

算法经典丛书

PEARSON  
Addison  
Wesley

# JAVA 算法 影印版

(第3版, 第1卷)

## Algorithms in Java Parts 1-4, 3e

[美] Robert Sedgewick 著

WITH JAVA CONSULTING BY MICHAEL SCHIDLOWSKY



清华大学出版社

算法经典丛书

# Java 算法

影印版（第3版，第1卷）

[美] Robert Sedgewick 著

清华大学出版社

· 北京 ·

## 内 容 简 介

《Java 算法》全面介绍了当今最重要的计算机算法在 Java 中的应用和实现。全书共分 3 卷, 第 1 卷内容包括基本概念(第 1 部分)、数据结构(第 2 部分)、排序算法(第 3 部分)和查找算法(第 4 部分)。

本书以 Java 语言作为算法描述语言, 易于理解、便于应用。可作高校计算机相关专业本科生和研究生的教材和补充读物, 也可供相关领域工程技术人员参考。

**EISBN: 0-201-36120-5**

**Algorithms in Java, Third Edition, Part 1-4**

**Robert Sedgewick**

**Copyright © 2003 by Pearson Education, Inc.**

**Original English language edition published by Pearson Education, Inc.**

**All right reserved.**

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

本书封面贴有 Pearson Education(培生教育出版集团)激光防伪标签, 无标签者不得销售。

北京市版权局著作权合同登记号: 图字 01-2003-0548 号

### 图书在版编目(CIP)数据

Java 算法. 第 1 卷 Algorithms in Java: 第 3 版/(美)西德威克著. 影印版.

—北京: 清华大学出版社, 2003

ISBN 7-302-06375-3

I.J... II.西... III.JAVA 语言—计算方法—英文 IV.TP312

中国版本图书馆 CIP 数据核字(2003)第 013033 号

出 版 者: 清华大学出版社(北京清华大学学研大厦, 邮编 100084)

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

<http://www.tup.tsinghua.edu.cn>

责任编辑: 尤晓东

印 刷 者: 北京牛山世兴印刷厂

发 行 者: 新华书店总店北京发行所

开 本: 787×960 1/16 印张: 48

版 次: 2003 年 3 月第 1 版 2003 年 3 月第 1 次印刷

书 号: ISBN 7-302-06375-3/TP·4810

印 数: 0001~3000

定 价: 68.00 元

# Preface

**T**HIS BOOK IS the first of three volumes that are intended to survey the most important computer algorithms in use today. This first volume (Parts 1–4) covers fundamental concepts (Part 1), data structures (Part 2), sorting algorithms (Part 3), and searching algorithms (Part 4); the second volume (Part 5) covers graphs and graph algorithms; and the (yet to be published) third volume (Parts 6–8) covers strings (Part 6), computational geometry (Part 7), and advanced algorithms and applications (Part 8).

The books are useful as texts early in the computer science curriculum, 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 books also are useful for self-study or as a reference for people engaged in the development of computer systems or applications programs because they contain implementations of useful algorithms and detailed information on these algorithms' performance characteristics. The broad perspective taken makes the series an appropriate introduction to the field.

Together the three volumes comprise the *Third Edition* of a book that has been widely used by students and programmers around the world for many years. I have completely rewritten the text for this edition, and I have added thousands of new exercises, hundreds of new figures, dozens of new programs, and 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 books makes the programs more broadly useful and relevant in modern object-oriented programming environments. People who have read previous editions will find a wealth of new information throughout; all readers will find a wealth of pedagogical material that provides effective access to essential concepts.

These books are not just for programmers and computer science students. Everyone who uses a computer wants it to run faster or to solve larger problems. The algorithms that we consider represent a body of knowledge developed during the last 50 years that is the basis for 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 basic algorithms. The goal of this book is to serve as a resource so that students and professionals can know and make intelligent use of these fundamental algorithms as the need arises in whatever computer application they might undertake.

## Scope

This book, *Algorithms in Java, Third Edition, Parts 1-4*, 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. 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. The second volume is devoted to graph algorithms, and the third consists of four additional parts that cover strings, geometry, and advanced topics. My primary goal in developing these books has been to bring together fundamental methods from these areas, to provide access to the best methods known for solving problems by computer.

You will most appreciate the material here 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 Java, C, 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. There is sufficient coverage of basic material for the book to be used to teach data structures to beginners, and there is sufficient detail and coverage of advanced material for the book to be used to teach the design and analysis of algorithms to upper-level students. Some instructors may wish to emphasize implementations and practical concerns; others may wish to emphasize analysis and theoretical concepts.

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 that are based on other programming languages have been used at scores of colleges and universities 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 on 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 put to good use immediately.

The exercises—nearly all of which are new to this third 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 algorithms to work 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, you can put the programs to practical use quickly. Program listings are available from the book's home page. You can use these working programs in many ways to help you study algorithms. Read them to check your understanding of the details of an algorithm, or to see one way to handle initializations, boundary conditions, and other awkward situations that often pose programming challenges. Run them to see the algorithms in action, to study performance empirically and check your results against the tables in the book, or to try your own modifications.

Characteristics of the algorithms and of the situations in which they might be useful are discussed in detail. Connections to the analysis of algorithms and theoretical computer science are developed in con-

text. 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 Java. The programs use a wide range of standard Java idioms, and the text includes concise descriptions of each construct.

Mike Schidlowsky and I developed a style of Java programming based on abstract data types that we feel is an effective way to present the algorithms and data structures as real programs. We have striven for elegant, compact, efficient, and portable implementations. The style is consistent whenever possible, so 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 Ada, Algol-60, Basic, C, C++, Fortran, Java, Mesa, Modula-3, Pascal, PostScript, Smalltalk, or countless other programming languages and environments where it has proved to be an effective sorting method. On the one hand, our code is informed by experience with implementing algorithms in these and numerous other languages (C and C++ versions of this book are also available); on the other hand, some of the properties of some of these languages are informed by their designers' experience with some of the algorithms and data structures that we consider in this book.

Chapter 1 constitutes a detailed example of this approach to developing efficient Java implementations of our algorithms, and Chapter 2 describes our approach to analyzing them. Chapters 3 and 4 are devoted to describing and justifying the basic mechanisms that we use for data type and ADT implementations. These four chapters set 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 and Andrew Appel for their willingness to answer all of my questions about programming languages. I would also like to thank the many readers who have provided me with 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 Helen Goldstein at Addison-Wesley, who have patiently shepherded this project as it has evolved. 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. In particular, Marilyn Rash did an outstanding job managing the book's production within a tightly compressed schedule.

I have gained three 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 Dupré 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. Third, Chris Van Wyk,

in a long series of spirited electronic mail exchanges, patiently defended the basic precepts of object-oriented programming and helped me develop a style of coding that exhibits the algorithms with clarity and precision while still taking advantage of what object-oriented programming has to offer. The basic approach that we developed for the C++ version of this book has substantially influenced the Java code here and will certainly influence future volumes in both languages (and C as well). I am extremely grateful for the opportunity to learn from Steve, Lyn, and Chris—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, 1983*  
*Princeton, New Jersey, 1990, 1992*  
*Jamestown, Rhode Island, 1997*  
*Princeton, New Jersey, 1998, 2002*

## Java Consultant's Preface

In the past decade, Java has become the language of choice for a variety of applications. But Java developers have found themselves repeatedly referring to references such as Sedgewick's *Algorithms in C* for solutions to common programming problems. There has long been an empty space on the bookshelf for a comparable reference work for Java; this book is here to fill that space.

We wrote the sample programs as utility methods to be used in a variety of contexts. To that end, we did not use the Java package mechanism. To focus on the algorithms at hand (and to expose the algorithmic basis of many fundamental library classes), we avoided the standard Java library in favor of more fundamental types. Proper error checking and other defensive practices would both substantially increase the amount of code and distract the reader from the core algorithms. Developers should introduce such code when using the programs in larger applications.

Although the algorithms we present are language independent, we have paid close attention to Java-specific performance issues. The timings throughout the book are provided as one context for comparing algorithms, and will vary depending on the virtual machine. As Java environments evolve, programs will perform as fast as natively compiled code, but such optimizations will not change the performance of algorithms relative to one another. We provide the timings as a useful reference for such comparisons.

I would like to thank Mike Zamansky, for his mentorship and devotion to the teaching of computer science, and Daniel Chaskes, Jason Sanders, and James Percy, for their unwavering support. I would also like to thank my family for their support and for the computer that bore my first programs. Bringing together Java with the classic algorithms of computer science was an exciting endeavor for which I am very grateful. Thank you, Bob, for the opportunity to do so.

*Michael Schidlowsky*  
*Oakland Gardens, New York, 2002*

---

*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 on 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 · 7
- 1.3 Union-Find Algorithms · 11
- 1.4 Perspective · 22
- 1.5 Summary of Topics · 24

### **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 · 60

## Data Structures

---

### Chapter 3. Elementary Data Structures 69

- 3.1 Building Blocks · 70
- 3.2 Arrays · 84
- 3.3 Linked Lists · 91
- 3.4 Elementary List Processing · 97
- 3.5 Memory Allocation for Lists · 107
- 3.6 Strings · 111
- 3.7 Compound Data Structures · 116

### Chapter 4. Abstract Data Types 127

- 4.1 Collections of Items · 137
- 4.2 Pushdown Stack ADT · 139
- 4.3 Examples of Stack ADT Clients · 142
- 4.4 Stack ADT Implementations · 148
- 4.5 Generic Implementations · 154
- 4.6 Creation of a New ADT · 157
- 4.7 FIFO Queues and Generalized Queues · 165
- 4.8 Duplicate and Index Items · 173
- 4.9 First-Class ADTs · 177
- 4.10 Application-Based ADT Example · 188
- 4.11 Perspective · 194

### Chapter 5. Recursion and Trees 197

- 5.1 Recursive Algorithms · 198
- 5.2 Divide and Conquer · 206
- 5.3 Dynamic Programming · 219
- 5.4 Trees · 227
- 5.5 Mathematical Properties of Trees · 236
- 5.6 Tree Traversal · 240
- 5.7 Recursive Binary-Tree Algorithms · 246
- 5.8 Graph Traversal · 251
- 5.9 Perspective · 257



# Sorting

---

## Chapter 6. Elementary Sorting Methods 263

- 6.1 Rules of the Game · 265
- 6.2 Generic Sort Implementations · 270
- 6.3 Selection Sort · 283
- 6.4 Insertion Sort · 285
- 6.5 Bubble Sort · 288
- 6.6 Performance Characteristics of Elementary Sorts · 289
- 6.7 Algorithm Visualization · 295
- 6.8 Shellsort · 300
- 6.9 Sorting Linked Lists · 308
- 6.10 Key-Indexed Counting · 312

## Chapter 7. Quicksort 315

- 7.1 The Basic Algorithm · 316
- 7.2 Performance Characteristics of Quicksort · 321
- 7.3 Stack Size · 325
- 7.4 Small Subfiles · 328
- 7.5 Median-of-Three Partitioning · 331
- 7.6 Duplicate Keys · 336
- 7.7 Strings and Vectors · 339
- 7.8 Selection · 341

## Chapter 8. Merging and Mergesort 347

- 8.1 Two-Way Merging · 348
- 8.2 Abstract In-Place Merge · 351
- 8.3 Top-Down Mergesort · 353
- 8.4 Improvements to the Basic Algorithm · 357
- 8.5 Bottom-Up Mergesort · 359
- 8.6 Performance Characteristics of Mergesort · 363
- 8.7 Linked-List Implementations of Mergesort · 366
- 8.8 Recursion Revisited · 370

## Chapter 9. Priority Queues and Heapsort 373

- 9.1 Elementary Implementations · 377
- 9.2 Heap Data Structure · 381