for all $\mathcal{A} \in \text{STRUC}[\sigma]$, $T(\text{bin}(\mathcal{A})) = \text{bin}(I(\mathcal{A}))$. Then we say that T computes I . \square

We use the notation $\text{DTIME}[t(n)]$ to denote the set of boolean queries that are computable by a deterministic, multi-tape Turing machine in $O(t(n))$ steps for inputs of universe size n .¹ Similarly we use $\text{NTIME}[t(n)]$, $\text{DSPACE}[s(n)]$, and $\text{NSPACE}[s(n)]$ to denote nondeterministic time, deterministic space and nondeterministic space, respectively. We assume that the reader is familiar with the following classical complexity classes: $\text{L} = \text{DSPACE}[\log n]$, $\text{NL} = \text{NSPACE}[\log n]$, $\text{P} = \text{polynomial time} = \bigcup_{k=1}^{\infty} \text{DTIME}[n^k]$, $\text{NP} = \text{nondeterministic polynomial time} = \bigcup_{k=1}^{\infty} \text{NTIME}[n^k]$, $\text{PSPACE} = \text{polynomial space} = \bigcup_{k=1}^{\infty} \text{DSPACE}[n^k] = \bigcup_{k=1}^{\infty} \text{NSPACE}[n^k]$, and $\text{EXPTIME} = \text{exponential time} = \bigcup_{k=1}^{\infty} \text{DTIME}[2^{n^k}]$. To talk about space $s(n)$, for $s(n) < \hat{n}$, the Turing machine is assumed to have a read-only input tape of length \hat{n} and some number of work tapes of total length $O(s(n))$.

In the definition of complexity classes, we consider only boolean queries. This is in order to be consistent with the standard definitions of complexity classes as sets of decision problems, i.e., boolean queries. For any complexity class \mathcal{C} , we use the notation $Q(\mathcal{C})$ to denote the set of all queries that are computable in the complexity class \mathcal{C} . Since \mathcal{C} just consists of boolean queries, what does it mean for

¹The usual definition in complexity theory writes $t(n)$ as the function $t'(\hat{n})$, a polynomially-related function of the length of the encoding of the input. We use n to be the size of the universe of the input structure and measure all sizes in this uniform way.

a general query to be “computable in \mathcal{C} ”? It means that each bit of $\text{bin}(I(\mathcal{A}))$ is uniformly computable in \mathcal{C} from $\text{bin}(\mathcal{A})$. In other words,

Definition 2.5 (*$Q(\mathcal{C})$, the Queries Computable in \mathcal{C}*) Let $I : \text{STRUC}[\sigma] \rightarrow \text{STRUC}[\tau]$ be a query. We say that I is *computable in \mathcal{C}* iff the boolean query I_b is an element of \mathcal{C} , where

$$I_b = \{(\mathcal{A}, i, a) \mid \text{The } i^{\text{th}} \text{ bit of } \text{bin}(I(\mathcal{A})) \text{ is “}a\text{”}\}. \quad (2.6)$$

Let $Q(\mathcal{C})$ be the set of all queries computable in \mathcal{C} :

$$Q(\mathcal{C}) = \mathcal{C} \cup \{I \mid I_b \in \mathcal{C}\} \quad \square$$

For each of the above resources (deterministic and nondeterministic time and space) there is a hierarchy theorem saying that more of the given resource enables us to compute more boolean queries (see Exercise 2.8). These theorems are proved by diagonalization arguments.

We say that a function $s : \mathbf{N} \rightarrow \mathbf{N}$ is *space constructible* (resp. *time constructible*) iff there is a deterministic Turing machine running in space $O(s(n))$, (resp. time $O(s(n))$) that on input 0^n , i.e., n in unary, computes $s(n)$ in binary. This is the same thing as saying that $s' \in Q(\text{DSPACE}[s(n)])$, resp. $s' \in Q(\text{DTIME}[s(n)])$ where s' is the function that on input 0^n computes $s(n)$ in binary.

Every reasonable function is constructible, as is every function one finds in this book. Many theorems in this book need to assume that the time and space bounds in question are constructible.

Exercise 2.7

1. Show that the following functions are time constructible: n , n^2 , $\lceil n \log n \rceil$, 2^n .
2. Show that the following are space constructible: n , n^2 , $\lceil n \log n \rceil$, 2^n , $\lceil \log n \rceil$.

□

Exercise 2.8 Prove the **Space Hierarchy Theorem**: For all space constructible $s(n) \geq \log n$, if $\lim_{n \rightarrow \infty} (t(n)/s(n)) = 0$, then $\text{DSPACE}[t(n)]$ is strictly contained in $\text{DSPACE}[s(n)]$.

[Hint: this is a diagonalization argument, but you have to be careful. On input M , the diagonalization program marks off $s(|M|)$ tape cells and then simulates machine M on input M . If $M(M)$ exceeds the given space or loops, then it should accept. Otherwise, do the opposite of what M would do.] \square

When comparing different resources, we are able to prove much less. For example, by Savitch's Theorem (Theorem 2.32), for $s(n) \geq \log n$,

$$\text{DSPACE}[s(n)] \subseteq \text{NSPACE}[s(n)] \subseteq \text{DSPACE}[(s(n))^2];$$

However, we know only the trivial relationships between nondeterministic and deterministic time:

$$\text{DTIME}[t(n)] \subseteq \text{NTIME}[t(n)] \subseteq \text{DTIME}[2^{O(t(n))}].$$

Consider the following series of containments. It follows from Savitch's Theorem and the Space Hierarchy Theorem that NL is not equal to PSPACE; but even now, more than twenty-five years after the introduction of the classes P and NP, no other inequalities, including that L is not equal to NP, are known.

$$L \subseteq NL \subseteq P \subseteq NP \subseteq \text{PSPACE}$$

2.3 Reductions and Complete Problems

There are several ways to compare the complexity of boolean queries. Perhaps the most natural is the Turing reduction. Complexity in this model is defined via oracles. Let A and B be boolean queries that may be difficult to compute. An *oracle* for B is a mythical device that when given a structure \mathcal{B} will answer in unit time whether or not \mathcal{B} satisfies query B .

Turing gave his original definition of oracles for the case of unsolvable problems. After proving that the halting problem, K , is undecidable, he considered the question of what can be decided by a Turing machine that had an oracle for K . In complexity theory, we use oracles for computable but difficult sets. Such an oracle can speed up some computations.

We say that A is *Turing reducible* to B if it is easy to compute query A given an oracle for B . The following makes this definition more precise.

Definition 2.9 An *oracle Turing machine* is a Turing machine equipped with an extra tape called the query tape. Let M be an oracle Turing machine and B be any

boolean query. We write M^B to denote the oracle Turing machine M equipped with an oracle for set B . M^B may write on its query tape like any other tape. At any time, M^B may enter the “query state”. Assume that the string $w = \text{bin}(\mathcal{A})$ is written on the query tape when M^B enters the query state. At the next step, “1” will appear on the query tape if $\mathcal{A} \in B$ and “0” otherwise. Thus, M^B may answer any membership question “Does \mathcal{A} satisfy B ?” in linear time: the time to copy the string $\text{bin}(\mathcal{A})$ to its query tape.

Given two boolean queries A, B and a complexity class \mathcal{C} , we say that A is \mathcal{C} -Turing reducible to B iff there exists an oracle Turing machine M such that M^B runs in complexity class \mathcal{C} and $\mathcal{L}(M^B) = A$. We denote this by $A \leq_{\mathcal{C}}^T B$. The superscript “ T ” stands for Turing reduction. An important example is the polynomial-time Turing reduction, \leq_p^T . \square

Example 2.10 As an example of a Turing reduction, define the boolean query CLIQUE to be the set of pairs $\langle G, k \rangle$ such that G is a graph having a complete subgraph of size k . The vocabulary for CLIQUE is $\tau_{gk} = \langle E^2, k \rangle$. As usual, we can identify the universe of a structure $\mathcal{A} \in \text{STRUC}[\tau_{gk}]$ with the set $\{0, 1, \dots, n-1\}$, where $n = \|\mathcal{A}\|$ is the number of vertices of \mathcal{A} , and the constant k thus represents not only a vertex, but a number between 0 and $n-1$. (We will see later that CLIQUE is an NP-complete problem.)

Define the query MAX-CLIQUE(G) to be the size of a largest clique in graph G . We show that the boolean version of MAX-CLIQUE is polynomial-time Turing reducible to CLIQUE. In symbols, this would be written,

$$\text{MAX-CLIQUE}_b \leq_p^T \text{CLIQUE} .$$

Where MAX-CLIQUE_b is defined as in Equation (2.6),

$$\text{MAX-CLIQUE}_b = \{(G, i, a) \mid \text{bit } i \text{ of MAX-CLIQUE}(G) \text{ is “a”}\} .$$

The reduction is as follows. Given (G, i, a) , perform binary search using an oracle for CLIQUE to determine the size s of the maximum clique for G . This is done by asking if $(G, n/2) \in \text{CLIQUE}$. If yes, ask about $(G, 3n/4)$, if no, ask about $(G, n/4)$. After $\log n$ queries to the oracle, s has been computed. Now accept iff bit i of s is “a”. \square

A simpler and more popular kind of reduction in complexity theory is the many-one reduction. (In descriptive complexity, we use first-order reductions. These are first-order queries that are at the same time many-one reductions.)

Definition 2.11 (Many-One Reduction) Let \mathcal{C} be a complexity class, and let $A \subseteq \text{STRUC}[\sigma]$ and $B \subseteq \text{STRUC}[\tau]$ be boolean queries. Suppose that the query $I : \text{STRUC}[\sigma] \rightarrow \text{STRUC}[\tau]$ is an element of $Q(\mathcal{C})$ with the property that for all $\mathcal{A} \in \text{STRUC}[\sigma]$,

$$\mathcal{A} \in A \quad \Leftrightarrow \quad I(\mathcal{A}) \in B$$

Then I is a \mathcal{C} -many-one reduction from A to B . We say that A is \mathcal{C} -many-one reducible to B , in symbols, $A \leq_{\mathcal{C}} B$. For example, when I is a first-order query (Definition 1.26), it is a first-order reduction (\leq_{fo}), when $I \in Q(\text{L})$, it is a logspace reduction (\leq_{log}); and when $I \in Q(\text{P})$, it is a polynomial-time reduction (\leq_{p}). \square

Observe that a many-one reduction is a particularly simple kind of Turing reduction: To decide whether \mathcal{A} is an element of A , compute $I(\mathcal{A})$ and ask the oracle whether $I(\mathcal{A})$ is an element of B . Many-one reductions are simpler than Turing reductions, and they seem to suffice in most situations.

Example 2.12 We give a first-order reduction from PARITY to MULT_b . PARITY is the boolean query on binary strings that is true iff the string has an odd number of ones. We will see later that PARITY is not first-order (Theorem 13.1). MULT, the multiplication query, maps a pair of boolean strings of length n to their product: a boolean string of length $2n$. Let $\tau_{ab} = \langle A^1, B^1 \rangle$ be the vocabulary of structures that are a pair A, B of boolean strings. Then $\text{MULT} : \text{STRUC}[\tau_{ab}] \rightarrow \text{STRUC}[\tau_s]$. Since reductions map boolean queries to boolean queries, we actually deal with the boolean version of MULT. MULT_b is a boolean query on structures of vocabulary $\tau_{abcd} = \langle A^1, B^1, c, d \rangle$ that is true iff bit c of the product of A and B is “d”.

Recall that $\tau_s = \langle S^1 \rangle$ is the vocabulary of boolean strings. The first-order reduction $I_{PM} : \text{STRUC}[\tau_s] \rightarrow \text{STRUC}[\tau_{abcd}]$ is given by the following formulas:

$$\begin{aligned} \varphi_A(x, y) &\equiv y = \text{max} \wedge S(x) \\ \varphi_B(x, y) &\equiv y = \text{max} \\ I_{PM} &\equiv \lambda_{xy} \langle \mathbf{true}, \varphi_A, \varphi_B, \langle 0, \text{max} \rangle, \langle 0, 1 \rangle \rangle \end{aligned}$$

Observe that the effect of this reduction is to line up all the bits of string \mathcal{A} into column $n - 1$ of the generated product. (See Figure 2.13.) It follows that

$$\mathcal{A} \in \text{PARITY} \quad \Leftrightarrow \quad I_{PM}(\mathcal{A}) \in \text{MULT}_b$$

as desired. Thus, $\text{PARITY} \leq_{\text{fo}} \text{MULT}_b$. It follows that if MULT were first-order, then PARITY would be as well. We will see later that PARITY is not first-order (Theorem 13.1), so we can conclude that MULT is not first order. \square

$$\begin{array}{rcccccccccccc}
A & & & \overbrace{\hspace{2cm}} & & \overbrace{\hspace{2cm}} & & & & \overbrace{\hspace{2cm}} & & & \\
& & & 0 \ \dots \ 0 & s_0 & 0 \ \dots \ 0 & s_1 & \dots & 0 \ \dots \ 0 & s_{n-1} & & & \\
B & \times & & 0 \ \dots \ 0 & 1 & 0 \ \dots \ 0 & 1 & \dots & 0 \ \dots \ 0 & 1 & & & \\
\hline
& & & 0 \ \dots \ 0 & s_0 & 0 \ \dots \ 0 & & & \dots & & & & \\
& & & 0 \ \dots \ 0 & s_1 & 0 \ \dots \ 0 & & & \dots & & & & \\
& & & 0 \ \dots \ 0 & \vdots & 0 \ \dots \ 0 & & & \dots & & & & \\
& & & 0 \ \dots \ 0 & s_{n-1} & 0 \ \dots \ 0 & & & \dots & & & & \\
\hline
& & & \dots & \boxed{P} & & & & \dots & & & & \\
\hline
\end{array}$$

Figure 2.13: First-order reduction of PARITY to MULT

When \mathcal{C} is a weak complexity class such as FO or L, the intuitive meaning of $A \leq_{\mathcal{C}} B$ is that the complexity of problem A is less than or equal to the complexity of problem B . The intuitive meaning of A being complete for \mathcal{C} is that A is a hardest query in \mathcal{C} and in fact every query in \mathcal{C} can be rephrased as an instance of A ; more precisely,

Definition 2.14 (Completeness for a Complexity Class) Let A be a boolean query, let \mathcal{C} be a complexity class, and let \leq_r be a reducibility relation. We say that A is *complete for \mathcal{C} via \leq_r* iff

1. $A \in \mathcal{C}$, and,
2. for all $B \in \mathcal{C}$, $B \leq_r A$.

In this book, when we say that a problem is complete for a complexity class and we do not say under what reduction, then we mean via first-order reductions \leq_{fo} . It will follow from Theorem 3.1 that if a problem is complete via first-order reductions, then it is complete via logspace and polynomial-time reductions. \square

For reasons that are not well understood, naturally occurring problems tend to be complete for important complexity classes such as P, NP, and NL. Completeness was originally defined via reductions such as polynomial-time many-one reductions (\leq_p) and later, logspace reductions (\leq_{log}). However, problems complete via \leq_p and \leq_{log} tend to remain complete via \leq_{fo} .

Most natural complexity classes include a large number of interesting complete problems. Here are a few boolean queries that are complete for their respective complexity classes. We state these problems very informally here, just to give

the reader an idea. More details on these problems and completeness proofs or references are provided later in the text.

Complete for L:

- CYCLE: Given an undirected graph, does it contain a cycle?
- REACH_d: Given a directed graph, is there a deterministic path from vertex s to vertex t ? (A deterministic path is such that for every edge (u, v) on the path, there is only one edge in the graph from u .)

Complete for NL:

- REACH: Given a directed graph, is there a path from vertex s to vertex t ?
- 2-SAT: Given a boolean formula in conjunctive normal form with only two literals per clause, is it satisfiable?

Complete for P:

- CIRCUIT-VALUE-PROBLEM (CVP): Given an acyclic boolean circuit, with inputs specified, does its output gate have value one?
- NETWORK-FLOW: Given a directed graph, with capacities on its edges, and a value V , is it possible to achieve a steady-state flow of value V through the graph?

Complete for NP:

- SAT: Given a boolean formula, is it satisfiable?
- 3-SAT: Given a boolean formula in conjunctive normal form with only three literals per clause, is it satisfiable?
- CLIQUE: Given an undirected graph and a value k , does the graph have a complete subgraph with k vertices?

Complete for PSPACE:

- QSAT: Given a quantified boolean formula, is it satisfiable?

- HEX, GEOGRAPHY, GO: Given a position in the generalized versions of the games hex, geography, or go, is there a forced win for the player whose move it is?

Exercise 2.15

1. Show that the relations \leq_{fo} , \leq_{log} , and \leq_p are transitive.
2. Let \leq_r be a transitive, many-one reduction, and let A be complete for complexity class \mathcal{C} via \leq_r . Let T be any boolean query. Show that T is complete for \mathcal{C} via \leq_r iff the following two conditions hold:

$$T \in \mathcal{C} \quad \text{and} \quad A \leq_r T$$

□

Exercise 2.16

1. Show the the value $\lceil \log n \rceil$ is first-order expressible.
2. Define the majority query as follows:

$$\text{MAJ} = \{ \mathcal{A} \in \text{STRUC}[\tau_s] \mid \text{string } \mathcal{A} \text{ contains more than } \|\mathcal{A}\|/2 \text{ “1”s} \}$$

Modify the reduction in Example 2.12 to show that $\text{MAJ} \leq_{fo} \text{MULT}_b$.

□

To prove a natural problem complete for a complexity class, one usually reduces another natural complete problem to it as in Exercise 2.15, part (2). Proving that the first natural complete problem is complete is more subtle, and we defer these proofs to later chapters. The following exercise, however, introduces an unnatural but universal kind of complete problem.

Exercise 2.17 Consider the following boolean query over strings from $\{0, 1, \#\}^*$.

$$U_{\text{time}} = \{ M\#w\#^r \mid M(w) \downarrow \text{ in } r \text{ steps} \}$$

Show that U_{time} is complete for P via first-order, many-one reductions (\leq_{fo}).
Hint: you must show that $U_{\text{time}} \in \text{P}$ and that for any problem $B \in \text{P}$, $B \leq_{fo} U_{\text{time}}$.

U_{time} . The latter is the easier part in this case: let M_B be a polynomial-time Turing machine that accepts B ; then the first-order reduction maps w to $M_B\#w\#^r$, where r is sufficiently large. \square

We next describe the problem SAT, which we will see later is NP-complete via first-order reductions (Theorem 7.16).

Example 2.18 SAT is the set of boolean formulas in conjunctive normal form (CNF) that admit a satisfying assignment, i.e., a way to set each boolean variable to **true** or **false** so that the whole formula evaluates to **true**. For example, consider the following boolean formulas:

$$\begin{aligned}\varphi_0 &= (x_1 \vee \overline{x_2} \vee x_3) \wedge (\overline{x_1} \vee \overline{x_2} \vee x_4) \wedge (x_2 \vee \overline{x_3} \vee x_5) \\ \varphi_1 &= (x_1 \vee \overline{x_2} \vee x_3) \wedge (x_1 \vee x_2 \vee \overline{x_3}) \wedge (x_1 \vee x_2 \vee x_3) \wedge (x_1 \vee \overline{x_2} \vee \overline{x_3}) \\ &\quad \wedge (\overline{x_1} \vee \overline{x_4} \vee x_5) \wedge (\overline{x_1} \vee x_4 \vee \overline{x_5}) \wedge (\overline{x_1} \vee x_4 \vee x_5) \wedge (\overline{x_1} \vee \overline{x_4} \vee \overline{x_5})\end{aligned}$$

Here the notation \overline{v} means $\neg v$. The reader should verify that $\varphi_0 \in \text{SAT}$ and $\varphi_1 \notin \text{SAT}$.

It is easy to see that SAT is in NP. In linear time, a nondeterministic algorithm can write down a “0” or a “1” for each boolean variable. Then it can deterministically check that each clause has been assigned at least one “1”, and if so, accept.

A boolean formula φ that is in CNF may be thought of as a set of clauses, each of which is a disjunction of literals. Recall that a literal is an atomic formula — in this case a boolean variable — or its negation. Thus, a natural way to encode φ is via the structure $\mathcal{A}_\varphi = \langle A, P, N \rangle$.

The universe A is a set of clauses and variables. The relation $P(c, v)$ means that variable v occurs positively in clause c and $N(c, v)$ means that v occurs negatively in c . We can think of every element of the universe as a variable and a clause. Thus, $n = \|\mathcal{A}_\varphi\|$ is equal to the maximum of the number of variables and the number of clauses occurring in φ . If v is really a variable but not a clause, we can harmlessly make it the clause $(v \vee \overline{v})$ by adding the pair (v, v) to the relations P and N . For example, a structure coding φ_0 in this way is:

$$\begin{aligned}\mathcal{A}_{\varphi_0} &= \langle \{1, 2, 3, 4, 5\}, P, N \rangle \\ P &= \{(1, 1), (1, 3), (2, 4), (3, 2), (3, 5), (4, 4), (5, 5)\} \\ N &= \{(1, 2), (2, 1), (2, 2), (3, 3), (4, 4), (5, 5)\}\end{aligned}$$

\square

We next show that $\text{SAT} \leq_{\text{fo}} \text{CLIQUE}$. It then follows from Theorem 7.16 that CLIQUE is also NP-complete.

Example 2.19 We show that SAT is first-order reducible to CLIQUE . Let \mathcal{A} be a boolean formula in CNF with clauses $C = \{c_1, \dots, c_n\}$ and variables $V = \{v_1, \dots, v_n\}$.

Let $L = \{v_1, \dots, v_n, \bar{v}_1, \dots, \bar{v}_n\}$. Define the instance of the clique problem $g(\mathcal{A}) = (V^{g(\mathcal{A})}, E^{g(\mathcal{A})}, k)$ as follows:

$$\begin{aligned} V^{g(\mathcal{A})} &= (C \times L) \cup \{w_0\} \\ E^{g(\mathcal{A})} &= \{(\langle c_1, \ell_1 \rangle, \langle c_2, \ell_2 \rangle) \mid c_1 \neq c_2 \text{ and } \bar{\ell}_1 \neq \ell_2\} \cup \\ &\quad \{(w_0, \langle c, \ell \rangle), (\langle c, \ell \rangle, w_0) \mid \ell \text{ occurs in } c\} \\ k^{g(\mathcal{A})} &= n + 1 = \|\mathcal{A}\| + 1 \end{aligned} \quad (2.20)$$

The graph $g(\mathcal{A})$ is an $n \times n$ array of vertices containing a row for every clause in \mathcal{A} and a column for every literal in L , plus a top vertex w_0 . (See Figure 2.21). There are edges between vertices $\langle c_1, \ell_1 \rangle$ and $\langle c_2, \ell_2 \rangle$ iff $c_1 \neq c_2$, i.e., the points come from different clauses; and, $\bar{\ell}_1 \neq \ell_2$, i.e., literals ℓ_1 and ℓ_2 are not the negations of each other. The other edges in the graph are between w_0 and those $\langle c, \ell \rangle$ such that literal ℓ occurs in clause c .

Observe that a clique of size $n + 1$ must involve w_0 and one vertex from each clause. This corresponds to a satisfying assignment, because no literal and its negation can be in a clique. Conversely, any satisfying assignment to \mathcal{A} determines an $n + 1$ -clique consisting of w_0 together with one literal per clause that is assigned “true”. It follows that mapping g is indeed a many-one reduction,

$$(\mathcal{A} \in \text{SAT}) \iff (g(\mathcal{A}) \in \text{CLIQUE}) \quad (2.22)$$

We now give the rather technical details of writing g as a first-order query, $g = \lambda_{x^1 x^2 x^3 y^1 y^2 y^3} \langle \varphi_0, \varphi_1, \psi_1 \rangle$. We encode the vertices as triples $\langle x^1, x^2, x^3 \rangle$, where x^1 corresponds to the clause, x^2 to the variable, and $x^3 = 1$ means the variable is positive and $x^3 = 2$ means the variable is negative. Vertex w_0 is $\langle 1, 1, 3 \rangle$, the only triple with $x^3 > 2$. See Figure 2.21, which shows part of $g(\mathcal{A})$ for a formula \mathcal{A} whose first clause is $(x_1 \vee \bar{x}_2 \vee \bar{x}_n)$. The numeric formula φ_0 , which describes the universe of $g(\mathcal{A})$, is the following,

$$\varphi_0 \equiv (x^3 \leq 2) \vee (x^1 x^2 x^3 = 113)$$

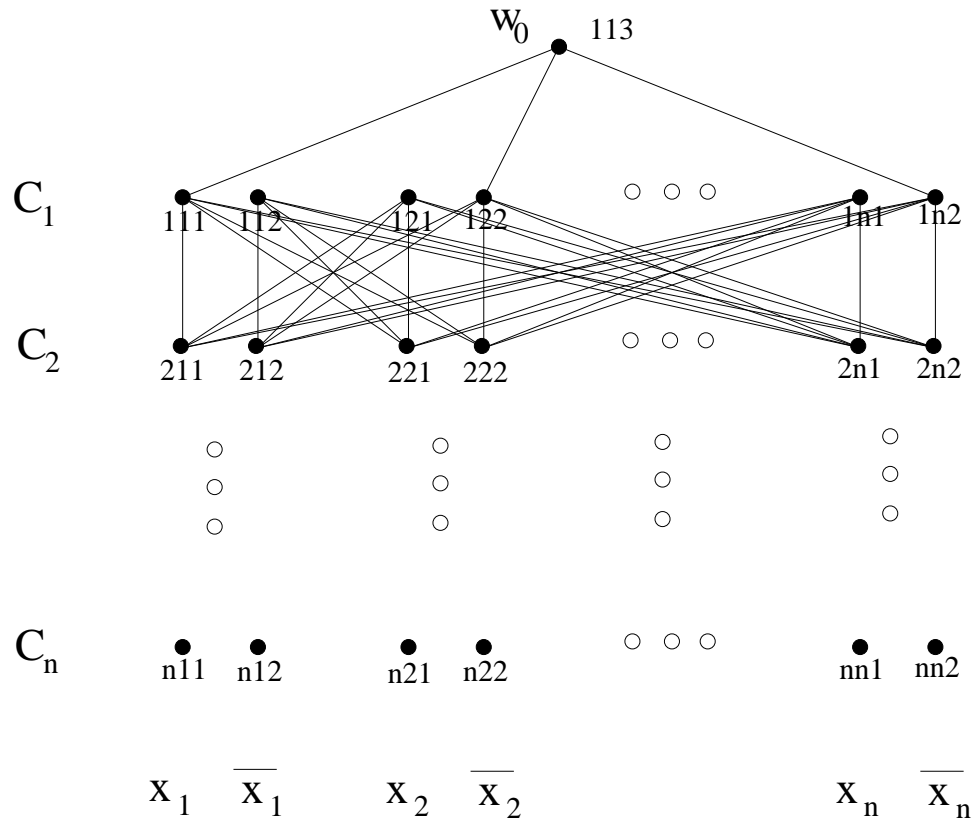


Figure 2.21: $SAT \leq_{fo} CLIQUE$, $C_1 = (x_1 \vee \overline{x_2} \vee \overline{x_n})$

To define the edge relation, let

$$\begin{aligned}\varphi'_1(\bar{x}, \bar{y}) &\equiv \alpha_1 \vee (\alpha_2 \wedge P(y^1, y^2)) \vee (\alpha_3 \wedge N(y^1, y^2)) \\ \alpha_1 &\equiv x^1 \neq y^1 \wedge x^3 < 3 \wedge y^3 < 3 \wedge (x^2 = y^2 \rightarrow x^3 = y^3) \\ \alpha_2 &\equiv x^3 = 3 \wedge y^3 = 1 \\ \alpha_3 &\equiv x^3 = 3 \wedge y^3 = 2\end{aligned}$$

Next, let φ_1 be the symmetric closure of φ'_1 ,

$$\varphi_1(x^1, x^2, x^3, y^1, y^2, y^3) \equiv \varphi'_1(x^1, x^2, x^3, y^1, y^2, y^3) \vee \varphi'_1(y^1, y^2, y^3, x^1, x^2, x^3)$$

Notice that φ_1 is a direct translation of Equation (2.20). We are thinking of the elements of the ordered universe as $1, 2, \dots, n$ instead of the usual $0, 1, \dots, n-1$. For this reason, the number $n+1$ which would usually be 011 in lexicographic order, is instead 122. Formula ψ_1 identifies k as $n+1$:

$$\psi_1(x^1, x^2, x^3) \equiv x^1 x^2 x^3 = 122$$

It follows that we have correctly encoded the desired first-order reduction g , and Equation (2.22) holds as desired. \square

Exercise 2.23 Construct sets S and T such that S is polynomial-time, Turing reducible to T , but S is not polynomial-time, many-one reducible to T , i.e., $S \leq_p^T T$, $S \not\leq_p T$.

[Hint: construct S and T systematically so that for all w , $w \in S$ iff $(0w \in T) \oplus (1w \in T)$ and systematically satisfy the conditions,

$$C_i \equiv \left(\begin{array}{l} \text{Turing machine } M_i \text{ running in time } i \cdot n^i \text{ does} \\ \text{not compute a many-one reduction from } S \text{ to } T. \end{array} \right)$$

If you are careful, you can carry out this construction so that S and T are in $\text{ETIME} = \text{DTIME}[2^{O(n)}]$. \square

2.4 Alternation

The concept of a nondeterministic acceptor of a boolean query has a long and rich history, going back to various kinds of nondeterministic automata. On the other

hand, it is important to remember that these are fictitious machines which cannot be built. We elaborate on the unreasonableness of nondeterministic machines.

For $A \subseteq \text{STRUC}[\tau]$ a boolean query, define its complement $\overline{A} = \text{STRUC}[\tau] - A$. Given a complexity class \mathcal{C} , one can define the complementary class as follows,

$$\text{co-}\mathcal{C} = \{\overline{A} \mid A \in \mathcal{C}\}$$

For example, since SAT is in NP, its complementary problem $\overline{\text{SAT}} = \text{UN-SAT}$ is in co-NP. The question whether NP is closed under complementation, i.e., is NP equal to co-NP? is open. Most people believe that these classes are different. Notice that if one could really build an NP machine, then one could also build a co-NP machine: All that is needed is a single gate to invert the former machine's answer. Thus from a very practical point of view, the complexity of a problem A and its complement, \overline{A} are identical.

One way to imagine a realization of an NP machine is via a parallel or biological machine with many processors. At each step, each processor p_i creates two copies of itself and sets them to work on two slightly different problems. If either of these offspring ever accepts, i.e., says “yes” to p_i , then p_i in turn says “yes” to its parent. These “yes”es travel up a binary tree to the root and the whole nondeterministic process accepts. In such a view of nondeterminism, in time $t(n)$ we can build about $2^{t(n)}$ processors. However, we do not make very good use of them. Their pattern of communication is very weak. Each processor can compute only the “or” of its children. Thus, the whole computation is one big “or” of its leaves.

There is a natural way to generalize the concept of nondeterminism so that it is closed under complementation and makes better use of its processors. Namely, we can let the processors compute either the “or” or the “and” of their children. This leads to the notion of *alternation*. An alternating Turing machine has both “or” states like a nondeterministic Turing machine and “and” states.

We now formally define and study alternating machines. We will see that in many ways the concept of alternation is more robust than the concept of nondeterminism.

Definition 2.24 An *alternating Turing machine* is a Turing machine whose states are divided into two groups: the existential states and the universal states. Recall that a *configuration* of any Turing machine — also called an *instantaneous description* (ID) — consists of the machine's state, work-tape contents, and head positions. The notion of when such a machine accepts an input is defined by induction: The alternating Turing machine in a given configuration c *accepts* iff

1. c is in a final accepting state, or
2. c is in an existential state and there exists a next configuration c' that accepts, or
3. c is in a universal state, there is at least one next configuration, and all next configurations accept.

Note that this is a generalization of the notion of acceptance for a nondeterministic Turing machine, which is an alternating Turing machine all of whose states are existential.

Turing machines have an awkward way of accessing their tapes, which makes it difficult for them to do anything in sublinear time. Since alternating Turing machines can sensibly use sublinear time, it is more reasonable to use machines that have a more random access nature. As a compromise, from now on we assume that our Turing machines have a *random access* read-only input. This works as follows: there is an *index tape*, which can be written and read like other tapes. Whenever the value v , written in binary, appears on the index tape, the read head automatically scans bit v of the input. \square

Define the complexity classes $\text{ASPACE}[s(n)]$ and $\text{ATIME}[t(n)]$ to be the set of boolean queries accepted by alternating Turing machines using a maximum of $O(s(n))$ tape cells, respectively a maximum of $O(t(n))$ time steps in any computation path on an input of length n . The main relationships between alternating and deterministic complexity are given by the following theorem.

Theorem 2.25 For $s(n) \geq \log n$, and for $t(n) \geq n$,

$$\begin{aligned} \bigcup_{k=1}^{\infty} \text{ATIME}[(t(n))^k] &= \bigcup_{k=1}^{\infty} \text{DSPACE}[(t(n))^k] \\ \text{ASPACE}[s(n)] &= \bigcup_{k=1}^{\infty} \text{DTIME}[k^{s(n)}] \end{aligned}$$

In particular, $\text{ASPACE}[\log n] = P$, and alternating polynomial time is equal to PSPACE.

Figure 2.26 shows the computation graph of an alternating machine. We assume for convenience that such machines have a unique accepting and a unique rejecting configuration and that each configuration has at most two possible next




Figure 2.26: Computation graph of an alternating Turing machine

moves. The start configuration is labeled “ s ” and the accept configuration is labeled “ t ”. We also assume that these machines have clocks that uniformly cause the machines to shut off at a fixed time that is a function of the length of the input. Shutting off means entering the reject configuration unless the machine is already in the accept configuration.

Observe Figure 2.26. The letters “E” and “A” below the vertices indicate whether the corresponding configurations are existential or universal. If they were all existential, then this would be a nondeterministic computation. The time $t(n)$ measures the depth of the computation graph. It is convenient to think of alternating Turing machines as a parallel model in which at each branching move an extra processor is created, and these two processors take the two branches. Eventually these two processors complete their tasks and report their answers to their parent. If the parent was existential, then it reports “accept” iff either of its children accept. If the parent is universal, then it reports “accept” iff both of its children accept. Notice that in time $t(n)$ potentially $2^{O(t(n))}$ processors are created. For purely nondeterministic machines these processors have a very poor system of communication: each parent can perform only the “or” of its children. The ability to perform

“and”s as well as “or”s lets alternating machines make more extensive use of their extra processors. We will see in Chapter 4 that alternating time $t(n)$ corresponds to a reasonable notion of parallel time $t(n)$ when $2^{O(t(n))}$ processors are available. As we will see, we can allow at most polynomially many processors by restricting the alternating machines to use no more than logarithmic space. The space used by an alternating machine is the maximum amount of space used in any path through its computation graph.

Before we prove some simulation results concerning alternating Turing machines, we give some examples of their use. The first example involves the circuit value problem (CVP). We will see in Exercise 3.28 that CVP is complete for P.

Definition 2.27 A *boolean circuit* is a directed, acyclic graph (DAG),

$$C = (V, E, G_{\wedge}, G_{\vee}, G_{\neg}, I, r) \in \text{STRUC}[\tau_c]; \quad \tau_c = \langle E^2, G_{\wedge}^1, G_{\vee}^1, G_{\neg}^1, I^1, r \rangle$$

Internal node w is an and-gate if $G_{\wedge}(w)$ holds, an or-gate if $G_{\vee}(w)$ holds, and a not-gate if $G_{\neg}(w)$ holds. The nodes v with no edges entering them are called leaves, and the input relation $I(v)$ represents the fact that the leaf v is on. Often we will be given a circuit C , and separately we will be given its input relation I .

Define the Circuit Value Problem (CVP) to consist of those circuits whose root gate r evaluates to one. Define the monotone, circuit value problem (MCVP) to be the subset of CVP in which no negation gates occur. \square

Exercise 2.28 Show that a boolean circuit can be evaluated in linear time.

[Hint: do this bottom up: from the leaves to the root. Each edge should be processed only once. By “linear time” we mean time $O(n + m)$ on a random-access machine (RAM). We assume that the edges of the input circuit are given as an adjacency list; m is the number of edges. Such machines correspond better to real computers than do multi-tape Turing machines. There is a polynomial-size memory and an $O(\log n)$ -bit word size. Here we assume that the edges of the input circuit are given as an adjacency list, m is the number of edges.] \square

Proposition 2.29 MCVP is recognizable in $\text{ASPACE}[\log n]$.

Proof Let G be a monotone boolean circuit as in Definition 2.27. Define the procedure “EVAL(a)”, where a is a vertex of G , as follows:

1. **if** $I(a)$ **then** accept

2. **else if** a has no outgoing edges **then** reject
3. **if** $G_\wedge(a)$ **then** in a universal state choose a child b of a
4. **else** in an existential state choose a child b of a
5. Return (EVAL(b))

The machine M simply calls EVAL(r). Observe that EVAL(a) returns “accept” iff gate a evaluates to one. Furthermore, the space used by EVAL is just the space to name two vertices a, b . Thus, M is an ASPACE[$\log n$] machine accepting MCVP, as desired. Observe that the alternating time required for this computation is the depth of circuit G — the length of the longest path in G starting at r . Recall that we have said that all alternating machines have timers. In this case, an appropriate time limit would be $n = \|G\|$, which is an upper bound on the length of the longest path. \square

Another boolean query that is well suited for alternating computation is the quantified satisfiability problem:

Definition 2.30 The *quantified satisfiability problem* (QSAT) is the set of true formulas of the following form:

$$\Psi = (Q_1 x_1)(Q_2 x_2) \cdots (Q_r x_r) \varphi$$

where φ is a boolean formula and each Q_i 's is each either \forall or \exists , and x_1, \dots, x_r are the boolean variables occurring in φ . \square

Observe that for any boolean formula φ on variables \bar{x} ,

$$\varphi \in \text{SAT} \iff (\exists \bar{x}) \varphi \in \text{QSAT} \quad \text{and} \quad \varphi \notin \text{SAT} \iff (\forall \bar{x}) \neg \varphi \in \text{QSAT}$$

Thus QSAT logically contains SAT and $\overline{\text{SAT}}$.

Proposition 2.31 QSAT is recognizable in ATIME[n].

Proof Construct an alternating machine A that works as follows. To evaluate the sentence

$$\Phi \equiv (\exists x_1)(\forall x_2) \cdots (Q_r x_r) \alpha(\bar{x})$$

in an existential state, A writes down a boolean value for x_1 , in a universal state it writes a bit for x_2 , and so on. Next A must evaluate the quantifier-free boolean formula α on these chosen values. This is especially easy for an alternating machine:

for each “ \wedge ” in α , A universally chooses which side to evaluate and for each “ \vee ”, A existentially chooses. Thus A only has to read one of the chosen bits x_i and accept iff it is true and occurs positively, or false and occurs negatively. Observe that A runs in linear time and accepts the sentence Φ iff Φ is true. \square

The next theorem explains the relationship between alternating time and deterministic and nondeterministic space.

Theorem 2.32 *Let $s(n) \geq \log n$ be space constructible. Then,*

$$\text{NSPACE}[s(n)] \subseteq \text{ATIME}[s(n)^2] \subseteq \text{DSPACE}[s(n)^2]$$

Proof For the first inclusion, let N be an $\text{NSPACE}[s(n)]$ Turing machine. Let w be an input to N . Let G_w denote the computation graph of N on input w . Note that N accepts w iff there is a path from s to t in G_w . We construct an $\text{ATIME}[s(n)^2]$ machine A that accepts the same language as N . A does this by calling the subroutine, $P(d, x, y)$, which accepts iff there is a path in G_w of length at most 2^d from x to y . For $d > 0$, P is defined as follows:

$$P(d, x, y) \equiv (\exists z)(P(d-1, x, z) \wedge P(d-1, z, y))$$

P works by existentially choosing a middle configuration z , universally choosing the first half or the second half, and then checking that the appropriate path of length 2^{d-1} exists. Thus, the time $T(d)$ taken to compute $P(d, x, y)$ is the time to write down a new, middle configuration plus the time to compute $P(d-1, x', y')$. The number of bits in a configuration of G_w is $O(s(n))$ where $n = |w|$. Thus,

$$T(d) = O(s(n)) + T(d-1) = O(d \cdot s(n))$$

The length of the maximum useful path in G_w is bounded by the number of configurations of N on input w , i.e. $2^{cs(n)}$ for an appropriate value of c . Thus, on input w , A calls $P(cs(n), s, t)$ and receives its answer in time $O(cs(n)s(n)) = O(s(n)^2)$, as desired.

For the second inclusion, let A be an $\text{ATIME}[t(n)]$ machine. On input w , A 's computation graph — pictured in Figure 2.26 — has depth $O(t(n))$ and size $2^{O(t(n))}$. A deterministic Turing machine can systematically search this entire and-or graph using space $O(t(n))$. This is done by keeping a string of length $O(t(n))$: $c_1 c_2 \dots c_r \star \dots \star$ denoting that we are currently simulating step r of A 's computation, having made choices $c_1 \dots c_r$ on all of the existential and universal branches up until this point. The rest of the simulation will report an **answer** as to whether choices $c_1 \dots c_r$ will lead to acceptance. This is done as follows:

If one of the following conditions holds:

1. $c_r = 1$, or
2. **answer** = “yes” and step r was existential, or
3. **answer** = “no” and step the r was universal,

then let $c_r = \star$ and report **answer** back to step $r - 1$. Otherwise, set $c_r = 1$ and continue. Note, that we do not have to store intermediate configurations of the simulation because the sequence $c_1 c_2 \dots c_r \star \dots \star$ uniquely determines which configuration of A to go to next. \square

An immediate corollary is that $\text{NSPACE}[s(n)]$ is contained in $\text{DSpace}[s(n)^2]$. This is Savitch’s theorem (Theorem 2.32), and it is the best known simulation of nondeterministic space by deterministic space. It is unknown whether equality holds in either or both of the inclusions of Theorem 2.32. Another corollary of Theorem 2.32 is the first part of Theorem 2.25.

We complete the proof of Theorem 2.25 by showing that $\text{ASPACE}[s(n)]$ is $\text{DTIME}[O(1)^{s(n)}]$. One direction of this is obvious: an $\text{ASPACE}[s(n)]$ machine has $O(1)^{s(n)}$ possible configurations. Thus, its entire computation graph is of size $O(1)^{s(n)}$ and thus may be traversed in $\text{DTIME}[O(1)^{s(n)}]$. The same traversal algorithm as in the second half of the proof of Theorem 2.32 works in this case.

In the other direction, we are given a $\text{DTIME}[k^{s(n)}]$ machine M . Let w be an input to M and let $n = |w|$. We can view M ’s computation as a $k^{s(n)} \times k^{s(n)}$ table — see Figure 2.34. Cell (t, p) of this table contains the symbol that is in position p of M ’s tape at time t of the computation. Furthermore, if M ’s head was at position p at time t , then this cell should also include M ’s state at time t .

Define an alternating procedure $C(t, p, a)$ that accepts iff the contents of cell p at time t in M ’s computation on input w is symbol a . $C(0, p, a)$ holds iff a is the correct symbol in position p of M ’s initial configuration on input w . This means that position 1 contains $\langle q_0, w_1 \rangle$ where q_0 is M ’s start state, and w_1 is the first symbol of w . Similarly, for $2 \leq p \leq n$, $C(0, p, a)$ holds iff $a = w_p$.

Inductively, $C(t + 1, p, a)$ holds iff the three symbols a_{-1}, a_0, a_1 in tape positions $p - 1, p, p + 1$ lead to an “a” in position p in one step of M ’s computation. We denote this symbolically as $(a_{-1}, a_0, a_1) \xrightarrow{M} a$. This condition can be read directly from M ’s transition table.

$$C(t + 1, p, a) \equiv (\exists a_{-1}, a_0, a_1) \left((a_{-1}, a_0, a_1) \xrightarrow{M} a \wedge (\forall i \in \{-1, 0, 1\})(C(t, p + i, a_i)) \right) \quad (2.33)$$

		Space											
		1	2	p	n	b	$T(n)$						
Time	0	$\langle q_0, w_1 \rangle$	w_2	\dots	w_n	b	b						
	1	w_1	$\langle q_1, w_2 \rangle$	\dots	w_n	b	b						
		\vdots	\vdots	\vdots			\vdots						
	t			<table border="1" style="margin: auto;"> <tr> <td>a_{-1}</td> <td>a_0</td> <td>a_1</td> </tr> <tr> <td></td> <td>a</td> <td></td> </tr> </table>	a_{-1}	a_0	a_1		a				
a_{-1}	a_0	a_1											
	a												
	$t + 1$												
		\vdots	\vdots	\vdots			\vdots						
	$T(n)$	$\langle q_f, 0 \rangle$	\dots	\dots		\dots							

Figure 2.34: A $\text{DTIME}[T(n)]$ computation on input $w_1 w_2 \dots w_n$

See Figure 2.34 which shows values of $C(t, p, \star)$ for a $\text{DTIME}[T(n)]$ Turing machine. In the present case, $T(n) = k^{s(n)}$.

Formula 2.33 can be evaluated by an alternating machine using the space to hold the values of t and p . This is $O(\log k^{s(n)}) = O(s(n))$. Note that M accepts w iff $C(k^{s(n)}, 1, a_f)$ holds, where a_f is the contents of the first cell of M 's accept configuration. For example, we can use $a_f = \langle q_f, 0 \rangle$, where q_f is M 's accept state.

This completes the proof of Theorem 2.25.

2.5 Simultaneous Resource Classes

Let the classes $\text{ASPACE-TIME}[s(n), t(n)]$ (resp. $\text{ATIME-ALT}[t(n), a(n)]$) be the sets of boolean queries accepted by alternating machines simultaneously using space $s(n)$ and time $t(n)$ (resp. time $t(n)$ and making at most $a(n)$ alternations between existential and universal states, and starting with existential. Thus $\text{ATIME-ALT}[n^{O(1)}, 1] = \text{NP}$). Two more important complexity classes may now be defined using these simultaneous alternating classes. Define the polynomial-time hierarchy (PH) to be the set of boolean queries accepted in polynomial time by alternating Turing machines making a bounded number of alternations between existential and universal states:

$$\text{PH} = \bigcup_{k=1}^{\infty} \text{ATIME-ALT}[n^k, k]. \tag{2.35}$$

Also define NC (Nick's Class) to be the set of boolean queries accepted by alternating Turing machines in $\log n$ space and poly log time:

$$\text{NC} = \bigcup_{k=1}^{\infty} \text{ASPACE-TIME}[\log n, \log^k n] \quad (2.36)$$

See Theorem 5.33 for the more usual definition of NC as the class of boolean queries accepted by a parallel random access machine using polynomially much hardware in poly log parallel time.

2.6 Summary

We conclude this section with a list of the complexity classes defined so far. These will be a main focus for much of what follows:

$$L \subseteq NL \subseteq NC \subseteq P \subseteq NP \subseteq PH \subseteq PSPACE \quad (2.37)$$

The above containments are easy to prove. (It is a good exercise for the reader to now show how to simulate each of these complexity classes by the next larger one.)

On the other hand, despite intense effort, there is very little known about the strictness of the above inclusions. It has not yet been proved that L is not equal to PH, or that P is not equal to PSPACE. It would probably be safe to assume the strictness of each of the inequalities in Equation (2.37) as an axiom and go on with the rest of one's life.

The fact that we cannot prove these inequalities reveals just the tip of the iceberg of what we do not know concerning the computational complexity of important computational problems. As an example, for the thousands of known NP complete problems, the best known algorithms to get an exact solution are all exponential time in the worst case. However, we have no proof that these are not computable in linear time.

See Figure 2.38 for a view of the computability and complexity world. The classes in the diagram that have not yet been defined will be described later in the text. The intuitive idea behind this diagram is that there is a set of boolean queries called "truly feasible". These are the queries that can be computed exactly with an "affordable" amount of time and hardware, on all "reasonably sized" instances.

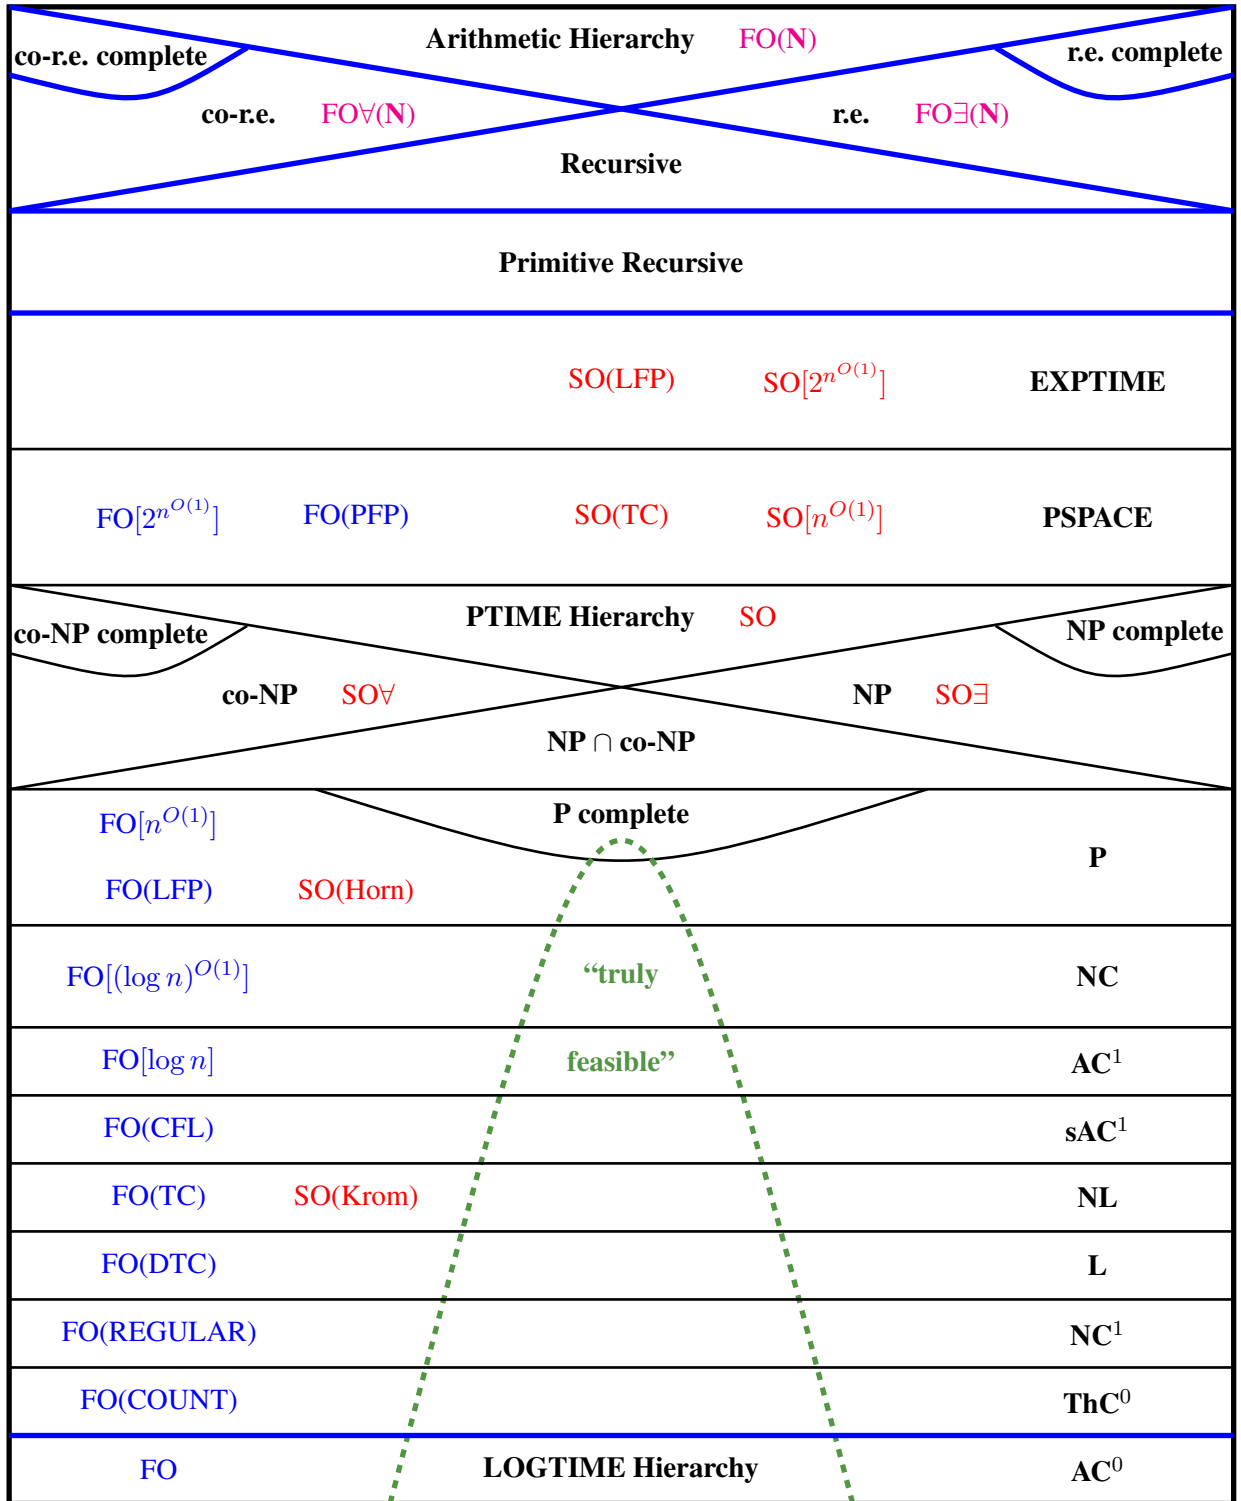


Figure 2.38: The World of Descriptive and Computational Complexity.

The truly feasible queries are a proper subset of P .² Many important problems that we need to solve are not truly feasible. For example, they may be NP-complete, EXPTIME-complete, or even r.e.-complete. The theory of algorithms and complexity helps us determine whether the problem we need to solve is feasible, and if not, how to choose a limited set of instances of the problem — or easier versions of them — that are feasible.

As the reader will learn, the theory of complexity via Turing machines is isomorphic to descriptive complexity, i.e., the theory of complexity via logic formulas. We will give descriptive characterizations of almost all of the classes in Figure 2.38. For example, the logarithmic-time hierarchy is equal to the set of first-order boolean queries ($LH = FO$, Theorem 5.30). The polynomial-time hierarchy is the set of second-order boolean queries ($PH = SO$, Theorem 7.21). The arithmetic hierarchy is an analogous but much larger class defined to be the set of boolean queries that are describable in the first-order theory of the natural numbers.

The descriptive characterizations of complexity classes in this book are all constructive. Efficient algorithms can be automatically translated into efficient descriptions and vice versa. The descriptive approach has added significant new insights and brought new methods to bear on the basic problems in complexity. Descriptive complexity affords an elegant and simple view of complexity. We hope that the reader will use this and perhaps other approaches to add a few more pieces to the puzzle.

Historical Notes and Suggestions for Further Reading

The reader is referred to [Pap94] for much more information on computational complexity. See also [BDG88] for more detail on “structural complexity” and [Str94] for a combination of algebraic, logical, and automata-theoretic approaches to complexity theory.

Chandra and Harel [CH80a] introduced the general notion of queries as in Definition 1.24. We got the idea for using queries as the general paradigm of computation from a lecture by Phokion Kolaitis.

Alternating machines were defined and Theorem 2.25 was proved independently by Kozen and by Chandra and Stockmeyer [CKS81]. Kozen titled his description of these machines, “Parallelism in Turing Machines”, which is very apt

²Those readers who believe that the class BPP properly contains P should change this sentence to, “. . . a proper subset of BPP.”

(cf. Theorem 5.33); but Chandra and Stockmeyer coined the name “alternating Turing machine”.

The QSAT boolean query was so named in [Pap94] because it is a natural generalization of SAT. Previously, it had been called the quantified boolean formula problem (QBF).

Many-one reductions have their name for an historical reason: to contrast them with 1:1 reductions in which the corresponding mapping is a 1:1 function.

Exercise 2.16 was suggested by Jose Balcázar, cf. [CSV84].

Nick’s class (NC) was originally defined by Nick Pippenger; see [Coo85].

Bibliography

- [AV91] S. Abiteboul and V. Vianu, “Generic Computation And Its Complexity,” *32nd IEEE Symposium on FOCS* (1991), 209-219.
- [AHV95] S. Abiteboul, R. Hull, and V. Vianu, *Foundations of Databases*, 1995, Addison-Wesley.
- [AVV97] S. Abiteboul, M.Y. Vardi, and V. Vianu, “Fixpoint Logics, Relational Machines, and Computational Complexity,” *J. Assoc. Comput. Mach.* 44(1) (1997), 30 – 56.
- [Agr01] M. Agrawal, “The First-Order Isomorphism Theorem,” *Foundations of Software Technology and Theoretical computer Science*, (2001), 58-69.
- [AAI97] M. Agrawal, E. Allender, R. Impagliazzo, T. Pitassi and S. Rudich, “Reducing the Complexity of Reductions,” *ACM Symp. Theory Of Comput.* (1997), 730–738.
- [AHU74] A.V. Aho, J.E. Hopcroft and J.D. Ullman, *The Design and Analysis of Computer Algorithms*, Addison- Wesley (1974).
- [Ajt83] M. Ajtai, “ Σ_1^1 Formulae on Finite Structures,” *Annals of Pure and Applied Logic* 24 (1983), 1-48.
- [Ajt87] M. Ajtai, “Recursive Construction for 3-Regular Expanders,” *28th IEEE Symp. on Foundations of Computer Science* (1987), 295-304.
- [Ajt89] M. Ajtai, “First-Order Definability on Finite Structures,” *Annals of Pure and Applied Logic* (1989), 211-225.
- [AF90] M. Ajtai and R. Fagin, “Reachability is Harder for Directed than for Undirected Graphs,” *J. Symb. Logic*, 55 (1990), 113-150.

- [AFS97] M. Ajtai, R. Fagin, and L. Stockmeyer, "The Closure of Monadic NP," IBM Research Report RJ 10092 (1997). The closure of monadic NP, with Miklos Ajtai and Larry Stockmeyer. *J. Computer and System Sciences* 60, June 2000, pp. 660-716. (Special issue for selected papers from the 1998 ACM Symp. on Theory of Computing).
- [AG87] M. Ajtai and Y. Gurevich, "Monotone versus Positive," *J. Assoc. Comput. Mach.* 34 (1987), 1004–1015.
- [AKL79] R. Aleliunas, R. Karp, R. Lipton, L. Lovász, and C. Rackoff, "Random Walks, Universal Traversal Sequences, and the Complexity of Maze Problems," *IEEE Found. of Comp. Sci. Symp.* (1979), 218-233.
- [ABI97] E. Allender, N. Immerman, J. Balcázar, "A First-Order Isomorphism Theorem," *SIAM J. Comput.* 26(2) (1997), 557-567.
- [ACF94] B. Alpern, L. Carter, E. Feig, and T. Selker, "The Uniform Memory Hierarchy Model of Computation," *Algorithmica*, 12(2-3) (1994), 72–109.
- [AG94] G. Almasi and A. Gottlieb, *Highly Parallel Computing (Second Edition)* 1994, Benjamin-Cummings.
- [ASE92] N. Alon, J. Spencer, and P. Erdős, *The Probabilistic Method*, 1992, John Wiley and Sons, Inc.
- [ADN95] H. Andréka, I. Düntsch, and I. Németi, "Expressibility of Properties of Relations," *J. Symbolic Logic* 60(3) (1995), 970 - 991.
- [AF97] S. Arora and R. Fagin, "On Winning Strategies in Ehrenfeucht-Fraïssé Games," *Theoret. Comp. Sci.* 174(1-2) (1997), 97-121.
- [ALMSS] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy, "Proof Verification and the Hardness of Approximation Problems," *J. Assoc. Comput. Mach.* 45(3) (1998), 501-555.
- [AS92] S. Arora and S. Safra, "Probabilistic Checking of Proofs: a New Characterization of NP," *J. Assoc. Comput. Mach.* 45(1) (1998), 70-122. A preliminary version appeared in *IEEE Found. of Comp. Sci. Symp.* (1992), 2-13.
- [Ba81] L. Babai, "Moderately Exponential Bound for Graph Isomorphism," in *Proc. Int. Conf. Fundamentals of Computation Theory* (1981), Springer LNCS, 34 – 50.

- [BK80] L. Babai and L. Kučera, Canonical Labelling of Graphs in Linear Average Time,” *20th IEEE Symp. on Foundations of Computer Science* (1980), 39-46.
- [BL83] L. Babai and E. Luks, “Canonical Labelling of Graphs,” *15th ACM STOC Symp.*, (1983), 171-183.
- [BDG88] J. Balcázar, J. Días, and J. Gabarró, *Structural Complexity*, Vols. I and II, EATCS Monographs on Theoretical Computer Science, 1988, Springer-Verlag.
- [BI97] D.M. Barrington and N. Immerman, “Time, Hardware, and Uniformity,” in *Complexity Theory Retrospective II*, L. Hemaspaandra and A. Selman, editors, 1997, Springer-Verlag, 1-22.
- [BIS88] D.M. Barrington, N. Immerman, and H. Straubing, “On Uniformity Within NC^1 ,” *Third Annual Structure in Complexity Theory Symp.* (1988), 47-59.
- [BBI97] D.M. Barrington, J. Buss, and N. Immerman, “Capturing Deterministic Space via Number of Variables,” in preparation.
- [Bar77] J. Barwise, “On Moschovakis Closure Ordinals,” *J. Symb. Logic* 42 (1977), 292-296.
- [Bea86] P. Beame, “Limits on the Power of Concurrent-Write Parallel Machines,” *18th ACM STOC* (1986), 169-176.
- [Bea96] P. Beame, “A Switching Lemma Primer,” manuscript, <http://www.cs.washington.edu/homes/beame/papers.html>.
- [Ben62] J. Bennett, “On Spectra” (1962), Ph.D. thesis, Princeton University.
- [BGK85] A. Blass, Y. Gurevich and D. Kozen, “A Zero–One Law for Logic With a Fixed Point Operator,” *Information and Computation* 67 (1985), 70-90.
- [Bol82] Béla Bollobás, *Random Graphs*, Academic Press (1982).
- [BH92] R. Boppana and M. Halldórsson, “Approximating Maximum Independent Sets by Excluding Subgraphs,” *BIT* 32(2) (1992), 180-196.
- [BS90] R. Boppana and M. Sipser, “The Complexity of Finite Functions,” in *Handbook of Theoretical Computer Science, Vol. A* 1990, Jan van Leeuwen, ed., Elsevier, Amsterdam and M.I.T. Press, Cambridge, MA.

- [B82] E. Börger, “Decision Problems in Predicate Logic,” in *Logic Colloquium '82*, G. Lolli, G. Longo and A. Marcia (editors) North-Holland, 1984, 263 – 301.
- [BCD88] A. Borodin, S.A. Cook, P.W. Dymond, W.L. Ruzzo, and M. Tompa, “Two Applications of Complementation via Inductive Counting,” *Third Annual Structure in Complexity Theory Symp.* (1988), 116-125.
- [BCH86] P. Beame, S. Cook, H.J. Hoover, “Log Depth Circuits for Division and Related Problems,” *SIAM J. Comput.* 15:4 (1986), 994-1003.
- [BCP83] A. Borodin, S. Cook, and N. Pippenger, “Parallel Computation for Well-Endowed Rings and Space-Bounded Probabilistic Machines,” *Information and Control*, 58 (1983), 113-136.
- [Bra96] J. Bradfield, “On the Expressivity of the Modal Mu-Calculus,” *Symp. Theoretical Aspects Comp. Sci.* (1996).
- [Büc60] R. Büchi, “Weak Second-Order Arithmetic and Finite Automata,” *Zeit. Math. Logik. Grund. Math.* 6 (1960), 66-92.
- [CFI92] J.-Y. Cai, M. Fürer, N. Immerman, “An Optimal Lower Bound on the Number of Variables for Graph Identification,” *Combinatorica* 12 (4) (1992) 389-410.
- [CH80a] Ashok Chandra and David Harel, “Computable Queries for Relational Databases,” *JCSS* 21(2) (1980), 156-178.
- [CH80b] Ashok Chandra and David Harel, “Structure and Complexity of Relational Queries,” *IEEE Found. of Comp. Sci. Symp.* (1980), 333-347. Also appeared in a revised as [CH82]
- [CH82] A. Chandra and D. Harel, “Structure and Complexity of Relational Queries,” *JCSS* 25 (1982), 99-128.
- [CKS81] A. Chandra, D. Kozen, and L. Stockmeyer, “Alternation,” *JACM*, 28, No. 1, (1981), 114-133.
- [CSV84] A. Chandra, L. Stockmeyer and U. Vishkin, “Constant Depth Reducibility,” *SIAM J. of Comp.* 13, No. 2, 1984, (423-439).
- [CE81] E. Clarke and E.A. Emerson, “Design and Synthesis of Synchronization Skeletons Using Branching Time Temporal Logic,” in *Proc. Workshop on Logic of Programs*, LNCS 131, 1981, Springer-Verlag, 52-71.

- [Coh66] P. Cohen, *Set Theory and the Continuum Hypothesis*, 1966, Benjamin.
- [Co88] K. Compton, "0-1 laws in logic and combinatorics," in *NATO Adv. Study Inst. on Algorithms and Order*, I. Rival, editor, 1988, D. Reidel, 353–383.
- [Coo71] S. Cook, "The Complexity of Theorem Proving Procedures," *Proc. Third Annual ACM STOC Symp.* (1971), 151-158.
- [Coo85] S. Cook, "A Taxonomy of Problems with Fast Parallel Algorithms," *Information and Control* 64 (1985), 2-22.
- [Cop94] D. Coppersmith, "A Left Coset Composed of n -cycles," Research Report RC19511 IBM (1994).
- [Cou90] B. Courcelle, "The Monadic Second-Order Logic of GraphsI: Recognizable Sets of Finite Graphs," *Information and Computation* 85 (1990), 12 - 75.
- [Cou97] B. Courcelle, "On the Expression of Graph Properties in Some Fragments of Monadic Second-Order Logic," in *Descriptive Complexity and Finite Models*, N. Immerman and Ph. Kolaitis, eds., 1997, American Mathematical Society, 33 - 62.
- [Da84] E. Dahlhaus, "Reduction to NP-Complete Problems by Interpretations," in *Logic and Machines: Decision Problems and Complexity*, Börger, Rödding, and Hasenjaeger eds., Lecture Notes In Computer Science 171, Springer-Verlag (1984), 357-365.
- [Daw93] A. Dawar, "Feasible Computation Through Model Theory," PhD Dissertation, University of Pennsylvania (1993).
- [DGH98] A. Dawar, G. Gottlob, L. Hella, "Capturing Relativized Complexity Classes without Order," to appear in *Mathematical Logic Quarterly*.
- [DH95] Anuj Dawar and Lauri Hella, "The Expressive Power of Finitely Many Generalized Quantifiers," *Information and Computation* 123(2) (1995), 172-184.
- [DDLW98] A. Dawar, K. Doets, S. Lindell, and S. Weinstein, "Elementary Properties of the Finite Ranks," *Mathematical Logic Quarterly* 44 (1998), 349-353.

- [DLW95] A. Dawar, S. Lindell, and S. Weinstein, "Infinitary logic and inductive definability over finite structures," *Information and Computation*, 119 (1995), 160-175.
- [DGS86] L. Denenberg, Y. Gurevich and S. Shelah, "Definability by Constant-Depth Polynomial-Size Circuits", *Information and Control* 70 (1986), 216-240.
- [deR84] M. de Rougemont, "Uniform Definability on Finite Structures with Successor," *16th ACM STOC Symp.*, (1984), 409-417.
- [DL98] W. Diffie and S. Landau, *Privacy on the Line: the Politics of Wiretapping and Encryption*, MIT Press, 1998.
- [DS95] G. Dong, J. Su, "Space-Bounded FOIES," *ACM Symp. Principles Database Systems* (1995), 139 -150.
- [DS93] G. Dong, J. Su, "Incremental and Decremental Evaluation of Transitive Closure by First-Order Queries," *Information and Computation*, 120(1) (1995), 101-106. Preliminary results presented at the 1993 Australian Computer Science Conference.
- [EF95] H.-D. Ebbinghaus, J. Flum, *Finite Model Theory* 1995, Springer 1995.
- [EFT94] H.-D. Ebbinghaus, J. Flum, and W. Thomas, *Mathematical Logic*, 2nd edition 1994, Springer-Verlag.
- [Ehr61] A. Ehrenfeucht, "An Application of Games to the Completeness Problem for Formalized Theories," *Fund. Math.* 49 (1961), 129-141.
- [End72] H. Enderton, *A Mathematical Introduction to Logic*, Academic Press, 1972.
- [ES74] P. Erdős and J. Spencer, *Probabilistic Methods in Combinatorics*, 1974, Academic Press.
- [Ete95] K. Etessami, "Counting Quantifiers, Successor Relations, and Logarithmic Space," *IEEE Structure in Complexity Theory Symp.* (1995), 2-11.
- [Ete95a] K. Etessami, "Ordering and Descriptive Complexity" Ph.D. thesis, 1995, UMass, Amherst.
- [EI95] K. Etessami and N. Immerman, "Reachability and the Power of Local Ordering," *Theoret. Comp. Sci.* 148(2) (1995), 261-279.

- [EI95a] K. Etessami and N. Immerman, "Tree Canonization and Transitive Closure," to appear in *Information and Computation*. A preliminary version appeared in *IEEE Symp. Logic In Comput. Sci.* (1995), 331-341.
- [Fag73] R. Fagin, "Contributions to the Model Theory of Finite Structures", Ph.D. Thesis (1973), U. C. Berkeley.
- [Fag74] R. Fagin, "Generalized First-Order Spectra and Polynomial-Time Recognizable Sets," in *Complexity of Computation*, (ed. R. Karp), *SIAM-AMS Proc.* 7 (1974), 43-73.
- [Fag75] R. Fagin, "Monadic generalized spectra," *Zeitschr. f. math. Logik und Grundlagen d. Math.* 21 (1975), 89-96.
- [Fag76] R. Fagin, "Probabilities on Finite Models," *J. Symbol. Logic* 41, No. 1 (1976), 50-58.
- [Fag93] R. Fagin, "Finite-Model Theory – a Personal Perspective," *Theoret. Comp. Sci.* 116 (1993), 3-31.
- [Fag97] R. Fagin, "Easier Ways to Win Logical Games," in *Descriptive Complexity and Finite Models*, N. Immerman and Ph. Kolaitis, eds., 1997, American Mathematical Society, 1 - 32.
- [FSV95] R. Fagin, L. Stockmeyer, and M.Y. Vardi, "On monadic NP vs. monadic co-NP," *Information and Computation* 120(1) (1995), 78-92.
- [FV98] T. Feder and M.Y. Vardi, "The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory," (1998).
- [Fel50] W. Feller, *An Introduction to Probability Theory and Its Applications*, Vol. 1, 1950, John Wiley, New York.
- [FRW84] F. Fich, Prabhakar Ragde, and, Avi Wigderson (1984), "Relations Between Concurrent-Write Models of Parallel Computation," *Third ACM Symp. on Principles of Distributed Computing*, 179-189.
- [FW78] S. Fortune and J. Wyllie, "Parallelism in Random Access Machines," *ACM Symp. Theory Of Comput.* (1978), 114-118.
- [Fra54] R. Fraïssé, "Sur les Classifications des Systems de Relations," *Publ. Sci. Univ. Alger I* (1954).

- [Fur87] M. Fürer, “A Counterexample In Graph Isomorphism Testing – Extended Abstract,” manuscript (October, 1987).
- [FSS84] M. Furst, J.B. Saxe, and M. Sipser, “Parity, Circuits, and the Polynomial-Time Hierarchy,” *Math. Systems Theory*, 17 (1984), 13-27.
- [Gab81] D. Gabbay, “Expressive Functional Completeness in Tense Logic,” in: *Aspects of Philosophical Logic*, 1981, ed. Monnich, D. Reidel, Dordrecht, 91-117.
- [Gai81] H. Gaifman, “On Local and Non-Local Properties,” *Proc. Herbrand Logic Colloq.* (1981), 105-135.
- [GH97] F. Gire and H. Hoang, “An Extension of Fixpoint Logic with a Symmetry-Based Choice Construct,” *Information and Computation* 144(1) (1998), 40-65.
- [Göd30] Gödel, K., “Die Vollständigkeit der Axiome des Logischen Funktionenkalküls,” *Monatshefte für Mathematik und Physik* 37 (1930), 349 - 360, (English translation in [vH67]).
- [Go82] L. Goldschlager, ”A Universal Interconnection Pattern for Parallel Computers,” *JACM*, October 1982.
- [Go77] L. Goldschlager, “The Monotone and Planar Circuit Value Problems are Log Space Complete for P,” *SIGACT News* 9(2) (1977).
- [Grä92] E. Grädel, “Capturing Complexity Classes by Fragments of Second Order Logic,” *Theoret. Comp. Sci.* 101 (1992), 35-57.
- [Grä92a] E. Grädel, “On Transitive Closure Logic,” *Computer Science Logic* (1992), LNCS, Springer, 149–163.
- [GM95] E. Grädel and G. McColm, “On the Power of Deterministic Transitive Closures,” *Information and Computation* 119 (1995), 129-135.
- [GM96] E. Grädel and G. McColm, “Hierarchies in Transitive Closure Logic, Stratified Datalog and Infinitary Logic,” *Annals of Pure and Applied Logic* 77 (1996), 166–199.
- [GO93] E. Grädel and M. Otto, “Inductive Definability with Counting on Finite Structures,” *Computer Science Logic* 1993, LNCS 702, Springer, 231–247.

- [Gra84] E. Grandjean, "The Spectra of First-Order Sentences and Computational Complexity," *SIAM J. of Comp.* 13, No. 2 (1984), 356-373.
- [Gra85] E. Grandjean, "Universal quantifiers and time complexity of Random Access Machines," *Math. Syst. Th.* (1985), 171-187.
- [Gra89] E. Grandjean, "First-order spectra with one variable," to appear in *J. Comput. Syst. Sci.*
- [Gro95] M. Grohe, "Complete Problems for Fixed-Point Logics," *J. Symbolic Logic* 60 (1995), 517-527.
- [Gro96] M. Grohe, "Equivalence in Finite-Variable Logics is Complete for Polynomial Time," *IEEE Found. of Comp. Sci. Symp.* (1996).
- [Gro96a] M. Grohe, "Arity Hierarchies," *Annals of Pure and Applied Logic* 82 (1996), 103-163.
- [Gro97] M. Grohe, "Large Finite Structures With Few L^k -Types," *IEEE Symp. Logic In Comput. Sci.* (1997).
- [GS98] M. Grohe and T. Schwentick, "Locality of order-invariant first-order formulas. MFCS'98, 437-445. Full version to appear in ACM TOCL.
- [Gro97a] M. Grohe, "Canonization for L^k -Invariants is Hard," *Annual Conference of the European Association for Computer Science Logic* (1997), M. Nielsen and W. Thomas, eds., 185-200.
- [Gro97b] M. Grohe, "Fixed-Point Logics on Planar Graphs," manuscript (1997).
- [Gur83] Y. Gurevich, "Algebras of feasible functions," *IEEE Found. of Comp. Sci. Symp.* (1983), 210-214.
- [Gur84] Y. Gurevich, "Toward Logic Tailored for Computational Complexity," *Computation and Proof Theory* (M.M. Richter et. al., eds.). Springer-Verlag Lecture Notes in Math. 1104 (1984), 175-216.
- [Gur88] Y. Gurevich, "Logic and the Challenge of Computer Science," in *Current Trends in Theoretical Computer Science*, ed. E. Börger, Computer Science Press (1988), 1-57.
- [Gur91] Y. Gurevich, "Evolving Algebras: A Tutorial Introduction," *Bulletin of EATCS*, 43 (1991), 264-284.

- [Gur93] Y. Gurevich, "Evolving Algebras 1993: Lipari Guide", *Specification and Validation Methods*, ed. E. Börger, Oxford University Press, 1995, 9–36.
- [GS85] Y. Gurevich and S. Shelah, "Fixed-Point Extensions of First-Order Logic," *Annals of Pure and Applied Logic* 32 (1986), 265–280.
- [GS96] Y. Gurevich and S. Shelah, "On Finite Rigid Structures," *J. Symbolic Logic* 61(2) (1996), 549 - 562.
- [HP93] P. Hajek and P. Pudlak, *Metamathematics of First-Order Arithmetic*, 1993, Springer, Berlin.
- [Ha65] W. Hanf, "Model-Theoretic Methods in the Study of Elementary Logic," in J. Addison, L. Henkin, and A. Tarski, eds., *The Theory of Models*, 1965, North Holland, 105-135.
- [HIM78] J. Hartmanis, N. Immerman, and S. Mahaney, "One-Way Log Tape Reductions," *IEEE Found. of Comp. Sci. Symp.* (1978), 65-72.
- [Has86] J. Hastad, "Almost Optimal Lower Bounds for Small Depth Circuits," *18th ACM STOC Symp.*, (1986), 6-20.
- [He96] L. Hella, "Logical Hierarchies in PTIME," *Information and Computation* 129(1) (1996), 1-19.
- [HKL97] L. Hella, Ph. Kolaitis, and K. Luosto, "How to Define a Linear Order on Finite Models," *Annals of Pure and Applied Logic* 87 (1997), 241-267.
- [HKL96] L. Hella, Ph. Kolaitis, and K. Luosto, "Almost Everywhere Equivalence of Logics in Finite Model Theory," *Bulletin of Symbolic Logic* 2(4) (1996), 422 - 443.
- [HLN97] L. Hella, L. Libkin, and J. Nurmonen, "Notions of locality and their logical characterizations over finite models," manuscript.
- [Hil85] D. Hillis, *The Connection Machine* 1985, MIT Press.
- [Hon82] J.-W. Hong, "On Some Deterministic Space Complexity Problems," *SIAM J. Comput.* 11 (1982), 591-601.
- [Ho86] J.-W. Hong, *Computation: Computability, Similarity, and Duality*, 1986, John Wiley & Sons.

- [HT72] J. Hopcroft and R. Tarjan, "Isomorphism of Planar Graphs," in *Complexity of Computer Computations*, R. Miller and J.W. Thatcher, eds., (1972), Plenum Press, 131-152.
- [HU79] J. Hopcroft and J. Ullman, *Introduction to Automata Theory, Languages, and Computation*, Addison-Wesley (1979).
- [I79] N. Immerman, "Length of Predicate Calculus Formulas as a New Complexity Measure," *20th IEEE FOCS Symp.* (1979), 337-347. Revised version: "Number of Quantifiers is Better than Number of Tape Cells," *JCSS* 22(3), June 1981, 65-72.
- [I80] N. Immerman, "Upper and Lower Bounds for First Order Expressibility," *21st IEEE FOCS Symp.* (1980), 74-82. Revised version: *JCSS* 25(1) (1982), 76-98.
- [I82] N. Immerman, "Relational Queries Computable in Polynomial Time," *14th ACM STOC Symp.* (1982), 147-152. Revised version: *Information and Control*, 68 (1986), 86-104.
- [I83] N. Immerman, "Languages Which Capture Complexity Classes," *15th ACM STOC Symp.* (1983), 347-354. Revised version: "Languages That Capture Complexity Classes," *SIAM J. Comput.* 16(4) (1987), 760-778.
- [I87] N. Immerman, "Expressibility as a Complexity Measure: Results and Directions," *Second Structure in Complexity Theory Conf.* (1987), 194-202.
- [I88] N. Immerman, "Nondeterministic Space is Closed Under Complementation," *SIAM J. Comput.* 17(5) (1988), 935-938. Also appeared in *Third Structure in Complexity Theory Conf.* (1988), 112-115.
- [I89] N. Immerman, "Descriptive and Computational Complexity," in *Computational Complexity Theory*, ed. J. Hartmanis, Lecture Notes for AMS Short Course on Computational Complexity Theory, *Proc. Symp. in Applied Math.* 38, American Mathematical Society (1989), 75-91.
- [I89a] N. Immerman, "Expressibility and Parallel Complexity," *SIAM J. of Comput.* 18 (1989), 625-638.
- [I91] N. Immerman, " $\text{DSPACE}[n^k] = \text{VAR}[k+1]$," *Sixth IEEE Structure in Complexity Theory Symp.* (July, 1991), 334-340.

- [IKL95] N. Immerman, Ph. Kolaitis, and J. Lynch, "A Tutorial on Finite Model Theory," DIMACS, August, 1995.
- [IK87] N. Immerman and D. Kozen, "Definability with Bounded Number of Bound Variables," *Second LICS Symp.* (1987), 236-244.
- [IL95] N. Immerman, S. Landau, "The Complexity of Iterated Multiplication," *Information and Computation* 116(1) (1995), 103-116.
- [IL90] N. Immerman and E. Lander, "Describing Graphs: A First-Order Approach to Graph Canonization," in *Complexity Theory Retrospective*, Alan Selman, ed., Springer-Verlag (1990), 59-81.
- [IPS96] N. Immerman, S. Patnaik and D. Stemple, "The Expressiveness of a Family of Finite Set Languages," *Theoretical Computer Science* 155(1) (1996), 111-140. A preliminary version appeared in *Tenth ACM Symposium on Principles of Database Systems* (1991), 37-52.
- [IV97] N. Immerman and M.Y. Vardi, "Model Checking and Transitive Closure Logic," *Proc. 9th Int'l Conf. on Computer-Aided Verification* (1997), Lecture Notes in Computer Science, Springer-Verlag, 291 - 302.
- [JL77] N. Jones and W. Laaser, "Complete Problems for Deterministic Polynomial Time," *Theoret. Comp. Sci.* 3 (1977), 105-117.
- [JLL76] N. Jones, E. Lien and W. Laaser, "New Problems Complete for Nondeterministic Logspace," *Math. Systems Theory* 10 (1976), 1-17.
- [JS74] N. Jones and A. Selman, "Turing Machines and the Spectra of First-Order Formulas," *J. Symbolic Logic* 39 (1974), 139-150.
- [Ka79] R. Karp, "Probabilistic Analysis of a Canonical Numbering Algorithm for Graphs," *Relations between combinatorics and other parts of mathematics*, Proceedings of Symposia in Pure Mathematics 34, 1979, American Mathematical Society, 365 - 378.
- [KL82] R. Karp and R. Lipton, "Turing Machines That Take Advice," *Ensiegn. Math.* 28 (1982), 192-209.
- [KV95] Ph. Kolaitis and J. Väänänen, "Generalized Quantifiers and Pebble Games on Finite Structures," *Annals of Pure and Applied Logic*, 74(1) (1995), 23-75.

- [KV98] Ph. Kolaitis and M.Y. Vardi, “Conjunctive-Query Containment and Constraint Satisfaction,” *ACM Symp. Principles Database Systems* (1998).
- [KV92a] Ph. Kolaitis and M.Y. Vardi, “0-1 Laws for Fragments of Second-Order Logic: an Overview,” in Y. Moschovakis, editor, *Logic From Computer Science 1992*, Springer-Verlag, 265–286.
- [Ku87] L. Kučera, “Canonical Labeling of Regular Graphs in Linear Average Time,” *28th IEEE FOCS Symp.* (1987), 271-279.
- [Kur64] S. Kuroda, “Classes of Languages and Linear-Bounded Automata,” *Information and Control* 7 (1964), 207-233.
- [Kur94] R. Kurshan, *Computer-Aided Verification of Coordinating Processes*, 1994, Princeton University Press, Princeton, NJ.
- [L75] R. Ladner, “The Circuit Value Problem is log space complete for P,” *SIGACT News*, 7(1) (1975), 18 – 20.
- [LR96] R. Lassaigne and M. de Rougemont, *Logique et Complexité*, 1996, Hermes.
- [LJK87] K.J. Lange, B. Jenner, and B. Kirsig, “The Logarithmic Hierarchy Collapses: $A\Sigma_2^L = A\Pi_2^L$,” *14th ICALP* (1987).
- [Lei87] D. Leivant, “Characterization of Complexity Classes in Higher-Order Logic,” *Second Structure in Complexity Theory Conf.* (1987), 203–217.
- [Lei89] D. Leivant, “Descriptive Characterizations of Computational Complexity,” *J. Comput. Sys. Sci.* 39 (1989), 51-83.
- [LP81] H. Lewis and C. Papadimitriou, *Elements of the Theory of Computation* 1982, Prentice-Hall.
- [LP82] H. Lewis and C. Papadimitriou, “Symmetric Space Bounded Computation,” *Theoret. Comput. Sci.* 19 (1982), 161-187.
- [LV93] M. Li and P. Vitányi, *An Introduction to Kolmogorov Complexity and its Applications*, 1993, Springer-Verlag, New York.
- [L92] S. Lindell, “A purely logical characterization of circuit uniformity,” *7th Structure in Complexity Theory Conf.* (1992), 185-192.

- [L] S. Lindell, "How to define exponentiation from addition and multiplication in first-order logic on finite structures," manuscript.
- [Lin66] P. Lindström, "First Order Predicate Logic with Generalized Quantifiers," *Theoria*, 32 (1966), 186-195.
- [LG77] L. Lovász and P. Gács, "Some Remarks on Generalized Spectra," *Zeitschr. f. math, Logik und Grundlagen d. Math*, 23 (1977), 547-554.
- [Luk82] E. Luks, "Isomorphism of Graphs of Bounded Valence Can be Tested in Polynomial Time," *J. Comput. Sys. Sci.* 25 (1982), pp. 42-65.
- [Lyn82] J. Lynch, "Complexity Classes and Theories of Finite Models," *Math. Sys. Theory* 15 (1982), 127-144.
- [MP94] J. Makowsky, Y. Pnueli, "Arity versus alternation in Second-Order Logic," in *Logical Foundations of Computer Science*, A. Nerode, Y. Matiyasevich eds., Springer LNCS 813 1994, 240 - 252.
- [McM93] K. McMillan, *Symbolic Model Checking*, 1993, Kluwer.
- [MP71] R. McNaughton and S. Papert, *Counter-Free Automata*, 1971, MIT Press, Cambridge, MA.
- [Man89] U. Manber, *Introduction to Algorithms: A Creative Approach*, Addison-Wesley, (1989).
- [MT97] O. Matz and W. Thomas, "The Monadic Quantifier Alternation Hierarchy Over Graphs is Infinite," *IEEE Symp. Logic In Comput. Sci.* (1997), 236-244.
- [MI94] J.A. Medina and N. Immerman, "A Syntactic Characterization of NP-Completeness," *IEEE Symp. Logic In Comput. Sci.* (1994), 241-250.
- [MI96] J.A. Medina and N. Immerman, "A Generalization of Fagin's Theorem," *IEEE Symp. Logic In Comput. Sci.* (1996), 2 - 12.
- [MSV94] P. Miltersen, S. Subramanian, J. Vitter, and R. Tamassia, "Complexity Models for Incremental Computation," *Theoret. Comp. Sci.* (130:1) (1994), 203-236.
- [Mos74] Y. Moschovakis, *Elementary Induction on Abstract Structures*, North Holland (1974).

- [Mos80] Y. Moschovakis, *Descriptive set theory*, 1980, North-Holland Pub. Co., Amsterdam, 637 p.
- [NT95] N. Nisan and A. Ta-Shma, "Symmetric Logspace is Closed Under Complement," *Chicago J. Theoret. Comp. Sci.* (1995).
- [OO1] M. Otto, "Two Variable first-Order Logic Over Ordered Domains," *J. Symbolic Logic* 66(2) (2001), 685-702.
- [Ott96] M. Otto, "The Expressive Power of Fixed-Point Logic with Counting," *J. Symbolic Logic* 61(1) (1996), 147 - 176
- [Ott97] M. Otto, *Bounded Variable Logics and Counting: A Study in Finite Models*, 1997, Lecture Notes in Logic, vol. 9, Springer-Verlag.
- [Pap94] C. Papadimitriou, *Computational Complexity* 1994, Addison-Wesley.
- [Pap85] C. Papadimitriou, "A Note on the Expressive Power of Prolog," *EATCS Bulletin* 26 (1985), 21-23.
- [PY91] C. Papadimitriou and M. Yannakakis, "Optimization, Approximation, and Complexity Classes," *J. Comput. Sys. Sci.* , 43 (1991), 425-440.
- [PI94] S. Patnaik and N. Immerman, "Dyn-FO: A Parallel, Dynamic Complexity Class," *J. Comput. Sys. Sci.* 55(2) (1997), 199-209. A preliminary version appeared in *ACM Symp. Principles Database Systems* (1994), 210-221.
- [Poi82] B. Poizat, "Deux ou trois choses que je sais de Ln" *JSL* 47 (1982), 641-658.
- [R96] G. Ramalingam *Bounded Incremental Computation*, 1996, Springer LNCS 1089.
- [Raz87] A. Razborov, "Lower Bounds on the Size of Bounded Depth Networks Over a Complete Basis With Logical Addition," *Matematischi Zametki* 41 (1987), 598-607 (in Russian). English translation in *Mathematical Notes of the Academy of Sciences of the USSR* 41, 333-338.
- [Rei87] J. Reif, "On Threshold Circuits and Polynomial Computation," *Second Annual Structure in Complexity Theory Symp.* (1987), 118-123.
- [Rei05] O. Reingold, "Undirected ST-Connectivity in Log-Space", *ACM Symp. Theory Of Comput.* 2005, 376 - 385.

- [RS72] D. Rödding and H. Schwichtenberg, "Bemerkungen zum Spektralproblem," *Zeitschrift für math. Logik und Grundlagen der Mathematik* 18 (1972), 1-12.
- [Ros82] J. Rosenstein, *Linear Orderings*, 1982, Academic Press.
- [Ruz81] L. Ruzzo, "On Uniform Circuit Complexity," *J. Comp. Sys. Sci.*, 21, No. 2 (1981), 365-383.
- [Sav70] W. Savitch, "Relationships Between Nondeterministic and Deterministic Tape Complexities," *J. Comput. System Sci.* 4 (1970), 177-192.
- [Sav73] W. Savitch, "Maze Recognizing Automata and Nondeterministic Tape Complexity," *J. Comput. Sys. Sci.* 7 (1973), 389-403.
- [Sch97] N. Schweikardt, "The Monadic Quantifier Alternation Hierarchy over Grids and Pictures," *Annual Conference of the European Association for Computer Science Logic* (1997), M. Nielsen and W. Thomas, eds., 383-397.
- [Sch94] T. Schwentick, "Graph Connectivity and Monadic NP," *IEEE Found. of Comp. Sci. Symp.* (1994), 614-622.
- [Sch97a] T. Schwentick, "Padding and the Expressive Power of Existential Second-Order Logics," *Annual Conference of the European Association for Computer Science Logic* (1997), M. Nielsen and W. Thomas, eds., 399-412.
- [SB98] T. Schwentick and K. Barthelmann, "Local Normal Forms for First-Order Logic with Applications to Games and Automata," to appear in *Symp. Theoretical Aspects Comp. Sci.* (1998).
- [See95] D. Seese, "FO-Problems and Linear Time Computability," Tech Report, Institut für Informatik und Formale Beschreibungsverfahren, Universität Karlsruhe, Germany (1995).
- [Sip83] M. Sipser, "Borel Sets and Circuit Complexity," *15th Symp. on Theory of Computation* (1983), 61-69.
- [Smo87] R. Smolensky, "Algebraic Methods in the Theory of Lower Bounds for Boolean Circuit Complexity," *19th ACM STOC* (1987), 77-82.
- [Spe93] J. Spencer, "Zero-One Laws With Variable Probability," *J. Symbolic Logic* 58 (1993), 1-14.

- [Ste94] I. Stewart, "On completeness for NP via projection translations," *Math. Syst. Theory* 27 (1994), 125–157.
- [Ste91] I. Stewart, "Comparing the Expressibility of Languages Formed Using NP-Complete Operators," *J. Logic and Computation* 1(3) (1991), 305-330.
- [Sto77] L. Stockmeyer, "The Polynomial-Time Hierarchy," *Theoretical Comp. Sci.* 3 (1977), 1-22.
- [SV84] L. Stockmeyer and U. Vishkin, "Simulation of Parallel Random Access Machines by Circuits," *SIAM J. of Comp.* 13, No. 2 (1984), 409-422.
- [Str94] H. Straubing, *Finite Automata, Formal Logic, and Circuit Complexity*, 1994, Birkhäuser.
- [Sze88] R. Szelepcsényi, "The Method of Forced Enumeration for Nondeterministic Automata," *Acta Informatica* 26 (1988), 279-284.
- [Tar36] A. Tarski, "Der Wahrheitsbegriff in den Formalisierten Sprachen," *Studia Philosophica I* (1936).
- [Tar55] A. Tarski, "A Lattice-Theoretical Fixpoint Theorem and its Applications," *Pacific. J. Math.*, 55 (1955), 285-309.
- [Tho86] S. Thomas, "Theories With Finitely Many Models," *J. Symbolic Logic*, 51, No. 2 (1986), 374-376.
- [Tra50] B. Trahtenbrot, "The Impossibility of an Algorithm for the Decision Problem for Finite Domains," *Doklady Akademii Nauk SSSR*, n.s., vol 70 (1950), 569-572 (in Russian).
- [Tur84] G. Turán, "On the Definability of Properties of Finite Graphs," *Discrete Math.* 49 (1984), 291-302.
- [TPP97] A. Turk, S. Probst, and G. Powers, "Verification of a Chemical Process Leak Test Procedure," in *Computer Aided Verification, 9th International Conf.*, O. Grumberg, ed. 1997, Springer, 84-94.
- [Tys97] J. Tyszkiewicz, "The Kolmogorov Expression Complexity of Logics," *Information and Computation* 135(2) (1997), 113-136.
- [Vaa99] . Väänänen, ed., *Generalized Quantifiers and Computation*, Ninth European Summer School in Logic, Language, and Information, 1997, Springer LNCS 1754.

- [Val82] L. Valiant, "Reducibility By Algebraic Projections," *L'Enseignement mathématique*, 28, 3-4 (1982), 253-68.
- [vD94] D. van Dalen, *Logic and Structure, Third Edition*, 1994, Springer-Verlag.
- [vH67] J. van Heijenoort, *From Frege to Gödel: A Source Book in Mathematical Logic, 1879 - 1931* 1967, Harvard University Press.
- [Var82] M.Y. Vardi, "Complexity of Relational Query Languages," *14th Symposium on Theory of Computation* (1982), 137-146.
- [Wei76] B. Weisfeiler, ed., *On Construction and Identification of Graphs*, Lecture Notes in Mathematics 558, Springer, 1976.
- [Wra76] C. Wrathall, "Complete Sets and the Polynomial Hierarchy," *Theoret. Comp. Sci.* 3 (1976).
- [Yao85] A. Yao, "Separating the Polynomial-Time Hierarchy by Oracles," *26th IEEE Symp. on Foundations of Comp. Sci.* (1985), 1-10.