


Algorithms in C Parts 1-4:
Fundamentals, Data Structures, Sorting, Searching
(Third Edition)

算法 I-IV (C 实现) ——

基础、数据结构、排序和搜索

(第三版 · 影印版)

[美] Robert Sedgewick 著



Algorithms in C Parts 1-4: Fundamentals, Data Structures, Sorting, Searching
(Third Edition) (ISBN 0-201-31452-5)

Robert Sedgewick

Copyright © 1998 by Addison-Wesley Publishing Company, Inc.

Original English Language Edition Published by Addison-Wesley Publishing Company, Inc.

All rights reserved.

Reprinting edition published by PEARSON EDUCATION ASIA LTD and CHINA ELECTRIC POWER PRESS,
Copyright © 2003.

本书影印版由 Pearson Education 授权中国电力出版社在中国境内（香港、澳门特别行政区和台湾地区除外）独家出版、发行。

未经出版者书面许可，不得以任何方式复制或抄袭本书的任何部分。

本书封面贴有 Pearson Education 防伪标签，无标签者不得销售。

For sale and distribution in the People's Republic of China exclusively (except Taiwan, Hong Kong SAR and Macao SAR).

仅限于中华人民共和国境内（不包括中国香港、澳门特别行政区和中国台湾地区）销售发行。

北京市版权局著作合同登记号：图字：01-2003-2441

图书在版编目（CIP）数据

算法 I-IV（C 实现）——基础、数据结构、排序和搜索：第 3 版 /（美）塞奇威克著。

—影印本。—北京：中国电力出版社，2003

（原版风暴系列）

ISBN 7-5083-1481-6

I. 算... II. 塞... III. ①电子计算机—算法设计—英文 ②C 语言—程序设计—英文 IV. TP301.6

中国版本图书馆 CIP 数据核字（2003）第 069623 号

责任编辑：朱恩从

丛 书 名：原版风暴系列

书 名：算法 I-IV（C 实现）——基础、数据结构、排序和搜索（第三版·影印版）

编 著：（美）Robert Sedgewick

出版发行：中国电力出版社

地址：北京市三里河路6号 邮政编码：100044

电话：（010）88515918 传真：（010）88518169

本书如有印装质量问题，我社负责退换

印 刷：汇鑫印务有限公司

发 行 者：新华书店总店北京发行所

开 本：787×1092 1/16 **印 张：**45.5

书 号：ISBN 7-5083-1481-6

版 次：2003年11月北京第一版

印 次：2003年11月第一次印刷

定 价：70.00 元

版权所有，翻印必究

Preface

THIS BOOK IS intended to survey the most important computer algorithms in use today, and to teach fundamental techniques to the growing number of people in need of knowing them. It can be used as a textbook for a second, third, or fourth course in computer science, after students have acquired basic programming skills and familiarity with computer systems, but before they have taken specialized courses in advanced areas of computer science or computer applications. The book also may be useful for self-study or as a reference for people engaged in the development of computer systems or applications programs, since it contains implementations of useful algorithms and detailed information on these algorithms' performance characteristics. The broad perspective taken makes the book an appropriate introduction to the field.

I have completely rewritten the text for this new edition, and I have added more than a thousand new exercises, more than a hundred new figures, and dozens of new programs. I have also added detailed commentary on all the figures and programs. This new material provides both coverage of new topics and fuller explanations of many of the classic algorithms. A new emphasis on abstract data types throughout the book makes the programs more broadly useful and relevant in modern object-oriented programming environments. People who have read old editions of the book will find a wealth of new information throughout; all readers will find a wealth of pedagogical material that provides effective access to essential concepts.

Due to the large amount of new material, we have split the new edition into two volumes (each about the size of the old edition) of which this is the first. This volume covers fundamental concepts, data structures, sorting algorithms, and searching algorithms; the second volume covers advanced algorithms and applications, building on the basic abstractions and methods developed here. Nearly all the material on fundamentals and data structures in this edition is new.

This book is not just for programmers and computer-science students. Nearly everyone who uses a computer wants it to run faster or to solve larger problems. The algorithms in this book represent a body of knowledge developed over the last 50 years that has become indispensable in the efficient use of the computer, for a broad variety of applications. From N -body simulation problems in physics to genetic-sequencing problems in molecular biology, the basic methods described here have become essential in scientific research; and from database systems to Internet search engines, they have become essential parts of modern software systems. As the scope of computer applications becomes more widespread, so grows the impact of many of the basic methods covered here. The goal of this book is to serve as a resource for students and professionals interested in knowing and making intelligent use of these fundamental algorithms as basic tools for whatever computer application they might undertake.

Scope

The book contains 16 chapters grouped into four major parts: fundamentals, data structures, sorting, and searching. The descriptions here are intended to give readers an understanding of the basic properties of as broad a range of fundamental algorithms as possible. Ingenious methods ranging from binomial queues to patricia tries are described, all related to basic paradigms at the heart of computer science. The second volume consists of four additional parts that cover strings, geometry, graphs, and advanced topics. My primary goal in developing these books has been to bring together the fundamental methods from these diverse areas, to provide access to the best methods known for solving problems by computer.

You will most appreciate the material in this book if you have had one or two previous courses in computer science or have had equivalent programming experience: one course in programming in a high-level language such as C, Java, or C++, and perhaps another course that teaches fundamental concepts of programming systems. This book is thus intended for anyone conversant with a modern programming language and with the basic features of modern computer systems. References that might help to fill in gaps in your background are suggested in the text.

Most of the mathematical material supporting the analytic results is self-contained (or is labeled as beyond the scope of this book), so little specific preparation in mathematics is required for the bulk of the book, although mathematical maturity is definitely helpful.

Use in the Curriculum

There is a great deal of flexibility in how the material here can be taught, depending on the taste of the instructor and the preparation of the students. The algorithms described here have found widespread use for years, and represent an essential body of knowledge for both the practicing programmer and the computer-science student. There is sufficient coverage of basic material for the book to be used for a course on data structures, and there is sufficient detail and coverage of advanced material for the book to be used for a course on algorithms. Some instructors may wish to emphasize implementations and practical concerns; others may wish to emphasize analysis and theoretical concepts.

A complete set of slide masters for use in lectures, sample programming assignments, interactive exercises for students, and other course materials may be found via the book's home page.

An elementary course on data structures and algorithms might emphasize the basic data structures in Part 2 and their use in the implementations in Parts 3 and 4. A course on design and analysis of algorithms might emphasize the fundamental material in Part 1 and Chapter 5, then study the ways in which the algorithms in Parts 3 and 4 achieve good asymptotic performance. A course on software engineering might omit the mathematical and advanced algorithmic material, and emphasize how to integrate the implementations given here into large programs or systems. A course on algorithms might take a survey approach and introduce concepts from all these areas.

Earlier editions of this book have been used in recent years at scores of colleges and universities around the world as a text for the second or third course in computer science and as supplemental reading for other courses. At Princeton, our experience has been that the breadth of coverage of material in this book provides our majors with an introduction to computer science that can be expanded upon in later courses on analysis of algorithms, systems programming and

theoretical computer science, while providing the growing group of students from other disciplines with a large set of techniques that these people can immediately put to good use.

The exercises—most of which are new to this edition—fall into several types. Some are intended to test understanding of material in the text, and simply ask readers to work through an example or to apply concepts described in the text. Others involve implementing and putting together the algorithms, or running empirical studies to compare variants of the algorithms and to learn their properties. Still others are a repository for important information at a level of detail that is not appropriate for the text. Reading and thinking about the exercises will pay dividends for every reader.

Algorithms of Practical Use

Anyone wanting to use a computer more effectively can use this book for reference or for self-study. People with programming experience can find information on specific topics throughout the book. To a large extent, you can read the individual chapters in the book independently of the others, although, in some cases, algorithms in one chapter make use of methods from a previous chapter.

The orientation of the book is to study algorithms likely to be of practical use. The book provides information about the tools of the trade to the point that readers can confidently implement, debug, and put to work algorithms to solve a problem or to provide functionality in an application. Full implementations of the methods discussed are included, as are descriptions of the operations of these programs on a consistent set of examples. Because we work with real code, rather than write pseudo-code, the programs can be put to practical use quickly. Program listings are available from the book's home page.

Indeed, one practical application of the algorithms has been to produce the hundreds of figures throughout the book. Many algorithms are brought to light on an intuitive level through the visual dimension provided by these figures.

Characteristics of the algorithms and of the situations in which they might be useful are discussed in detail. Although not emphasized, connections to the analysis of algorithms and theoretical computer science are developed in context. When appropriate, empirical and

analytic results are presented to illustrate why certain algorithms are preferred. When interesting, the relationship of the practical algorithms being discussed to purely theoretical results is described. Specific information on performance characteristics of algorithms and implementations is synthesized, encapsulated, and discussed throughout the book.

Programming Language

The programming language used for all of the implementations is C. Any particular language has advantages and disadvantages; we use C because it is widely available and provides the features needed for our implementations. The programs can be translated easily to other modern programming languages, since relatively few constructs are unique to C. We use standard C idioms when appropriate, but this book is not intended to be a reference work on C programming.

There are many new programs in this edition, and many of the old ones have been reworked, primarily to make them more readily useful as abstract-data-type implementations. Extensive comparative empirical tests on the programs are discussed throughout the text.

Previous editions of the book have presented basic programs in Pascal, C++, and Modula-3. This code is available through the book home page on the web; code for new programs and code in new languages such as Java will be added as appropriate.

A goal of this book is to present the algorithms in as simple and direct a form as possible. The style is consistent whenever possible, so that programs that are similar look similar. For many of the algorithms in this book, the similarities hold regardless of the language: Quicksort is quicksort (to pick one prominent example), whether expressed in Algol-60, Basic, Fortran, Smalltalk, Ada, Pascal, C, PostScript, Java, or countless other programming languages and environments where it has proved to be an effective sorting method.

We strive for elegant, compact, and portable implementations, but we take the point of view that efficiency matters, so we try to be aware of the performance characteristics of our code at all stages of development. Chapter 1 constitutes a detailed example of this approach to developing efficient C implementations of our algorithms, and sets the stage for the rest of the book.

Acknowledgments

Many people gave me helpful feedback on earlier versions of this book. In particular, hundreds of students at Princeton and Brown have suffered through preliminary drafts over the years. Special thanks are due to Trina Avery and Tom Freeman for their help in producing the first edition; to Janet Incerpi for her creativity and ingenuity in persuading our early and primitive digital computerized typesetting hardware and software to produce the first edition; to Marc Brown for his part in the algorithm visualization research that was the genesis of so many of the figures in the book; and to Dave Hanson for his willingness to answer all of my questions about C. I would also like to thank the many readers who have provided me with detailed comments about various editions, including Guy Almes, Jon Bentley, Marc Brown, Jay Gischer, Allan Heydon, Kennedy Lemke, Udi Manber, Dana Richards, John Reif, M. Rosenfeld, Stephen Seidman, Michael Quinn, and William Ward.

To produce this new edition, I have had the pleasure of working with Peter Gordon and Debbie Lafferty at Addison-Wesley, who have patiently shepherded this project as it has evolved from a standard update to a massive rewrite. It has also been my pleasure to work with several other members of the professional staff at Addison-Wesley. The nature of this project made the book a somewhat unusual challenge for many of them, and I much appreciate their forbearance.

I have gained two new mentors in writing this book, and particularly want to express my appreciation to them. First, Steve Summit carefully checked early versions of the manuscript on a technical level, and provided me with literally thousands of detailed comments, particularly on the programs. Steve clearly understood my goal of providing elegant, efficient, and effective implementations, and his comments not only helped me to provide a measure of consistency across the implementations, but also helped me to improve many of them substantially. Second, Lyn Dupre also provided me with thousands of detailed comments on the manuscript, which were invaluable in helping me not only to correct and avoid grammatical errors, but also—more important—to find a consistent and coherent writing style that helps bind together the daunting mass of technical material here. I am extremely grateful

for the opportunity to learn from Steve and Lyn—their input was vital in the development of this book.

Much of what I have written here I have learned from the teaching and writings of Don Knuth, my advisor at Stanford. Although Don had no direct influence on this work, his presence may be felt in the book, for it was he who put the study of algorithms on the scientific footing that makes a work such as this possible. My friend and colleague Philippe Flajolet, who has been a major force in the development of the analysis of algorithms as a mature research area, has had a similar influence on this work.

I am deeply thankful for the support of Princeton University, Brown University, and the Institut National de Recherche en Informatique et Automatique (INRIA), where I did most of the work on the book; and of the Institute for Defense Analyses and the Xerox Palo Alto Research Center, where I did some work on the book while visiting. Many parts of the book are dependent on research that has been generously supported by the National Science Foundation and the Office of Naval Research. Finally, I thank Bill Bowen, Aaron Lemonick, and Neil Rudenstine for their support in building an academic environment at Princeton in which I was able to prepare this book, despite my numerous other responsibilities.

Robert Sedgewick
Marly-le-Roi, France, February, 1983
Princeton, New Jersey, January, 1990
Jamestown, Rhode Island, August, 1997

*To Adam, Andrew, Brett, Robbie,
and especially Linda*

Notes on Exercises

Classifying exercises is an activity fraught with peril, because readers of a book such as this come to the material with various levels of knowledge and experience. Nonetheless, guidance is appropriate, so many of the exercises carry one of four annotations, to help you decide how to approach them.

Exercises that *test your understanding* of the material are marked with an open triangle, as follows:

- ▷ 9.57 Give the binomial queue that results when the keys EASY QUESTION are inserted into an initially empty binomial queue.

Most often, such exercises relate directly to examples in the text. They should present no special difficulty, but working them might teach you a fact or concept that may have eluded you when you read the text.

Exercises that *add new and thought-provoking* information to the material are marked with an open circle, as follows:

- 14.20 Write a program that inserts N random integers into a table of size $N/100$ using separate chaining, then finds the length of the shortest and longest lists, for $N = 10^3, 10^4, 10^5$, and 10^6 .

Such exercises encourage you to think about an important concept that is related to the material in the text, or to answer a question that may have occurred to you when you read the text. You may find it worthwhile to read these exercises, even if you do not have the time to work them through.

Exercises that are intended to *challenge you* are marked with a black dot, as follows:

- 8.46 Suppose that mergesort is implemented to split the file at a *random* position, rather than exactly in the middle. How many comparisons are used by such a method to sort N elements, on the average?

Such exercises may require a substantial amount of time to complete, depending upon your experience. Generally, the most productive approach is to work on them in a few different sittings.

A few exercises that are *extremely difficult* (by comparison with most others) are marked with two black dots, as follows:

- 15.29 Prove that the height of a trie built from N random bit-strings is about $2 \lg N$.

These exercises are similar to questions that might be addressed in the research literature, but the material in the book may prepare you to enjoy trying to solve them (and perhaps succeeding).

The annotations are intended to be neutral with respect to your programming and mathematical ability. Those exercises that require expertise in programming or in mathematical analysis are self-evident. All readers are encouraged to test their understanding of the algorithms by implementing them. Still, an exercise such as this one is straightforward for a practicing programmer or a student in a programming course, but may require substantial work for someone who has not recently programmed:

1.23 Modify Program 1.4 to generate random pairs of integers between 0 and $N - 1$ instead of reading them from standard input, and to loop until $N - 1$ *union* operations have been performed. Run your program for $N = 10^3, 10^4, 10^5$, and 10^6 and print out the total number of edges generated for each value of N .

In a similar vein, all readers are encouraged to strive to appreciate the analytic underpinnings of our knowledge about properties of algorithms. Still, an exercise such as this one is straightforward for a scientist or a student in a discrete mathematics course, but may require substantial work for someone who has not recently done mathematical analysis:

1.13 Compute the *average* distance from a node to the root in a worst-case tree of 2^n nodes built by the weighted quick-union algorithm.

There are far too many exercises for you to read and assimilate them all; my hope is that there are enough exercises here to stimulate you to strive to come to a broader understanding on the topics that interest you than you can glean by simply reading the text.

Contents

Fundamentals

Chapter 1. Introduction **3**

- 1.1 Algorithms · 4
- 1.2 A Sample Problem—Connectivity · 6
- 1.3 Union-Find Algorithms · 11
- 1.4 Perspective · 22
- 1.5 Summary of Topics · 23

Chapter 2. Principles of Algorithm Analysis **27**

- 2.1 Implementation and Empirical Analysis · 28
- 2.2 Analysis of Algorithms · 33
- 2.3 Growth of Functions · 36
- 2.4 Big-Oh notation · 44
- 2.5 Basic Recurrences · 49
- 2.6 Examples of Algorithm Analysis · 53
- 2.7 Guarantees, Predictions, and Limitations · 59

Data Structures

Chapter 3. Elementary Data Structures	69
3.1 Building Blocks · 70	
3.2 Arrays · 82	
3.3 Linked Lists · 90	
3.4 Elementary List Processing · 96	
3.5 Memory Allocation for Lists · 105	
3.6 Strings · 109	
3.7 Compound Data Structures · 115	
 Chapter 4. Abstract Data Types	 127
4.1 Abstract Objects and Collections of Objects · 131	
4.2 Pushdown Stack ADT · 135	
4.3 Examples of Stack ADT Clients · 138	
4.4 Stack ADT Implementations · 144	
4.5 Creation of a New ADT · 149	
4.6 FIFO Queues and Generalized Queues · 153	
4.7 Duplicate and Index Items · 161	
4.8 First-Class ADTs · 165	
4.9 Application-Based ADT Example · 178	
4.10 Perspective · 184	
 Chapter 5. Recursion and Trees	 187
5.1 Recursive Algorithms · 188	
5.2 Divide and Conquer · 196	
5.3 Dynamic Programming · 208	
5.4 Trees · 216	
5.5 Mathematical Properties of Trees · 226	
5.6 Tree Traversal · 230	
5.7 Recursive Binary-Tree Algorithms · 235	
5.8 Graph Traversal · 241	
5.9 Perspective · 247	

Sorting

Chapter 6. Elementary Sorting Methods 253

- 6.1 Rules of the Game · 255
- 6.2 Selection Sort · 261
- 6.3 Insertion Sort · 262
- 6.4 Bubble Sort · 265
- 6.5 Performance Characteristics of Elementary Sorts · 267
- 6.6 Shellsort · 273
- 6.7 Sorting Other Types of Data · 281
- 6.8 Index and Pointer Sorting · 287
- 6.9 Sorting of Linked Lists · 294
- 6.10 Key-Indexed Counting · 298

Chapter 7. Quicksort 303

- 7.1 The Basic Algorithm · 304
- 7.2 Performance Characteristics of Quicksort · 309
- 7.3 Stack Size · 313
- 7.4 Small Subfiles · 316
- 7.5 Median-of-Three Partitioning · 319
- 7.6 Duplicate Keys · 324
- 7.7 Strings and Vectors · 327
- 7.8 Selection · 329

Chapter 8. Merging and Mergesort 335

- 8.1 Two-Way Merging · 336
- 8.2 Abstract In-place Merge · 339
- 8.3 Top-Down Mergesort · 341
- 8.4 Improvements to the Basic Algorithm · 344
- 8.5 Bottom-Up Mergesort · 347
- 8.6 Performance Characteristics of Mergesort · 351
- 8.7 Linked-List Implementations of Mergesort · 354
- 8.8 Recursion Revisited · 357

Chapter 9. Priority Queues and Heapsort 361

- 9.1 Elementary Implementations · 365
- 9.2 Heap Data Structure · 368

9.3 Algorithms on Heaps · 371	
9.4 Heapsort · 376	
9.5 Priority-Queue ADT · 383	
9.6 Priority Queues for Index Items · 389	
9.7 Binomial Queues · 392	
Chapter 10. Radix Sorting	403
10.1 Bits, Bytes, and Words · 405	
10.2 Binary Quicksort · 409	
10.3 MSD Radix Sort · 413	
10.4 Three-Way Radix Quicksort · 421	
10.5 LSD Radix Sort · 425	
10.6 Performance Characteristics of Radix Sorts · 428	
10.7 Sublinear-Time Sorts · 433	
Chapter 11. Special-Purpose Sorts	439
11.1 Batcher's Odd-Even Mergesort · 441	
11.2 Sorting Networks · 446	
11.3 External Sorting · 454	
11.4 Sort-Merge Implementations · 460	
11.5 Parallel Sort/Merge · 466	

Searching

Chapter 12. Symbol Tables and BSTs	477
12.1 Symbol-Table Abstract Data Type · 479	
12.2 Key-Indexed Search · 485	
12.3 Sequential Search · 489	
12.4 Binary Search · 497	
12.5 Binary Search Trees (BSTs) · 502	
12.6 Performance Characteristics of BSTs · 508	
12.7 Index Implementations with Symbol Tables · 511	
12.8 Insertion at the Root in BSTs · 516	
12.9 BST Implementations of Other ADT Functions · 519	

Chapter 13. Balanced Trees **529**

- 13.1 Randomized BSTs · 533
- 13.2 Splay BSTs · 540
- 13.3 Top-Down 2-3-4 Trees · 546
- 13.4 Red-Black Trees · 551
- 13.5 Skip Lists · 561
- 13.6 Performance Characteristics · 569

Chapter 14. Hashing **573**

- 14.1 Hash Functions · 574
- 14.2 Separate Chaining · 583
- 14.3 Linear Probing · 588
- 14.4 Double Hashing · 594
- 14.5 Dynamic Hash Tables · 599
- 14.6 Perspective · 603

Chapter 15. Radix Search **609**

- 15.1 Digital Search Trees · 610
- 15.2 Tries · 614
- 15.3 Patricia Tries · 623
- 15.4 Multiway Tries and TSTs · 632
- 15.5 Text String Index Algorithms · 648

Chapter 16. External Searching **655**

- 16.1 Rules of the Game · 657
- 16.2 Indexed Sequential Access · 660
- 16.3 B Trees · 662
- 16.4 Extendible Hashing · 676
- 16.5 Perspective · 688

Index **693**