


数据结构与算法分析 (C++版) (第三版)

Data Structures and Algorithm Analysis in C++
Third Edition



英文版

[美] Clifford A. Shaffer 著



电子工业出版社
PUBLISHING HOUSE OF ELECTRONICS INDUSTRY

<http://www.phei.com.cn>

TP311.12/Y16

2013.

国外计算机科学教材系列

数据结构与算法分析

(C++ 版)(第三版)

(英文版)

Data Structures and Algorithm Analysis in C++

Third Edition

[美] Clifford A. Shaffer 著



電子工業出版社

Publishing House of Electronics Industry

北京·BEIJING

内 容 简 介

本书采用程序员最爱用的面向对象 C++ 语言来描述数据结构和算法,并把数据结构原理和算法分析技术有机地结合在一起,系统介绍了各种类型的数据结构和排序、检索的各种方法。作者非常注意对每一种数据结构的存储方法及有关算法进行分析比较。书中还引入了一些比较高级的数据结构与先进的算法分析技术,并介绍了可计算性理论的一般知识。书中分别给出了 C++ 实现方法和伪码实现方法,便于读者根据情况选择。本书作者维护的网站上可下载相关代码、编程项目和辅助练习资料。

本书适合作为大专院校计算机软件专业与计算机应用专业学生的教材和参考书,也适合计算机工程技术人员参考。

ISBN13: 978-0-486-48582-9

ISBN10: 0-486-48582-X

Data Structures and Algorithm Analysis in C++, Third Edition

Copyright © 2011 by Clifford A. Shaffer

All right reserved.

Dover 出版社所出版的这个版本首印于 2011 年,是对 A Practical Introduction to Data Structures and Algorithm Analysis 的第三版(可从作者网站 <http://people.csvt.edu/~shaffer/Book/> 下载)进行过修订与更新后首次以纸质书形式出版。

版权贸易合同登记号 图字: 01-2012-7754

图书在版编目(CIP)数据

数据结构与算法分析: C++ 版: 第 3 版 = Data Structures and Algorithm Analysis in C++: 英文/(美)谢弗(Shaffer, C. A.)著. —北京: 电子工业出版社, 2013.1

国外计算机科学教材系列

ISBN 978-7-121-19260-9

I. ①数… II. ①谢… III. ①数据结构-高等学校-教材-英文 ②算法分析-高等学校-教材-英文 ③C 语言-程序设计-高等学校-教材-英文 IV. ①TP311.12 ②TP312

中国版本图书馆 CIP 数据核字(2012)第 302702 号

策划编辑: 马 岚

责任编辑: 马 岚

印 刷: 涿州市京南印刷厂

装 订: 涿州市京南印刷厂

出版发行: 电子工业出版社

北京市海淀区万寿路 173 信箱 邮编 100036

开 本: 787×1092 1/16 印张: 38.25 字数: 1273 千字

印 次: 2013 年 1 月第 1 次印刷

定 价: 69.00 元

凡所购买电子工业出版社图书有缺损问题, 请向购买书店调换。若书店售缺, 请与本社发行部联系, 联系及邮购电话: (010)88254888。

质量投诉请发邮件至 zltz@phei.com.cn, 盗版侵权举报请发邮件至 dbqq@phei.com.cn。

服务热线: (010)88258888。

Preface

We study data structures so that we can learn to write more efficient programs. But why must programs be efficient when new computers are faster every year? The reason is that our ambitions grow with our capabilities. Instead of rendering efficiency needs obsolete, the modern revolution in computing power and storage capability merely raises the efficiency stakes as we attempt more complex tasks.

The quest for program efficiency need not and should not conflict with sound design and clear coding. Creating efficient programs has little to do with “programming tricks” but rather is based on good organization of information and good algorithms. A programmer who has not mastered the basic principles of clear design is not likely to write efficient programs. Conversely, concerns related to development costs and maintainability should not be used as an excuse to justify inefficient performance. Generality in design can and should be achieved without sacrificing performance, but this can only be done if the designer understands how to measure performance and does so as an integral part of the design and implementation process. Most computer science curricula recognize that good programming skills begin with a strong emphasis on fundamental software engineering principles. Then, once a programmer has learned the principles of clear program design and implementation, the next step is to study the effects of data organization and algorithms on program efficiency.

Approach: This book describes many techniques for representing data. These techniques are presented within the context of the following principles:

1. Each data structure and each algorithm has costs and benefits. Practitioners need a thorough understanding of how to assess costs and benefits to be able to adapt to new design challenges. This requires an understanding of the principles of algorithm analysis, and also an appreciation for the significant effects of the physical medium employed (e.g., data stored on disk versus main memory).
2. Related to costs and benefits is the notion of tradeoffs. For example, it is quite common to reduce time requirements at the expense of an increase in space requirements, or vice versa. Programmers face tradeoff issues regularly in all

phases of software design and implementation, so the concept must become deeply ingrained.

3. Programmers should know enough about common practice to avoid reinventing the wheel. Thus, programmers need to learn the commonly used data structures, their related algorithms, and the most frequently encountered design patterns found in programming.
4. Data structures follow needs. Programmers must learn to assess application needs first, then find a data structure with matching capabilities. To do this requires competence in Principles 1, 2, and 3.

As I have taught data structures through the years, I have found that design issues have played an ever greater role in my courses. This can be traced through the various editions of this textbook by the increasing coverage for design patterns and generic interfaces. The first edition had no mention of design patterns. The second edition had limited coverage of a few example patterns, and introduced the dictionary ADT and comparator classes. With the third edition, there is explicit coverage of some design patterns that are encountered when programming the basic data structures and algorithms covered in the book.

Using the Book in Class: Data structures and algorithms textbooks tend to fall into one of two categories: teaching texts or encyclopedias. Books that attempt to do both usually fail at both. This book is intended as a teaching text. I believe it is more important for a practitioner to understand the principles required to select or design the data structure that will best solve some problem than it is to memorize a lot of textbook implementations. Hence, I have designed this as a teaching text that covers most standard data structures, but not all. A few data structures that are not widely adopted are included to illustrate important principles. Some relatively new data structures that should become widely used in the future are included.

Within an undergraduate program, this textbook is designed for use in either an advanced lower division (sophomore or junior level) data structures course, or for a senior level algorithms course. New material has been added in the third edition to support its use in an algorithms course. Normally, this text would be used in a course beyond the standard freshman level “CS2” course that often serves as the initial introduction to data structures. Readers of this book should typically have two semesters of the equivalent of programming experience, including at least some exposure to C++. Readers who are already familiar with recursion will have an advantage. Students of data structures will also benefit from having first completed a good course in Discrete Mathematics. Nonetheless, Chapter 2 attempts to give a reasonably complete survey of the prerequisite mathematical topics at the level necessary to understand their use in this book. Readers may wish to refer back to the appropriate sections as needed when encountering unfamiliar mathematical material.

A sophomore-level class where students have only a little background in basic data structures or analysis (that is, background equivalent to what would be had from a traditional CS2 course) might cover Chapters 1-11 in detail, as well as selected topics from Chapter 13. That is how I use the book for my own sophomore-level class. Students with greater background might cover Chapter 1, skip most of Chapter 2 except for reference, briefly cover Chapters 3 and 4, and then cover chapters 5-12 in detail. Again, only certain topics from Chapter 13 might be covered, depending on the programming assignments selected by the instructor. A senior-level algorithms course would focus on Chapters 11 and 14-17.

Chapter 13 is intended in part as a source for larger programming exercises. I recommend that all students taking a data structures course be required to implement some advanced tree structure, or another dynamic structure of comparable difficulty such as the skip list or sparse matrix representations of Chapter 12. None of these data structures are significantly more difficult to implement than the binary search tree, and any of them should be within a student's ability after completing Chapter 5.

While I have attempted to arrange the presentation in an order that makes sense, instructors should feel free to rearrange the topics as they see fit. The book has been written so that once the reader has mastered Chapters 1-6, the remaining material has relatively few dependencies. Clearly, external sorting depends on understanding internal sorting and disk files. Section 6.2 on the UNION/FIND algorithm is used in Kruskal's Minimum-Cost Spanning Tree algorithm. Section 9.2 on self-organizing lists mentions the buffer replacement schemes covered in Section 8.3. Chapter 14 draws on examples from throughout the book. Section 17.2 relies on knowledge of graphs. Otherwise, most topics depend only on material presented earlier within the same chapter.

Most chapters end with a section entitled "Further Reading." These sections are not comprehensive lists of references on the topics presented. Rather, I include books and articles that, in my opinion, may prove exceptionally informative or entertaining to the reader. In some cases I include references to works that should become familiar to any well-rounded computer scientist.

Use of C++: The programming examples are written in C++, but I do not wish to discourage those unfamiliar with C++ from reading this book. I have attempted to make the examples as clear as possible while maintaining the advantages of C++. C++ is used here strictly as a tool to illustrate data structures concepts. In particular, I make use of C++'s support for hiding implementation details, including features such as classes, private class members, constructors, and destructors. These features of the language support the crucial concept of separating logical design, as embodied in the abstract data type, from physical implementation as embodied in the data structure.

To keep the presentation as clear as possible, some important features of C++ are avoided here. I deliberately minimize use of certain features commonly used by experienced C++ programmers such as class hierarchy, inheritance, and virtual functions. Operator and function overloading is used sparingly. C-like initialization syntax is preferred to some of the alternatives offered by C++.

While the C++ features mentioned above have valid design rationale in real programs, they tend to obscure rather than enlighten the principles espoused in this book. For example, inheritance is an important tool that helps programmers avoid duplication, and thus minimize bugs. From a pedagogical standpoint, however, inheritance often makes code examples harder to understand since it tends to spread the description for one logical unit among several classes. Thus, my class definitions only use inheritance where inheritance is explicitly relevant to the point illustrated (e.g., Section 5.3.1). This does not mean that a programmer should do likewise. Avoiding code duplication and minimizing errors are important goals. Treat the programming examples as illustrations of data structure principles, but do not copy them directly into your own programs.

One painful decision I had to make was whether to use templates in the code examples. In the first edition of this book, the decision was to leave templates out as it was felt that their syntax obscures the meaning of the code for those not familiar with C++. In the years following, the use of C++ in computer science curricula has greatly expanded. I now assume that readers of the text will be familiar with template syntax. Thus, templates are now used extensively in the code examples.

My implementations are meant to provide concrete illustrations of data structure principles, as an aid to the textual exposition. Code examples should not be read or used in isolation from the associated text because the bulk of each example's documentation is contained in the text, not the code. The code complements the text, not the other way around. They are not meant to be a series of commercial-quality class implementations. If you are looking for a complete implementation of a standard data structure for use in your own code, you would do well to do an Internet search.

For instance, the code examples provide less parameter checking than is sound programming practice, since including such checking would obscure rather than illuminate the text. Some parameter checking and testing for other constraints (e.g., whether a value is being removed from an empty container) is included in the form of a call to **Assert**. The inputs to **Assert** are a Boolean expression and a character string. If this expression evaluates to **false**, then a message is printed and the program terminates immediately. Terminating a program when a function receives a bad parameter is generally considered undesirable in real programs, but is quite adequate for understanding how a data structure is meant to operate. In real programming applications, C++'s exception handling features should be used to deal with input data errors. However, assertions provide a simpler mechanism for indi-

cating required conditions in a way that is both adequate for clarifying how a data structure is meant to operate, and is easily modified into true exception handling. See the Appendix for the implementation of **Assert**.

I make a distinction in the text between “C++ implementations” and “pseudocode.” Code labeled as a C++ implementation has actually been compiled and tested on one or more C++ compilers. Pseudocode examples often conform closely to C++ syntax, but typically contain one or more lines of higher-level description. Pseudocode is used where I perceived a greater pedagogical advantage to a simpler, but less precise, description.

Exercises and Projects: Proper implementation and analysis of data structures cannot be learned simply by reading a book. You must practice by implementing real programs, constantly comparing different techniques to see what really works best in a given situation.

One of the most important aspects of a course in data structures is that it is where students really learn to program using pointers and dynamic memory allocation, by implementing data structures such as linked lists and trees. It is often where students truly learn recursion. In our curriculum, this is the first course where students do significant design, because it often requires real data structures to motivate significant design exercises. Finally, the fundamental differences between memory-based and disk-based data access cannot be appreciated without practical programming experience. For all of these reasons, a data structures course cannot succeed without a significant programming component. In our department, the data structures course is one of the most difficult programming course in the curriculum.

Students should also work problems to develop their analytical abilities. I provide over 450 exercises and suggestions for programming projects. I urge readers to take advantage of them.

Contacting the Author and Supplementary Materials: A book such as this is sure to contain errors and have room for improvement. I welcome bug reports and constructive criticism. I can be reached by electronic mail via the Internet at **shaffer@vt.edu**. Alternatively, comments can be mailed to

Cliff Shaffer
Department of Computer Science
Virginia Tech
Blacksburg, VA 24061

The electronic posting of this book, along with a set of lecture notes for use in class can be obtained at^①

<http://www.cs.vt.edu/~shaffer/book.html>.

① 本书已根据作者于2012年11月19日更新的勘误表进行了更正。——编者注

The code examples used in the book are available at the same site. Online Web pages for Virginia Tech's sophomore-level data structures class can be found at

<http://courses.cs.vt.edu/~cs3114>.

This book was typeset by the author using L^AT_EX. The bibliography was prepared using B^IB_TE_X. The index was prepared using **makeindex**. The figures were mostly drawn with **Xfig**. Figures 3.1 and 9.10 were partially created using Mathematica.

Acknowledgments: It takes a lot of help from a lot of people to make a book. I wish to acknowledge a few of those who helped to make this book possible. I apologize for the inevitable omissions.

Virginia Tech helped make this whole thing possible through sabbatical research leave during Fall 1994, enabling me to get the project off the ground. My department heads during the time I have written the various editions of this book, Dennis Kafura and Jack Carroll, provided unwavering moral support for this project. Mike Keenan, Lenny Heath, and Jeff Shaffer provided valuable input on early versions of the chapters. I also wish to thank Lenny Heath for many years of stimulating discussions about algorithms and analysis (and how to teach both to students). Steve Edwards deserves special thanks for spending so much time helping me on various redesigns of the C++ and Java code versions for the second and third editions, and many hours of discussion on the principles of program design. Thanks to Layne Watson for his help with Mathematica, and to Bo Begole, Philip Isenhour, Jeff Nielsen, and Craig Struble for much technical assistance. Thanks to Bill McQuain, Mark Abrams and Dennis Kafura for answering lots of silly questions about C++ and Java.

I am truly indebted to the many reviewers of the various editions of this manuscript. For the first edition these reviewers included J. David Bezek (University of Evansville), Douglas Campbell (Brigham Young University), Karen Davis (University of Cincinnati), Vijay Kumar Garg (University of Texas – Austin), Jim Miller (University of Kansas), Bruce Maxim (University of Michigan – Dearborn), Jeff Parker (Agile Networks/Harvard), Dana Richards (George Mason University), Jack Tan (University of Houston), and Lixin Tao (Concordia University). Without their help, this book would contain many more technical errors and many fewer insights.

For the second edition, I wish to thank these reviewers: Gurdip Singh (Kansas State University), Peter Allen (Columbia University), Robin Hill (University of Wyoming), Norman Jacobson (University of California – Irvine), Ben Keller (Eastern Michigan University), and Ken Bosworth (Idaho State University). In addition, I wish to thank Neil Stewart and Frank J. Thesen for their comments and ideas for improvement.

Third edition reviewers included Randall Lechlitner (University of Houston, Clear Lake) and Brian C. Hipp (York Technical College). I thank them for their comments.

Prentice Hall was the original print publisher for the first and second editions. Without the hard work of many people there, none of this would be possible. Authors simply do not create printer-ready books on their own. Foremost thanks go to Kate Hargett, Petra Rector, Laura Steele, and Alan Apt, my editors over the years. My production editors, Irwin Zucker for the second edition, Kathleen Caren for the original C++ version, and Ed DeFelippis for the Java version, kept everything moving smoothly during that horrible rush at the end. Thanks to Bill Zobrist and Bruce Gregory (I think) for getting me into this in the first place. Others at Prentice Hall who helped me along the way include Truly Donovan, Linda Behrens, and Phyllis Bregman. Thanks to Tracy Dunkelberger for her help in returning the copyright to me, thus enabling the electronic future of this work. I am sure I owe thanks to many others at Prentice Hall for their help in ways that I am not even aware of.

I am thankful to Shelley Kronzek at Dover publications for her faith in taking on the print publication of this third edition. Much expanded, with both Java and C++ versions, and many inconsistencies corrected, I am confident that this is the best edition yet. But none of us really knows whether students will prefer a free online textbook or a low-cost, printed bound version. In the end, we believe that the two formats will be mutually supporting by offering more choices. Production editor James Miller and design manager Marie Zaczekiewicz have worked hard to ensure that the production is of the highest quality.

I wish to express my appreciation to Hanan Samet for teaching me about data structures. I learned much of the philosophy presented here from him as well, though he is not responsible for any problems with the result. Thanks to my wife Terry, for her love and support, and to my daughters Irena and Kate for pleasant diversions from working too hard. Finally, and most importantly, to all of the data structures students over the years who have taught me what is important and what should be skipped in a data structures course, and the many new insights they have provided. This book is dedicated to them.

Cliff Shaffer
Blacksburg, Virginia

Contents

I Preliminaries	1
1 Data Structures and Algorithms	3
1.1 A Philosophy of Data Structures	4
1.1.1 The Need for Data Structures	4
1.1.2 Costs and Benefits	6
1.2 Abstract Data Types and Data Structures	8
1.3 Design Patterns	12
1.3.1 Flyweight	13
1.3.2 Visitor	13
1.3.3 Composite	14
1.3.4 Strategy	15
1.4 Problems, Algorithms, and Programs	16
1.5 Further Reading	18
1.6 Exercises	20
2 Mathematical Preliminaries	25
2.1 Sets and Relations	25
2.2 Miscellaneous Notation	29
2.3 Logarithms	31
2.4 Summations and Recurrences	32
2.5 Recursion	36
2.6 Mathematical Proof Techniques	38
2.6.1 Direct Proof	39
2.6.2 Proof by Contradiction	39

2.6.3	Proof by Mathematical Induction	40
2.7	Estimation	46
2.8	Further Reading	47
2.9	Exercises	48
3	Algorithm Analysis	55
3.1	Introduction	55
3.2	Best, Worst, and Average Cases	61
3.3	A Faster Computer, or a Faster Algorithm?	62
3.4	Asymptotic Analysis	65
3.4.1	Upper Bounds	65
3.4.2	Lower Bounds	67
3.4.3	Θ Notation	68
3.4.4	Simplifying Rules	69
3.4.5	Classifying Functions	70
3.5	Calculating the Running Time for a Program	71
3.6	Analyzing Problems	76
3.7	Common Misunderstandings	77
3.8	Multiple Parameters	79
3.9	Space Bounds	80
3.10	Speeding Up Your Programs	82
3.11	Empirical Analysis	85
3.12	Further Reading	86
3.13	Exercises	86
3.14	Projects	90
II	Fundamental Data Structures	93
4	Lists, Stacks, and Queues	95
4.1	Lists	96
4.1.1	Array-Based List Implementation	100
4.1.2	Linked Lists	103
4.1.3	Comparison of List Implementations	112
4.1.4	Element Implementations	114
4.1.5	Doubly Linked Lists	115

4.2	Stacks	120
4.2.1	Array-Based Stacks	121
4.2.2	Linked Stacks	124
4.2.3	Comparison of Array-Based and Linked Stacks	125
4.2.4	Implementing Recursion	125
4.3	Queues	129
4.3.1	Array-Based Queues	129
4.3.2	Linked Queues	134
4.3.3	Comparison of Array-Based and Linked Queues	134
4.4	Dictionaries	134
4.5	Further Reading	145
4.6	Exercises	145
4.7	Projects	149
5	Binary Trees	151
5.1	Definitions and Properties	151
5.1.1	The Full Binary Tree Theorem	153
5.1.2	A Binary Tree Node ADT	155
5.2	Binary Tree Traversals	155
5.3	Binary Tree Node Implementations	160
5.3.1	Pointer-Based Node Implementations	160
5.3.2	Space Requirements	166
5.3.3	Array Implementation for Complete Binary Trees	168
5.4	Binary Search Trees	168
5.5	Heaps and Priority Queues	178
5.6	Huffman Coding Trees	185
5.6.1	Building Huffman Coding Trees	186
5.6.2	Assigning and Using Huffman Codes	192
5.6.3	Search in Huffman Trees	195
5.7	Further Reading	196
5.8	Exercises	196
5.9	Projects	200
6	Non-Binary Trees	203
6.1	General Tree Definitions and Terminology	203
6.1.1	An ADT for General Tree Nodes	204

6.1.2	General Tree Traversals	205
6.2	The Parent Pointer Implementation	207
6.3	General Tree Implementations	213
6.3.1	List of Children	214
6.3.2	The Left-Child/Right-Sibling Implementation	215
6.3.3	Dynamic Node Implementations	215
6.3.4	Dynamic “Left-Child/Right-Sibling” Implementation	218
6.4	K -ary Trees	218
6.5	Sequential Tree Implementations	219
6.6	Further Reading	223
6.7	Exercises	223
6.8	Projects	226
III	Sorting and Searching	229
7	Internal Sorting	231
7.1	Sorting Terminology and Notation	232
7.2	Three $\Theta(n^2)$ Sorting Algorithms	233
7.2.1	Insertion Sort	233
7.2.2	Bubble Sort	235
7.2.3	Selection Sort	237
7.2.4	The Cost of Exchange Sorting	238
7.3	Shellsort	239
7.4	Mergesort	241
7.5	Quicksort	244
7.6	Heapsort	251
7.7	Binsort and Radix Sort	252
7.8	An Empirical Comparison of Sorting Algorithms	259
7.9	Lower Bounds for Sorting	261
7.10	Further Reading	265
7.11	Exercises	265
7.12	Projects	269
8	File Processing and External Sorting	273
8.1	Primary versus Secondary Storage	273

8.2	Disk Drives	276
8.2.1	Disk Drive Architecture	276
8.2.2	Disk Access Costs	280
8.3	Buffers and Buffer Pools	282
8.4	The Programmer's View of Files	290
8.5	External Sorting	291
8.5.1	Simple Approaches to External Sorting	294
8.5.2	Replacement Selection	296
8.5.3	Multiway Merging	300
8.6	Further Reading	303
8.7	Exercises	304
8.8	Projects	307
9	Searching	311
9.1	Searching Unsorted and Sorted Arrays	312
9.2	Self-Organizing Lists	317
9.3	Bit Vectors for Representing Sets	323
9.4	Hashing	324
9.4.1	Hash Functions	325
9.4.2	Open Hashing	330
9.4.3	Closed Hashing	331
9.4.4	Analysis of Closed Hashing	339
9.4.5	Deletion	344
9.5	Further Reading	345
9.6	Exercises	345
9.7	Projects	348
10	Indexing	351
10.1	Linear Indexing	353
10.2	ISAM	356
10.3	Tree-based Indexing	358
10.4	2-3 Trees	360
10.5	B-Trees	364
10.5.1	B ⁺ -Trees	368
10.5.2	B-Tree Analysis	374
10.6	Further Reading	375

10.7 Exercises	375
10.8 Projects	377
IV Advanced Data Structures	379
11 Graphs	381
11.1 Terminology and Representations	382
11.2 Graph Implementations	386
11.3 Graph Traversals	390
11.3.1 Depth-First Search	393
11.3.2 Breadth-First Search	394
11.3.3 Topological Sort	394
11.4 Shortest-Paths Problems	399
11.4.1 Single-Source Shortest Paths	400
11.5 Minimum-Cost Spanning Trees	402
11.5.1 Prim's Algorithm	404
11.5.2 Kruskal's Algorithm	407
11.6 Further Reading	409
11.7 Exercises	409
11.8 Projects	411
12 Lists and Arrays Revisited	413
12.1 Multilists	413
12.2 Matrix Representations	416
12.3 Memory Management	420
12.3.1 Dynamic Storage Allocation	422
12.3.2 Failure Policies and Garbage Collection	429
12.4 Further Reading	433
12.5 Exercises	434
12.6 Projects	435
13 Advanced Tree Structures	437
13.1 Tries	437
13.2 Balanced Trees	442
13.2.1 The AVL Tree	443
13.2.2 The Splay Tree	445

13.3	Spatial Data Structures	448
13.3.1	The K-D Tree	450
13.3.2	The PR quadtree	455
13.3.3	Other Point Data Structures	459
13.3.4	Other Spatial Data Structures	461
13.4	Further Reading	461
13.5	Exercises	462
13.6	Projects	463
V	Theory of Algorithms	467
14	Analysis Techniques	469
14.1	Summation Techniques	470
14.2	Recurrence Relations	475
14.2.1	Estimating Upper and Lower Bounds	475
14.2.2	Expanding Recurrences	478
14.2.3	Divide and Conquer Recurrences	480
14.2.4	Average-Case Analysis of Quicksort	482
14.3	Amortized Analysis	484
14.4	Further Reading	487
14.5	Exercises	487
14.6	Projects	491
15	Lower Bounds	493
15.1	Introduction to Lower Bounds Proofs	494
15.2	Lower Bounds on Searching Lists	496
15.2.1	Searching in Unsorted Lists	496
15.2.2	Searching in Sorted Lists	498
15.3	Finding the Maximum Value	499
15.4	Adversarial Lower Bounds Proofs	501
15.5	State Space Lower Bounds Proofs	504
15.6	Finding the i th Best Element	507
15.7	Optimal Sorting	509
15.8	Further Reading	512
15.9	Exercises	512
15.10	Projects	515