

# Brief Contents

|  |            |
|--|------------|
| Preface . . . . .                                  | viii       |
| Prologue . . . . .                                 | 1          |
| <b>I TIME AND SPACE</b>                            | <b>7</b>   |
| 1 Efficient Computation . . . . .                  | 9          |
| 2 Time Complexity . . . . .                        | 79         |
| 3 Space Complexity . . . . .                       | 148        |
| 4 A Time-Space Lower Bound . . . . .               | 176        |
| 5 Nonuniformity . . . . .                          | 212        |
| 6 Parallelism . . . . .                            | 253        |
| <b>II RANDOMNESS</b>                               | <b>289</b> |
| 7 Randomized Computation . . . . .                 | 291        |
| 8 Hashing . . . . .                                | 337        |
| 9 Error Correcting Codes . . . . .                 | 360        |
| 10 Heuristics . . . . .                            | 390        |
| 11 Pseudorandomness . . . . .                      | 416        |
| <b>III SCOPE OF COMPLEXITY</b>                     | <b>459</b> |
| 12 Cryptography . . . . .                          | 461        |
| 13 Learning . . . . .                              | 494        |
| 14 Optimization . . . . .                          | 546        |
| 15 Verification . . . . .                          | 585        |
| <b>IV CONCRETE MODELS</b>                          | <b>621</b> |
| 16 Decision Trees and Branching Programs . . . . . | 623        |
| 17 Communication Protocols . . . . .               | 665        |
| 18 Circuits and Formulas . . . . .                 | 695        |
| Epilogue . . . . .                                 | 727        |
| Discrete Math Background . . . . .                 | 735        |
| References . . . . .                               | 759        |
| Index . . . . .                                    | 791        |

# Contents

|  |             |
|--|-------------|
| <b>Preface</b>   | <b>viii</b> |
| <b>Prologue</b>  | <b>1</b>    |
| Part I: Time and space . . . . .                           | 1           |
| Part II: Randomness . . . . .                              | 3           |
| Part III: Scope of complexity . . . . .                    | 4           |
| Part IV: Concrete models . . . . .                         | 5           |
| <b>I TIME AND SPACE</b>                                    | <b>7</b>    |
| <b>1 Efficient Computation</b>                             | <b>9</b>    |
| 1.1 Problems and algorithms . . . . .                      | 9           |
| 1.2 Efficiency . . . . .                                   | 14          |
| 1.3 Input/output encodings . . . . .                       | 20          |
| 1.4 The model of computation . . . . .                     | 25          |
| 1.5 Robustness of the model . . . . .                      | 33          |
| 1.6 Complexity classes . . . . .                           | 43          |
| 1.7 Time versus space . . . . .                            | 47          |
| 1.8 Time-efficient reductions . . . . .                    | 54          |
| 1.9 Space-efficient reductions . . . . .                   | 67          |
| Exercises . . . . .  | 73          |
| Notes . . . . .  | 77          |
| <b>2 Time Complexity</b>                                   | <b>79</b>   |
| 2.1 Tractable and seemingly intractable problems . . . . . | 79          |
| 2.2 P versus NP . . . . .                                  | 87          |
| 2.3 NP-completeness . . . . .                              | 94          |
| 2.4 Independent sets . . . . .                             | 106         |
| 2.5 Long paths . . . . .                                   | 111         |
| 2.6 Graph colorings . . . . .                              | 115         |
| 2.7 Subset sums . . . . .                                  | 119         |
| 2.8 Satisfying assignments . . . . .                       | 122         |
| 2.9 Variants of P versus NP . . . . .                      | 137         |
| Exercises . . . . .  | 139         |
| Notes . . . . .  | 146         |

|          |   |            |
|----------|---|------------|
| <b>3</b> | <b>Space Complexity</b>                       | <b>148</b> |
| 3.1      | Logic, games, and recursion . . . . .         | 148        |
| 3.2      | PSPACE-completeness . . . . .                 | 154        |
| 3.3      | NL and NL-completeness . . . . .              | 158        |
| 3.4      | Turning verifiers into solvers . . . . .      | 163        |
| 3.5      | NL is closed under complement . . . . .       | 167        |
|          | Exercises . . . . .                           | 170        |
|          | Notes . . . . .                               | 174        |
| <br>     |   |            |
| <b>4</b> | <b>A Time-Space Lower Bound</b>               | <b>176</b> |
| 4.1      | Padding . . . . .                             | 176        |
| 4.2      | Time and space hierarchies . . . . .          | 181        |
| 4.3      | The polynomial hierarchy . . . . .            | 190        |
| 4.4      | Trading space for time . . . . .              | 196        |
| 4.5      | Time-space-efficient reductions . . . . .     | 199        |
| 4.6      | Putting it all together . . . . .             | 207        |
|          | Exercises . . . . .                           | 208        |
|          | Notes . . . . .                               | 210        |
| <br>     |   |            |
| <b>5</b> | <b>Nonuniformity</b>                          | <b>212</b> |
| 5.1      | Algorithms that take advice . . . . .         | 212        |
| 5.2      | Logic circuits . . . . .                      | 218        |
| 5.3      | Time versus circuit size . . . . .            | 228        |
| 5.4      | P-completeness . . . . .                      | 235        |
| 5.5      | Branching programs . . . . .                  | 239        |
| 5.6      | Space versus branching program size . . . . . | 245        |
|          | Exercises . . . . .                           | 249        |
|          | Notes . . . . .                               | 252        |
| <br>     |   |            |
| <b>6</b> | <b>Parallelism</b>                            | <b>253</b> |
| 6.1      | Parallel time . . . . .                       | 253        |
| 6.2      | Arithmetic . . . . .                          | 260        |
| 6.3      | Space versus parallel time . . . . .          | 263        |
| 6.4      | L-completeness . . . . .                      | 266        |
| 6.5      | Parallelization . . . . .                     | 271        |
| 6.6      | Narrow branching programs . . . . .           | 275        |
|          | Exercises . . . . .                           | 286        |
|          | Notes . . . . .                               | 287        |

|           |   |            |
|-----------|---|------------|
| <b>II</b> | <b>RANDOMNESS</b>                                       | <b>289</b> |
| <b>7</b>  | <b>Randomized Computation</b>                           | <b>291</b> |
| 7.1       | Time-efficient randomized algorithms . . . . .          | 291        |
| 7.2       | The model and complexity classes . . . . .              | 298        |
| 7.3       | Amplification . . . . .                                 | 301        |
| 7.4       | Derandomizing BPP . . . . .                             | 305        |
| 7.5       | Implications of randomized algorithms for NP . . . . .  | 307        |
| 7.6       | Concentration bounds . . . . .                          | 309        |
| 7.7       | Space-efficient randomized algorithms . . . . .         | 315        |
| 7.8       | Derandomizing BPL . . . . .                             | 320        |
| 7.9       | Promise problems . . . . .                              | 326        |
| 7.10      | Map of complexity classes . . . . .                     | 330        |
|           | Exercises . . . . .                                     | 330        |
|           | Notes . . . . .   | 335        |
| <b>8</b>  | <b>Hashing</b>  | <b>337</b> |
| 8.1       | Randomized hashing . . . . .                            | 337        |
| 8.2       | Constructing pairwise uniformity . . . . .              | 340        |
| 8.3       | Isolating witnesses . . . . .                           | 343        |
| 8.4       | Approximately counting and sampling witnesses . . . . . | 346        |
| 8.5       | Hash mixing and applications . . . . .                  | 352        |
|           | Exercises . . . . .                                     | 356        |
|           | Notes . . . . .   | 359        |
| <b>9</b>  | <b>Error Correcting Codes</b>                           | <b>360</b> |
| 9.1       | Binary codes . . . . .                                  | 360        |
| 9.2       | Bounds for binary codes . . . . .                       | 366        |
| 9.3       | Polynomial codes . . . . .                              | 371        |
| 9.4       | Bit string fields . . . . .                             | 375        |
| 9.5       | Codes using bit string fields . . . . .                 | 379        |
|           | Exercises . . . . .                                     | 386        |
|           | Notes . . . . .   | 389        |
| <b>10</b> | <b>Heuristics</b>                                       | <b>390</b> |
| 10.1      | Distributional problems . . . . .                       | 390        |
| 10.2      | Average-case analogue of P versus NP . . . . .          | 394        |
| 10.3      | Worst-case versus average-case . . . . .                | 397        |
| 10.4      | Hardness amplification . . . . .                        | 405        |
|           | Exercises . . . . .                                     | 412        |
|           | Notes . . . . .   | 414        |

|            |  |            |
|------------|--|------------|
| <b>11</b>  | <b>Pseudorandomness</b>                                    | <b>416</b> |
| 11.1       | Statistical distance and pseudorandom generators . . . . . | 416        |
| 11.2       | Fooling parities, clauses, and terms . . . . .             | 423        |
| 11.3       | Fooling space-efficient computations . . . . .             | 426        |
| 11.4       | Fooling time-efficient computations . . . . .              | 435        |
| 11.5       | Extracting randomness . . . . .                            | 442        |
| 11.6       | Seeded extractors . . . . .                                | 445        |
|            | Exercises . . . . .  | 452        |
|            | Notes . . . . .  | 457        |
| <br>       |  |            |
| <b>III</b> | <b>SCOPE OF COMPLEXITY</b>                                 | <b>459</b> |
| <br>       |  |            |
| <b>12</b>  | <b>Cryptography</b>  | <b>461</b> |
| 12.1       | Perfectly secure encryption . . . . .                      | 462        |
| 12.2       | Computationally secure encryption . . . . .                | 465        |
| 12.3       | Pseudorandom generators from one-way functions . . . . .   | 469        |
| 12.4       | Pseudorandom functions . . . . .                           | 478        |
| 12.5       | Applications of pseudorandom functions . . . . .           | 485        |
|            | Exercises . . . . .  | 490        |
|            | Notes . . . . .  | 493        |
| <br>       |  |            |
| <b>13</b>  | <b>Learning</b>  | <b>494</b> |
| 13.1       | Probably approximately correct learners . . . . .          | 494        |
| 13.2       | Consistent learners . . . . .                              | 500        |
| 13.3       | Time complexity . . . . .                                  | 506        |
| 13.4       | Sample complexity . . . . .                                | 512        |
| 13.5       | Amplification and boosting . . . . .                       | 521        |
| 13.6       | Exact learners . . . . .                                   | 526        |
| 13.7       | Learning circuits . . . . .                                | 534        |
|            | Exercises . . . . .  | 539        |
|            | Notes . . . . .  | 544        |
| <br>       |  |            |
| <b>14</b>  | <b>Optimization</b>  | <b>546</b> |
| 14.1       | Approximation algorithms . . . . .                         | 547        |
| 14.2       | Hardness of approximation . . . . .                        | 553        |
| 14.3       | Label cover . . . . .                                      | 562        |
| 14.4       | Probabilistically checkable proofs . . . . .               | 566        |
|            | Exercises . . . . .  | 579        |
|            | Notes . . . . .  | 584        |

|           |  |            |
|-----------|--|------------|
| <b>15</b> | <b>Verification</b>                              | <b>585</b> |
| 15.1      | Arthur and Merlin . . . . .                      | 586        |
| 15.2      | Interactive proofs . . . . .                     | 594        |
| 15.3      | Delegation and program checking . . . . .        | 607        |
| 15.4      | Zero-knowledge proofs . . . . .                  | 613        |
|           | Exercises . . . . .                              | 617        |
|           | Notes . . . . .                                  | 619        |
| <b>IV</b> | <b>CONCRETE MODELS</b>                           | <b>621</b> |
| <b>16</b> | <b>Decision Trees and Branching Programs</b>     | <b>623</b> |
| 16.1      | Deterministic query complexity . . . . .         | 624        |
| 16.2      | Randomized query complexity . . . . .            | 630        |
| 16.3      | Certificate complexity . . . . .                 | 639        |
| 16.4      | Sensitivity . . . . .                            | 644        |
| 16.5      | Deterministic versus randomized . . . . .        | 649        |
| 16.6      | Decision tree size . . . . .                     | 653        |
| 16.7      | Branching program size . . . . .                 | 655        |
|           | Exercises . . . . .                              | 659        |
|           | Notes . . . . .                                  | 664        |
| <b>17</b> | <b>Communication Protocols</b>                   | <b>665</b> |
| 17.1      | Deterministic communication complexity . . . . . | 665        |
| 17.2      | Randomized communication complexity . . . . .    | 673        |
| 17.3      | Asymmetry . . . . .                              | 679        |
| 17.4      | Applications . . . . .                           | 682        |
| 17.5      | Search problems . . . . .                        | 687        |
|           | Exercises . . . . .                              | 690        |
|           | Notes . . . . .                                  | 694        |
| <b>18</b> | <b>Circuits and Formulas</b>                     | <b>695</b> |
| 18.1      | Gate elimination . . . . .                       | 695        |
| 18.2      | Circuits and polynomials . . . . .               | 697        |
| 18.3      | Circuits and random restrictions . . . . .       | 700        |
| 18.4      | Formulas and random restrictions . . . . .       | 708        |
| 18.5      | Formulas and communication protocols . . . . .   | 712        |
| 18.6      | Monotone circuits . . . . .                      | 717        |
| 18.7      | Natural proofs barrier . . . . .                 | 720        |
|           | Exercises . . . . .                              | 723        |
|           | Notes . . . . .                                  | 725        |

---

|  |            |
|--|------------|
| <b>Epilogue</b>                              | <b>727</b> |
| Fine-grained complexity . . . . .            | 727        |
| Parameterized complexity . . . . .           | 727        |
| Proof complexity . . . . .                   | 728        |
| Expander graphs . . . . .                    | 729        |
| Harmonic analysis . . . . .                  | 730        |
| Program size complexity . . . . .            | 730        |
| Complexity of counting . . . . .             | 730        |
| Arithmetic complexity . . . . .              | 731        |
| Property testing . . . . .                   | 731        |
| Quantum computing . . . . .                  | 731        |
| Complexity beyond computer science . . . . . | 733        |
| Notes . . . . .                              | 733        |
| <br>   |            |
| <b>Discrete Math Background</b>              | <b>735</b> |
| 0.1 Combinatorics . . . . .                  | 735        |
| 0.2 Graphs . . . . .                         | 738        |
| 0.3 Probability . . . . .                    | 743        |
| 0.4 Algebra . . . . .                        | 748        |
| <br>   |            |
| <b>References</b>                            | <b>759</b> |
| <br>   |            |
| <b>Index</b>                                 | <b>791</b> |

# Preface

## Why this book?

Computational complexity theory is part of the foundation of computer science (CS). It provides fundamental insights about many areas of CS, including security, artificial intelligence, big data processing, verification, optimization, parallel computing, and distributed computing. Because computational complexity is sometimes abstract and technical, it might appear too isolated to have practical relevance. To convey how computational complexity is germane to CS as a whole, this book highlights connections to other areas of CS and emphasizes broadly useful intuitions and techniques for reasoning about computing.

Some students are reluctant to study computational complexity because it's math-heavy. This book aims to make the subject accessible to readers with limited math background. A music teacher once lamented, "They don't write easy music in hard keys." Students struggle to adapt to key signatures with lots of sharps or flats because most music written in those keys is technically demanding. For many people, mathematical rigor is a "hard key." This book provides some "easy music" to help readers acclimate and see that computational complexity is not so formidable. The appendix reviews relevant discrete math background.

## Intended audience

This book is aimed at general CS graduate students, but it's also suitable for advanced undergraduates, students in related fields, working professionals, and armchair computer scientists.

The reader is assumed to have a typical undergraduate education in CS, though an intrepid reader can appreciate much of the book without such a background. The reader should have experience writing programs and a little familiarity with data structures (including stacks, queues, and dictionaries), algorithms (such as binary search and merge sort), discrete math, assembly language, and digital logic. The text sometimes alludes to basic concepts of operating systems, compilers, computer architecture, security, and artificial intelligence.

An important prerequisite is mathematical maturity: the ability to mentally connect notations to the abstract concepts they represent, to recognize valid logical reasoning, and to distinguish key ideas from technical details.

## Goals of this book

The foremost reason to study computational complexity is to gain insight about the essence of computing. This book offers intuition with "proof idea" sections and clarifies definitions with plenty of concrete examples. The book sometimes only presents interesting special cases—rather than general results—when they're easier to absorb. To help the reader solidify the material, each chapter ends with exercises ranging from calisthenics to guided advanced material.

This book is “batteries included”: Everything is explained from first principles, without requiring the reader to fill in details. The book is both a self-contained tutorial and a thorough reference with full proofs. The reader is advised when certain proofs are ancillary and safe to skip. A few advanced “museum exhibit” theorems are stated without proof. (Full disclosure: An auxiliary “battery” that’s unfortunately not included is the existence of prime numbers and irreducible polynomials of suitable sizes. These standard theorems are a bit too deep to prove here, but we need them for constructing finite fields, so we accept them as gifts from pure math.)

This book tries to appeal to a broad CS audience by mentioning pragmatic motivations and by connecting complexity theoretic ideas to analogues elsewhere in CS. Most areas of CS benefit from computational complexity’s rigorous perspective on the power and limitations of computing. The book covers subareas of complexity that form theoretical foundations of other areas of CS, such as cryptography and machine learning.

For general CS students, the book can be an accessible survey of complexity theory and provide inspiration about the benefits of rigorously formulating questions about computing. For students specializing in theoretical computer science, the book covers standard material and some advanced topics. The reader will emerge ready to study research papers in complexity.

Another goal is to rein in the necessary math background. The basic background includes logic, combinatorics, graph theory, discrete probability, and algebra limited to prime fields, polynomials, vectors, and matrices—nothing more abstract is needed. The appendix reviews notation, terminology, proof techniques, and basic lemmas. Slightly more advanced tools—including concentration bounds and power-of-2 fields—are developed in the main text. The book entirely avoids calculus, information theory, linear programming, continuous probability, and complex numbers. Inductive proofs are phrased as iterative or recursive algorithms that maintain invariants (rather than “raw induction”), befitting the CS context.

Perhaps controversially, the main text doesn’t mention researchers’ names that got attached to results and definitions in the literature. For example, some people may decry this book’s avoidance of the name “Markov’s inequality.” The goals are to avoid perpetuating any misallocation of credit (since the history of CS and math is littered with attributions of questionable veracity) and to emphasize that the results and ideas are fundamental parts of reality, independent of the circumstances of their discovery. Bibliographic references and historical context (including traditional names) are relegated to the “Notes” section in each chapter. A few exceptions are when someone’s name has become a common adjective with a suffix, such as “boolean” and “cartesian.”

The overarching goal is to tell a story, with plotlines woven together like a satisfying novel. This book presents computational complexity as a cohesive subject whose many threads are ultimately interconnected, rather than a collection of disparate topics. Each of the book’s four parts ends with a “plot twist.”

## Manifesto on the model of computation

Turing machines are the most traditional model of computation for teaching complexity theory. This book uses the Word RAM model instead. Advantages of Word RAM include:

- It resembles how actual computers work. Turing machines' back-and-forth sequential-access tapes may give students reservations about the relevance to real-world computing. But students generally don't need convincing that Word RAM—which is an idealized assembly language—appropriately distills the essence of computation. It's possible to define random-access Turing machines with special tapes for constructing addresses to teleport tape heads, but this is awkward and doesn't fully purge sequential access from the model.
- Implementing algorithms in the Word RAM model is reasonably straightforward (and only slightly tedious). For example, the breadth-first search algorithm naturally corresponds to a modest amount of Word RAM code. Implementing algorithms as Turing machines can be nightmarish. The author shudders to envision the state diagram of a random-access Turing machine for breadth-first search. Turing machines must fuss about determining the input size, delineating boundaries between pieces of data, doing simple operations such as addition at the bit level, and constructing memory addresses bit-by-bit. These boring details distract from the ideas and get swept under the rug (“clearly it can be done”). Word RAM helps students feel anchored because a high-level algorithm is never far from a low-level implementation.
- With Word RAM, the input size and running time conform to our intuition about algorithms. For example, consider summing  $N$  integers that fit in  $O(\log N)$ -bit words. With Word RAM, the input size is  $N$  words, and the running time is  $O(N)$ , as expected. With Turing machines, the input size is  $O(N \log N)$  bits, and the running time must account for doing each addition at the bit level. Analyses of Turing machines have omnipresent polylogarithmic factors to soak up these low-level nuisances. Analyses of Word RAM implementations are consistent with our expectations about algorithms.

Disadvantages of Word RAM include:

- It's not a “theoretically minimal” model. The instruction set is chosen more for convenience than for mathematical purity. Some people find the austerity of Turing machines appealing.
- A Word RAM program must specify how the word size depends on the input size. This is extra baggage attached to a program's code. It sometimes throws a small wrench in the works, such as in the proofs of the time and space hierarchy theorems. Also, defining the set of permissible word sizes is subtle. (Some versions of Word RAM let the word size change during a computation, but the author considers that too unrealistic.)
- Combining programs that have different word sizes—or modifying programs in a way that changes the word size—requires attention to preserve the programs' intended behavior. This minor annoyance doesn't crop up with Turing machines because their “word size” is constant.

## Content and organization

This book has more than enough material for two semesters. A typical one-semester graduate complexity course would cover much of the first seven chapters and selections from other chapters. The most fundamental material is in the first seven chapters. The next two chapters cover widely useful tools (pairwise uniform hashing and error correcting codes) and are independent of each

other. The rest of the chapters often use concepts from the first nine chapters but are largely independent of each other (and can be read in any order), except minor dependencies such as:

- Locally decodable codes (§10.3) are relevant to probabilistically checkable proofs (§14.4).
- Pseudorandom generators (§11.1) are relevant to encryption schemes (§12.2).
- Pseudorandom functions (§12.4) are relevant to hardness of learning circuits (§13.7) and to the natural proofs barrier (§18.7).
- Commitment schemes (§12.3) are relevant to zero-knowledge proofs (§15.4).
- Communication complexity of search problems (§17.5) is relevant to formula size complexity (§18.5).

Here are some of the author’s editorial decisions about the selection and organization of material:

- The time-space lower bound for SATISFIABILITY (Chapter 4) serves as a unified excuse to introduce several fundamental concepts such as diagonalization and the polynomial hierarchy. Since the polynomial hierarchy is treated as more of a “means” than an “end,” material about it is distributed across relevant chapters rather than in a dedicated chapter.
- The book avoids esoteric complexity classes and antiquated topics in structural complexity such as relativization, exact counting of witnesses, the isomorphism conjecture, and sparse decision problems.
- The chapters on concrete nonuniform models of computation are the last part of the book because the uniform model (Word RAM programs) is discarded at this point, and because the circuit complexity approach to P versus NP in the final chapter brings the book “full circle.”
- When carving a sculpture, some perfectly good stone will be lost: The epilogue is a teaser for important topics that didn’t make it into the text for various reasons, such as being math-heavier or harder to motivate in the broader context of CS.

## Acknowledgments

My perspective on computational complexity was profoundly influenced by my advisors throughout my education—Dieter van Melkebeek, Luca Trevisan, and Toniann Pitassi—and by all my collaborators, especially Mika Göös. I thank Dieter van Melkebeek and anonymous reviewers for helpful feedback. I thank the people at Cambridge University Press for help during the publication process. The dblp website was very useful for curating references. I thank the US National Science Foundation for financial support. Finally, I thank my family for essential nontechnical support.