

To Maria and Miriam,
with my deepest gratitude

Preface

The most valuable acquisitions in a scientific or technical education are the general-purpose mental tools which remain serviceable for a life-time.

—George Forsythe, "What to do until computer scientist comes." (1968)

Algorithms play the central role both in science and practice of computing. The recognition of this fact has led to the appearance of a considerable number of textbooks on the subject. By and large, they follow one of the two alternatives in presenting it. The first one classifies algorithms according to a problem type. Such a book would have separate chapters on algorithms for sorting, searching, graphs, and so on. The advantage of this approach is that it allows an immediate comparison of, say, the efficiency of different algorithms for the same problem. The drawback of this approach is that it emphasizes problem types at the expense of algorithm design techniques.

The second alternative is to organize the presentation around algorithm design techniques. Under this organization, algorithms from different areas of computing are grouped together if they are based on the same design approach. I share the belief of many (e.g., [BaY95]) that this organization is more appropriate for the basic course on the design and analysis of algorithms. There are three principal reasons for this emphasis on algorithm design techniques. First, these techniques provide a student with tools for designing algorithms for new problems. This makes learning algorithm design technique a very valuable endeavor from the practical standpoint. Second, they seek to classify multitudes of known algorithms according to an underlying design idea. Learning to see such commonality among algorithms from different application areas should be a major goal of computer science education. After all, every science considers classification of its principal subject as a major if not the central point of its discipline. Third, in my opinion, algorithm design techniques have considerable utility as general problem solving strategies, which are applicable to problems beyond computing.

There are, of course, several textbooks that are organized around algorithm design techniques. The problem with these books is that they uncritically follow the same classification of design techniques. This classification has several serious shortcomings, from both theoretical and educational points of view. The most significant of the shortcomings is its failure to classify many important algorithms. This limitation has forced the authors of existing textbooks to depart from the design technique organization and to include chapters dealing with specific problem types. Unfortunately, such a switch leads to a loss of course coherence and almost unavoidably creates a confusion in students' mind.

The new taxonomy of algorithm design techniques

It was my frustration with shortcomings of the existing classification of algorithm design techniques that has motivated me to develop a new taxonomy of them [Lev99], which is the basis of this book. Here are the principal advantages of the new taxonomy:

- The new taxonomy is more comprehensive than the traditional one. It includes several strategies—brute-force, decrease-and-conquer, transform-and-conquer, and time and space tradeoffs—that are rarely if ever recognized as important design paradigms.
- The new taxonomy covers naturally many classic algorithms (Euclid’s algorithm, heapsort, search trees, hashing, topological sorting, Gaussian elimination, Horner’s rule—to name a few examples) that the traditional taxonomy cannot classify. As a result, it makes it possible to present the standard body of classic algorithms in a unified and coherent fashion.
- It naturally accommodates the existence of important varieties of several design techniques. (For example, it recognizes three variations of decrease-and-conquer and three variations of transform-and-conquer.)
- It is better aligned with analytical methods for the efficiency analysis (see Appendix B).

Design techniques as general problem solving strategies

Most applications of the design techniques in the book are to classic problems of computer science. The only innovation here is an inclusion of some material on numerical algorithms, which are covered within the same general framework. (An inclusion of numerical algorithms is encouraged by *Computing Curricula 2001* [CC01]—a new model curriculum for computer science programs.) But the design techniques can be considered general problem solving tools, whose applications are not limited to traditional computing and mathematical problems. Two factors make this point particularly important. First, more and more computing applications go beyond its traditional domain, and there are all reasons to believe that this trend will only strengthen in the future. Second, developing students’ problem solving skills has come to be recognized as a major goal of the college education. Among all the courses in a computer science curriculum, a course on the design and analysis of algorithms is uniquely suitable for this task, because it can offer a student specific strategies for solving problems. I am not proposing that a course on the design and analysis of algorithms should become a course on general problem solving. But I do believe that the unique opportunity provided by studying the design and analysis of algorithms should not be missed. Towards this goal, the book includes applications to puzzles and puzzle-like games. Also using puzzles in teaching algorithms is certainly

not a new idea, the book tries to do this systematically by going beyond a few standard examples..

Textbook pedagogy

My goal was to write a text that would not trivialize the subject but still would be readable by most students on their own. Here are some of the things done in the book toward this objective.

- Sharing the opinion of George Forsythe (see the epigraph), I have sought to stress major ideas underlying the design and analysis of algorithms. In choosing specific algorithms to illustrate these ideas, I limited the number of covered algorithms to those that demonstrate an underlying design technique or an analysis method most clearly. Fortunately, most classic algorithms satisfy this criterion.
- In Chapter 2, which is devoted to the efficiency analysis, the methods used for analyzing nonrecursive algorithms are separated from those typically used for analyzing recursive algorithms. The chapter also includes sections devoted to empirical analysis and algorithm visualization.
- The narrative is systematically interrupted by questions to the reader. Some of them are asked rhetorically, in anticipation of a concern or doubt, and are answered immediately. The goal of the others is to prevent the reader from drifting through the text without a satisfactory level of comprehension.
- Each chapter ends with a summary recapping the most important concepts and results discussed in the chapter.
- The book contains about 600 exercises. Some of them are of the drill nature; others make important points about the material covered in the body of the text or introduce algorithms not covered there at all; a few exercises take advantage of Internet resources. Several exercises are designed to prepare the reader for the material covered later in the book. More difficult problems—there are not many of them—are marked by special symbols. (Because marking problems as difficult may discourage some students from trying to tackle them, this marking is made in the Instructor’s Manual but not in the book itself.) On the other hand, puzzles, games, and puzzle-like questions are marked in the exercises with a special icon.
- The book provides hints to all the exercises. Detailed solutions, except for programming projects, are provided in the Instructor’s Manual available to qualified adopters from the publisher. (Contact your Addison-Wesley representative, or email aw.cse@aw.com.) The supplements available to all readers of this book are at www.aw.com/cssupport.

Prerequisites

The book assumes that a reader has gone through an introductory programming course and a standard course on discrete structures. With such a background, he or she should be able to handle the book's material without undue difficulty. Still, fundamental data structures, necessary summation formulas, and recurrence relations are reviewed in Section 1.4, Appendix A, and Appendix B, respectively. Calculus is used in only three sections (Section 2.2, 10.4, and 11.4), and to a very limited degree; if students lack calculus as an assured part of their background, the portions of these three sections can be omitted without hindering their understanding of the rest of the material.

Use in the Curriculum

The book can serve as a textbook for the basic course on design and analysis of algorithms, which is organized around algorithm design techniques. It might contain too much material for a typical one-semester course. By and large, portions of Chapters 3 through 11 can be skipped without a danger of making later parts of the book incomprehensible to the reader. Any portion of the book can be assigned for self-study. In particular, Sections 2.6 and 2.7 on empirical analysis and algorithm visualization, respectively, can be assigned in conjunction with projects.

Here is a possible plan for a one-semester course; it assumes a 40-class meeting format.

Lecture	Topic	Sections
1-2	Introduction	1.1–1.3
3-4	Analysis framework; O , Θ , Ω notations	2.1–2.2
5	Mathematical analysis of nonrecursive algorithms	2.3
6-7	Mathematical analysis of recursive algorithms	2.4–2.5(+App.B)
8	Brute-force algorithms	3.1–3.2(+3.3)
9	Exhaustive search	3.4
10-12	Divide-and-conquer: mergesort, quicksort, binary search	4.1–4.3
13	Other divide-and-conquer examples	4.4 or 4.5 or 4.6
14-16	Decrease-by-one: insertion sort, DFS & BFS, topological sorting	5.1–5.3
17	Decrease-by-a-constant-factor algorithms	5.5
18	Variable-size-decrease algorithms	5.6
19-21	Presorting, Gaussian elimination, balanced search trees	6.1–6.3
22	Representation change: heaps and heapsort	6.4
23	Representation change: Horner's rule and binary exponentiation	6.5
24	Problem reduction	6.6
25-27	Space-time tradeoffs: string matching, hashing, B-trees	7.2–7.4
28-30	Dynamic programming algorithms	3 from 8.1-8.4
31-33	Greedy algorithms: Prim's, Kruskal's, Dijkstra's, Huffman's	9.1–9.4
34	Lower-bound arguments	10.1
35	Decision trees	10.2
36	P , NP , and NP -complete problems	10.3
37	Numerical algorithms	10.4(+11.4)
38	Backtracking	11.1
39	Branch-and-bound	11.2
40	Approximation algorithms for NP -hard problems	11.3

Acknowledgments

I would like to start by acknowledging the authors of other algorithm textbooks from whose insights and presentation ideas I have benefited both directly and indirectly. The advice and criticism of the book's reviewers have made the book better than it would have been otherwise. I am thankful to Simon Berkovich (George Washington University), Richard Borie (University of Alabama), Douglas M. Campbell (Brigham Young University), Bin Cong (California State University, Fullerton), Steve Homer (Boston University), Roland Hübscher (Auburn University), Sukhamay Kundu (Louisiana State University), Sheau-Dong Lang (University of Central Florida), John C. Lusth (University of Arkansas), John F. Meyer (University of Michigan), Steven R. Seidel (Michigan Technological University), Ali Shokoufandeh (Drexel University), and George H. Williams (Union College).

I am grateful to Mary-Angela Papalaskari, who used the manuscript in teaching a course on algorithms at Villanova and suggested several improvements to the text and the exercises. She also enthusiastically supported the idea of a systematic utilization of puzzles in the book. My other colleague, John Matulis,

used the manuscript in his sections, too, and provided me with a useful feedback. My former student, Andiswa Heinegg, helped to prepare the manuscript and, through her critique, made it more clear, both in content and style.

Students at Villanova have suffered inconvenience of using the manuscript as their textbook over the past few semesters. I acknowledge their patience, a useful feedback, and corrections to the errors and typos they found. The remaining errors are not, of course, their fault: I introduced them after they had taken the course.

I thank all the people at Addison-Wesley and its associates who worked on my book. I am especially grateful to my editor Maite Suarez-Rivas and her former assistant Lisa Hogue for sharing and supporting my enthusiasm for this project.

Finally, I am indebted to two members of my family. Living with a spouse writing a book is probably more trying than doing an actual writing. My wife Maria lived through two years of this, helping me any way she could. And help she did: all 250 or so figures in the book were created by her. My daughter Miriam has been my English prose guru over many years. Not only she read large portions of the book; she was instrumental in finding the epigraphs to the book's chapters. I affectionately dedicate this book to her and Maria.

Anany Levitin
anany.levitin@villanova.edu
August, 2002