


Data Structures

RICK DECKER



TP33
D295

9161849

Data Structures

RICK DECKER

*Hamilton College
Clinton, New York*



PRENTICE HALL, Englewood Cliffs, New Jersey 07632

Library of Congress Cataloging-in-Publication Data

Decker, Rick

Data structures / Rick Decker.

p. cm.

Includes index.

ISBN 0-13-198813-1

1. Data structures (Computer science) I. Title.

QA76.9.D35D43. 1989

005.7'3—dc19

88-37633

CIP#

Editorial/production supervision and
interior design: David Ershun

Cover design: Diane Saxe

Cover art: Thomas Cole: *The Voyage of Life: Youth*, 1840. Oil on canvas,
52½ × 78½ inches. Credit: The Munson Williams Proctor Institute
Museum of Art, Utica, New York.

Manufacturing buyer: Mary Noonan



© 1989 by Prentice-Hall, Inc.
A Division of Simon & Schuster
Englewood Cliffs, New Jersey 07632

All rights reserved. No part of this book may be
reproduced, in any form or by any means,
without permission in writing from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-198813-1

PRENTICE-HALL INTERNATIONAL (UK) LIMITED, *London*
PRENTICE-HALL OF AUSTRALIA PTY. LIMITED, *Sydney*
PRENTICE-HALL CANADA INC., *Toronto*
PRENTICE-HALL HISPANOAMERICANA, S.A., *Mexico*
PRENTICE-HALL OF INDIA PRIVATE LIMITED, *New Delhi*
PRENTICE-HALL OF JAPAN INC., *Tokyo*
SIMON & SCHUSTER ASIA PTE. LTD., *Singapore*
EDITORIA PRENTICE-HALL DO BRASIL LTDA., *Rio de Janeiro*

Preface

The overwhelming majority of authors have very little to say. If we suppose, rather charitably, that in a typical book of fifteen chapters there are only eight passages worthy of quotation, then simple mathematics will convince us that in short order there will be no original quotations left for chapter headings. The implication is obvious. ...

ARMAND BLAGUE, *How to Write*

Over the years, I've had a number of students who have said, in one form or another, "I want to be a computer scientist because I really like programming and am very good at it." Of course, computer scientists, both novices and seasoned veterans, are often called upon to write programs, but to equate computer science with programming is to confuse the product with the process. Being an excellent draftsman who can faithfully represent a scene on paper is no guarantee that your works will eventually hang in the Metropolitan Museum. It's a step in the right direction, but an artist must also have an intimate familiarity with the more general principles of composition, perspective, color, and so on.

In essence, programming is little more than the efficient management of a particular kind of large intellectual process, and the guidelines for good programming are nothing but the application of commonsense principles that apply to any complex creative task. It goes without saying, however, that before you can think efficiently you must have something to think *about*, which for our purposes means that in order to write good programs, you must have an idea about how information may be represented in a program.

Computer science is a young discipline, but it has developed enough over the past few decades that there is a growing consensus about what should constitute the core data structures. In this book, I tried to capture this core by providing what might be called the "classic" data structures, the most commonly applied methods for

representing information in a computer program, along with the algorithms for manipulating this information. In terms of something to think *about* when thinking about programming, you have here a collection of tools that should be part of the working knowledge of any programmer.

This book is not about programming, however. Computer science is a science and, in common with other sciences, consists in the main of seeking a theoretical framework that can be used to describe the behavior of the objects it studies—in our case, computers and their programs. One of the themes that determined the form of the book is to provide a broad view of what a data structure really is. I have taken the approach that data structures are not just a collection of ad hoc declarations and procedure definitions, but, rather, that any data structure is a particular instance of an *abstract data type*, which consists of (1) a set of objects, with (2) a common logical structure defined for each of the objects, along with (3) a collection of structure-preserving operations on the objects.

I chose to define the structure of an abstract data type by specifying a *structural relation* on each set of positions. Doing so provides a natural progression of the chapters, where each new abstract data type is introduced by removing some of the structural restrictions from a prior type. Thus, we begin with lists, whose structure is defined by a linear order, and progress to trees by removing the requirement that each position have a unique successor, then to directed graphs by removing the requirement of unique predecessor, then to sets by removing all restrictions on the structural relation. Throughout this process, we see that each new abstract data type still can be described by the threefold view of a collection of sets with a structural relation and a collection of operations on those sets.

THE AUDIENCE

Though I did not set out to tailor this book to any preexisting curriculum, it turned out that it covers essentially all of CS2 and part of CS7, as described in the ACM *Curriculum '78*, and a subset of the union of CS2 and CO2, set forth in Norman Gibbs's and Alan Tucker's 1985 *Model Curriculum for a Liberal Arts Degree in Computer Science*. The material contained here should be covered early in any computer science curriculum, and I wrote this book for an audience of first- and second-year students in computer science who were familiar with a high-level language such as Pascal. A course in discrete mathematics is desirable as a pre- or corequisite for this material, but the relevant mathematical background in relations, probability, and other topics is summarized in the appendixes for those who need it.

THE CONTENTS

My intent was to write a book that could be used as the basis for a semester-length course in data structures or advanced programming. Realizing that the subject matter of

this book comes at an early stage in the education of a computer scientist, I included a number of mentions, necessarily brief, of some of the interesting topics that await the student down the road. Most of the canonical sorting algorithms are covered, along with mentions of computational complexity, compiler design, unsolvable problems, NP-completeness and the fundamental paradigms for algorithms.

Chapters 1 and 2 form the introduction of the text. Chapter 1 covers the necessary preliminaries, such as program design, the definition of an abstract data type, a review of Pascal-based pseudocode and pointers, big-O estimates, and the notion of time and space estimates. Appendix A, on relations, provides a more detailed study of the material of Chapter 1. Chapter 2 is typical of subsequent chapters, in that we introduce an abstract data type, `STRING`, describe how it may be implemented, compare the implementations, and provide extensive application, in this case the Boyer-Moore string-matching algorithm.

In Chapters 3 and 4 we continue the investigation of linear data structures. Chapter 3 covers lists, along with related list-like structures such as doubly linked lists, sorted lists, and braids, and concludes with a discussion of memory management. Chapter 4 covers the rest of standard linear structures, stacks and queues, motivating these by applications to manipulate postfix expressions. Since a considerable amount of queue applications involve simulation, Appendix C would be appropriate at this point.

Chapter 5 provides a segue into nonlinear structures by providing an introduction to recursion and recursively defined data structures. Timing estimates for recursive algorithms are covered in depth, along with an introduction to LISP. Appendix B covers logs and exponentials, induction, and elementary combinatorics, and would be appropriate supplementary material at this stage.

Chapters 6 and 7 cover trees. Chapter 6 provides the necessary background on binary trees and their implementations, traversal algorithms, treesort and heapsort, and discusses minimal-length codes and tries. Chapter 7, which could be omitted if necessary, covers two extensions of binary search trees, namely AVL trees and B-trees.

Chapter 8 could also be optional. It covers graphs and digraphs, along with a representative sample of graph algorithms for traversal, spanning trees, minimal-cost paths, minimal spanning trees, and an introduction to complexity theory through the Traveling Salesperson Problem.

Chapter 9, on sets, describes bit vector and list implementations of sets, as well as dictionaries, and provides a comprehensive introduction to hashing. The last section of Chapter 9 discusses the UNION-FIND data structure, and could be omitted, if worse came to worse.

In Chapter 10 we consider a problem of regenerating text from a large sample and trace the development of programs to solve this problem, using a real computer/compiler system to show how time and space constraints arise in practice from choices of data structure.

ACKNOWLEDGMENTS

Stephen King will probably make more money from his next book than I'll see in the rest of my life. It's worth every penny, folks: writing is just plain hard work. A lot of

people deserve praise for seeing this book of mine through to completion. Thanks go to my colleagues for reading advance copies, especially to Peter Allen of Columbia University, Brian Kernaghan of AT&T/Bell Labs, and Seymour Pollack of Washington University for their review of the manuscript; to my students for catching the countless errors in the earlier versions; to Eloise Starkweather for producing a cover that knocked my socks off; to the nice folks at Prentice Hall for being so tolerant of my first effort with them; and especially to my family and friends, none of whom will ever read *Data Structures*, but all of whom were supportive and understanding when I came dragging in after a long day's work.

Contents

PREFACE	xi
Part 1 INTRODUCTION	1
1 PRELIMINARIES	3
1.1 Program Design: Algorithms and Data Structures 4	
<i>Algorithms</i> 4	
<i>Data Structures</i> 6	
1.2 Abstract Data Types: The Threefold Way 9	
<i>Relations</i> 11	
1.3 Pseudocode and Pointers 15	
<i>Pointers</i> 16	
1.4 Complexity Measures 20	
<i>Big-O</i> 20	
<i>Order Arithmetic</i> 22	
<i>Timing Functions</i> 24	
1.5 Summary 28	
2 STRINGS	34
2.1 The Abstract Data Type STRING 34	
2.2 Array Implementations of STRING 39	
<i>Static Arrays</i> 39	
<i>String Tables</i> 43	
2.3 Pointer Implementation of STRING 48	
2.4 Application: Pattern Matching 54	
2.5 Summary 60	

Part 2 LINEAR STRUCTURES**65****3 LISTS****67**

3.1 The Abstract Data Type LIST 68

3.2 Implementations of LIST 71

Arrays 71*Linked Lists* 73*Cursors* 77

3.3 Other List Structures 80

Two-Way Lists 81*Sorted Lists* 82*Braids* 86

3.4 Application: Memory Management 87

Allocation 88*Deallocation* 91*Compaction* 95

3.5 Summary 108

4 OTHER LINEAR STRUCTURES**114**

4.1 The Abstract Data Type STACK 114

Array Implementation of STACK 118*Stacks as Linked Lists* 119*Application 1: Postfix Arithmetic* 121*Application 2: The Electronic Labyrinth* 123

4.2 Queues 127

Array Implementations of QUEUE 128*Pointer Implementation of QUEUE* 130

4.3 Application: Infix to Postfix Conversion 132

4.4 Summary 134

Part 3 NONLINEAR STRUCTURES**141**

5 RECURSION	142
5.1 Recursive Algorithms 142	
<i>Induction and Recursion</i> 148	
<i>Timing Recursive Algorithms</i> 149	
5.2 Case Study: Design of Algorithms 154	
5.3 Recursive Data Structures 159	
<i>General Lists and LISP</i> 161	
5.4 Summary 167	
6 TREES	174
6.1 The Structure of Trees 175	
6.2 Binary Trees 179	
6.3 Binary Tree Traversals 182	
<i>Application: Parse Trees</i> 186	
6.4 Implementations of BINARY TREE 189	
<i>Pointers</i> 189	
<i>Threaded Trees</i> 191	
<i>Cursors</i> 195	
6.5 Data-Ordered Binary Trees 197	
<i>Binary Search Trees</i> 198	
<i>Application: Treesort</i> 204	
<i>The PRIORITY QUEUE ADT and Heaps</i> 205	
<i>Application: Heapsort</i> 210	
6.6 Two Applications of Trees 212	
<i>Huffman Codes</i> 213	
<i>Tries</i> 216	
6.7 Summary 219	
7 SPECIALIZED TREES	225
7.1 Balanced Trees 226	
7.2 B-Trees 233	
<i>k-ary Trees, Again</i> 234	
<i>B-Trees, Explained</i> 235	
<i>Application: External Storage</i> 244	
7.3 Summary 248	

8	GRAPHS AND DIGRAPHS	251
8.1	Graphs	252
8.2	Representations of GRAPH	256
	<i>Adjacency Matrices</i>	256
	<i>Adjacency Lists and Edge Lists</i>	259
8.3	Graph Traversals	263
	<i>Depth-First Traversals</i>	263
	<i>Breadth-First Traversals</i>	265
	<i>Spanning Trees</i>	266
8.4	Further Applications of GRAPH	269
	<i>Counting Paths</i>	269
	<i>Minimum Spanning Trees</i>	272
8.5	Directed Graphs	274
	<i>Application: Cheapest Paths (II)</i>	275
8.6	Computational Complexity and Graph Algorithms	280
8.7	Summary	284
9	SETS	290
9.1	The SET ADT	291
9.2	Bit Vectors and Ordinal Sets	292
	<i>Sets Represented by Sorted Lists</i>	295
9.3	Hashing	297
	<i>Open Hashing</i>	303
9.4	Hashing, Continued	304
	<i>Hash Functions</i>	304
	<i>Collision Resolution, Again</i>	306
	<i>Time and Space Estimates</i>	308
	<i>Application: Multiple Hashing, Encoding, and Spelling Checkers</i>	311
9.5	The UNION-FIND ADT	314
	<i>Tree Representations of UNION-FIND</i>	315
	<i>Weighting</i>	318
	<i>Path Compression</i>	320
	<i>Application: Minimum Spanning Trees, Revisited</i>	322
9.6	Summary	324

10 TRAVESTY—PUTTING IT ALL TOGETHER	329
10.1 The Problem	329
10.2 The Solutions	334
<i>Arrays</i>	335
<i>Hashing</i>	340
<i>Tries</i>	347
<i>Long Strings</i>	354
10.3 Applications	360
<i>Reactive Keyboards</i>	360
<i>Coding, Once Again</i>	361
10.4 Summary	362
Appendix A RELATIONS	365
Successors and Restrictions	369
Application: Linear Orders and Closures	372
Appendix B TOPICS IN MATHEMATICS	378
B.1 Exponential and Logarithmic Functions	378
<i>Logarithms</i>	381
B.2 Induction	383
B.3 Counting Techniques	385
<i>Permutations</i>	385
<i>Combinations</i>	386
Appendix C RANDOM NUMBERS AND SIMULATION	392
C.1 Random Numbers	393
C.2 Probability Distributions	395
<i>Cumulative Probability Functions</i>	396
<i>Uniform Distributions</i>	398
<i>Normal Distributions</i>	399
<i>Exponential Distributions</i>	400

C.3 Selection Algorithms	402
<i>Selecting Subsets</i>	402
<i>Selecting a Subset with k Elements</i>	403
<i>Enumerating Permutations</i>	406

INDEX

409

PART 1

Introduction

Preliminaries

*When making haggis, the easiest part is obtaining the ingredients.
Anyone can get a sheep's stomach and a quantity of oatmeal,
but the real trick is knowing how to combine them.*

Roberta Campbell, *The Cuisine of Scotland*

The discipline of computer science is concerned with the study of problem solving with computers. Notice that we did not say that computer science *consists* of problem solving with computers, any more than mathematics consists of solving equations or music consists of producing notes. It is not enough to be able to answer the question, "How do we solve a particular problem with the help of a computer?" If it were, the study of computer science could stop after one or two introductory programming courses. Instead, the proper subject matter includes questions like

1. What are the possible different ways to solve a problem?
2. How are the solutions for a particular problem related?
3. What technique is best for a particular problem?
4. What do we mean by a "best" solution for a problem?
5. In what ways are solutions for different problems related?
6. How do we verify that we have a solution for a problem?
7. What problems can, and cannot, be solved with a computer?

Although all of these questions contain the word *problem*, they all seek answers in a context that is broader than simply solving a particular problem. In fact, all of these

questions, and all of the questions of computer science, are different aspects of the same fundamental question:

What *general principles* underlie the notion of problem solving with computers?

In this text, we will concern ourselves primarily with those aspects of this fundamental question that deal with the structure of the data in a program and, to a lesser extent, with the techniques of manipulating that data.

1.1 PROGRAM DESIGN: ALGORITHMS AND DATA STRUCTURES

We can view the subject matter of computer science as the result of the process of generalizing from specific problem-solving instances, which is to say that computer science seeks to find properties common to many instances of problem solving. Program design—the writing of programs to solve a problem—proceeds in the opposite direction: from a vague notion of what needs to be done, to the writing of a program in a specific language for a specific computer. The aim of program design is captured in the most appropriate title yet invented for a book on computer science, Nicholas Wirth's *Algorithms + Data Structures = Programs*. Wirth, the developer of the Pascal language, chose his title to point out the twofold nature of a computer program: that a program consists of an algorithm describing how to manipulate information with a computer, along with a data structure that provides a logical basis for organization of that information in the computer. These two aspects of a program are intimately intertwined: Making a decision about one of the aspects often profoundly affects the other.

Algorithms

An **algorithm** is a finite list of unambiguous instructions that can be performed on a computer in such a way that the process is guaranteed to halt in a finite amount of time. "Add up the integers from 1 to 100" almost qualifies as an algorithm, except that the single instruction it uses is ambiguous—it does not provide sufficient detail for us to decide how to perform the required operation. The instruction does provide us with a useful starting point, however: Reading it, we have a clear idea of what problem we have to solve. Indeed, just getting to the point where we know what the problem is can often represent the major part of a programming task. Knowing the problem, we can now try to refine the problem into a suitable algorithm.

This simple addition problem occupies a hallowed place in mathematical folklore, and will serve as a good example of a situation in which there is more than one algorithm to solve a given problem. Karl Frederick Gauss was born in Germany in 1777, and grew to be, if not the best, then certainly one of the best mathematicians who ever lived. The story goes* that when Gauss was a boy in what would be the eighteenth-

*This story has about the same amount of truth to it as the tale of George Washington and the cherry tree, and has survived for about the same reasons.