


DATA STRUCTURES AND ALGORITHMS

ALFRED V. AHO
JOHN E. HOPCROFT
JEFFREY D. ULLMAN



Data Structures and Algorithms

ALFRED V. AHO

*Bell Laboratories
Murray Hill, New Jersey*

JOHN E. HOPCROFT

*Cornell University
Ithaca, New York*

JEFFREY D. ULLMAN

*Stanford University
Stanford, California*



ADDISON-WESLEY PUBLISHING COMPANY
*Reading, Massachusetts • Menlo Park, California
London • Amsterdam • Don Mills, Ontario • Sydney*

This book is in the
ADDISON-WESLEY SERIES IN
COMPUTER SCIENCE AND INFORMATION PROCESSING

Michael A. Harrison
Consulting Editor

Library of Congress Cataloging in Publication Data

Aho, Alfred V.

Data structures and algorithms.

1. Data structures (Computer science) 2. Algorithms.

I. Hopcroft, John E., 1939-

II. Ullman,

Jeffrey D., 1942-

III. Title.

QA76.9.D35A38 1982 001.64 82-11596

ISBN 0-201-00023-7

Reproduced by Addison-Wesley from camera-ready copy supplied by the authors.

Reprinted with corrections April, 1987

Copyright © 1983 by Bell Telephone Laboratories, Incorporated.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. Printed in the United States of America. Published simultaneously in Canada.

ISBN: 0-201-00023-7

MNOP-DO-8987

Preface

This book presents the data structures and algorithms that underpin much of today's computer programming. The basis of this book is the material contained in the first six chapters of our earlier work, *The Design and Analysis of Computer Algorithms*. We have expanded that coverage and have added material on algorithms for external storage and memory management. As a consequence, this book should be suitable as a text for a first course on data structures and algorithms. The only prerequisite we assume is familiarity with some high-level programming language such as Pascal.

We have attempted to cover data structures and algorithms in the broader context of solving problems using computers. We use abstract data types informally in the description and implementation of algorithms. Although abstract data types are only starting to appear in widely available programming languages, we feel they are a useful tool in designing programs, no matter what the language.

We also introduce the ideas of step counting and time complexity as an integral part of the problem solving process. This decision reflects our long-held belief that programmers are going to continue to tackle problems of progressively larger size as machines get faster, and that consequently the time complexity of algorithms will become of even greater importance, rather than of less importance, as new generations of hardware become available.

The Presentation of Algorithms

We have used the conventions of Pascal to describe our algorithms and data structures primarily because Pascal is so widely known. Initially we present several of our algorithms both abstractly and as Pascal programs, because we feel it is important to run the gamut of the problem solving process from problem formulation to a running program. The algorithms we present, however, can be readily implemented in any high-level programming language.

Use of the Book

Chapter 1 contains introductory remarks, including an explanation of our view of the problem-to-program process and the role of abstract data types in that process. Also appearing is an introduction to step counting and "big-oh" and "big-omega" notation.

Chapter 2 introduces the traditional list, stack and queue structures, and the mapping, which is an abstract data type based on the mathematical notion of a function. The third chapter introduces trees and the basic data structures

that can be used to support various operations on trees efficiently.

Chapters 4 and 5 introduce a number of important abstract data types that are based on the mathematical model of a set. Dictionaries and priority queues are covered in depth. Standard implementations for these concepts, including hash tables, binary search trees, partially ordered trees, tries, and 2-3 trees are covered, with the more advanced material clustered in Chapter 5.

Chapters 6 and 7 cover graphs, with directed graphs in Chapter 6 and undirected graphs in 7. These chapters begin a section of the book devoted more to issues of algorithms than data structures, although we do discuss the basics of data structures suitable for representing graphs. A number of important graph algorithms are presented, including depth-first search, finding minimal spanning trees, shortest paths, and maximal matchings.

Chapter 8 is devoted to the principal internal sorting algorithms: quick-sort, heapsort, binsort, and the simpler, less efficient methods such as insertion sort. In this chapter we also cover the linear-time algorithms for finding medians and other order statistics.

Chapter 9 discusses the asymptotic analysis of recursive procedures, including, of course, recurrence relations and techniques for solving them.

Chapter 10 outlines the important techniques for designing algorithms, including divide-and-conquer, dynamic programming, local search algorithms, and various forms of organized tree searching.

The last two chapters are devoted to external storage organization and memory management. Chapter 11 covers external sorting and large-scale storage organization, including B-trees and index structures.

Chapter 12 contains material on memory management, divided into four subareas, depending on whether allocations involve fixed or varying sized blocks, and whether the freeing of blocks takes place by explicit program action or implicitly when garbage collection occurs.

Material from this book has been used by the authors in data structures and algorithms courses at Columbia, Cornell, and Stanford, at both undergraduate and graduate levels. For example, a preliminary version of this book was used at Stanford in a 10-week course on data structures, taught to a population consisting primarily of Juniors through first-year graduate students. The coverage was limited to Chapters 1-4, 9, 10, and 12, with parts of 5-7.

Exercises

A number of exercises of varying degrees of difficulty are found at the end of each chapter. Many of these are fairly straightforward tests of the mastery of the material of the chapter. Some exercises require more thought, and these have been singly starred. Doubly starred exercises are harder still, and are suitable for more advanced courses. The bibliographic notes at the end of each chapter provide references for additional reading.

Acknowledgments

We wish to acknowledge Bell Laboratories for the use of its excellent UNIXTM-based text preparation and data communication facilities that significantly eased the preparation of a manuscript by geographically separated authors. Many of our colleagues have read various portions of the manuscript and have given us valuable comments and advice. In particular, we would like to thank Ed Beckham, Jon Bentley, Kenneth Chu, Janet Coursey, Hank Cox, Neil Immerman, Brian Kernighan, Steve Mahaney, Craig McMurray, Alberto Mendelzon, Alistair Moffat, Jeff Naughton, Kerry Nemovicher, Paul Niamkey, Yoshio Ohno, Rob Pike, Chris Rouen, Maurice Schlumberger, Stanley Selkow, Chengya Shih, Bob Tarjan, W. Van Snyder, Peter Weinberger, and Anthony Yeracaris for helpful suggestions. Finally, we would like to give our warmest thanks to Mrs. Claire Metzger for her expert assistance in helping prepare the manuscript for typesetting.

A.V.A.

J.E.H.

J.D.U.

Contents

Chapter 1 Design and Analysis of Algorithms

1.1	From Problems to Programs	1
1.2	Abstract Data Types	10
1.3	Data Types, Data Structures, and Abstract Data Types	13
1.4	The Running Time of a Program	16
1.5	Calculating the Running Time of a Program	21
1.6	Good Programming Practice	27
1.7	Super Pascal	29

Chapter 2 Basic Data Types

2.1	The Data Type "List"	37
2.2	Implementation of Lists	40
2.3	Stacks	53
2.4	Queues	56
2.5	Mappings	61
2.6	Stacks and Recursive Procedures	64

Chapter 3 Trees

3.1	Basic Terminology	75
3.2	The ADT TREE	82
3.3	Implementations of Trees	84
3.4	Binary Trees	93

Chapter 4 Basic Operations on Sets

4.1	Introduction to Sets	107
4.2	An ADT with Union, Intersection, and Difference	109
4.3	A Bit-Vector Implementation of Sets	112
4.4	A Linked-List Implementation of Sets	115
4.5	The Dictionary	117
4.6	Simple Dictionary Implementations	119
4.7	The Hash Table Data Structure	122
4.8	Estimating the Efficiency of Hash Functions	129
4.9	Implementation of the Mapping ADT	135
4.10	Priority Queues	135
4.11	Implementations of Priority Queues	138
4.12	Some Complex Set Structures	145

Chapter 5 Advanced Set Representation Methods

5.1	Binary Search Trees	155
5.2	Time Analysis of Binary Search Tree Operations	160
5.3	Tries	163
5.4	Balanced Tree Implementations of Sets	169
5.5	Sets with the MERGE and FIND Operations	180
5.6	An ADT with MERGE and SPLIT	189

Chapter 6 Directed Graphs

6.1	Basic Definitions	198
6.2	Representations for Directed Graphs	199
6.3	The Single-Source Shortest Paths Problem	203
6.4	The All-Pairs Shortest Path Problem	208
6.5	Traversals of Directed Graphs	215
6.6	Directed Acyclic Graphs	219
6.7	Strong Components	222

Chapter 7 Undirected Graphs

7.1	Definitions	230
7.2	Minimum-Cost Spanning Trees	233
7.3	Traversals	239
7.4	Articulation Points and Biconnected Components	244
7.5	Graph Matching	246

Chapter 8 Sorting

8.1	The Internal Sorting Model	253
8.2	Some Simple Sorting Schemes	254
8.3	Quicksort	260
8.4	Heapsort	271
8.5	Bin Sorting	274
8.6	A Lower Bound for Sorting by Comparisons	282
8.7	Order Statistics	286

Chapter 9 Algorithm Analysis Techniques

9.1	Efficiency of Algorithms	293
9.2	Analysis of Recursive Programs	294
9.3	Solving Recurrence Equations	296
9.4	A General Solution for a Large Class of Recurrences	298

Chapter 10 Algorithm Design Techniques

10.1	Divide-and-Conquer Algorithms	306
10.2	Dynamic Programming	311
10.3	Greedy Algorithms	321
10.4	Backtracking	324
10.5	Local Search Algorithms	336

Chapter 11 Data Structures and Algorithms for External Storage

11.1	A Model of External Computation.....	347
11.2	External Sorting	349
11.3	Storing Information in Files	361
11.4	External Search Trees	368

Chapter 12 Memory Management

12.1	The Issues in Memory Management.....	378
12.2	Managing Equal-Sized Blocks.....	382
12.3	Garbage Collection Algorithms for Equal-Sized Blocks	384
12.4	Storage Allocation for Objects with Mixed Sizes	392
12.5	Buddy Systems	400
12.6	Storage Compaction	404

Bibliography	411
---------------------------	-----

Index	419
--------------------	-----

CHAPTER 1

Design and Analysis of Algorithms

There are many steps involved in writing a computer program to solve a given problem. The steps go from problem formulation and specification, to design of the solution, to implementation, testing and documentation, and finally to evaluation of the solution. This chapter outlines our approach to these steps. Subsequent chapters discuss the algorithms and data structures that are the building blocks of most computer programs.

1.1 From Problems to Programs

Half the battle is knowing what problem to solve. When initially approached, most problems have no simple, precise specification. In fact, certain problems, such as creating a “gourmet” recipe or preserving world peace, may be impossible to formulate in terms that admit of a computer solution. Even if we suspect our problem can be solved on a computer, there is usually considerable latitude in several problem parameters. Often it is only by experimentation that reasonable values for these parameters can be found.

If certain aspects of a problem can be expressed in terms of a formal model, it is usually beneficial to do so, for once a problem is formalized, we can look for solutions in terms of a precise model and determine whether a program already exists to solve that problem. Even if there is no existing program, at least we can discover what is known about this model and use the properties of the model to help construct a good solution.

Almost any branch of mathematics or science can be called into service to help model some problem domain. Problems essentially numerical in nature can be modeled by such common mathematical concepts as simultaneous linear equations (e.g., finding currents in electrical circuits, or finding stresses in frames made of connected beams) or differential equations (e.g., predicting population growth or the rate at which chemicals will react). Symbol and text processing problems can be modeled by character strings and formal grammars. Problems of this nature include compilation (the translation of programs written in a programming language into machine language) and information retrieval tasks such as recognizing particular words in lists of titles owned by a library.

Algorithms

Once we have a suitable mathematical model for our problem, we can attempt to find a solution in terms of that model. Our initial goal is to find a solution in the form of an *algorithm*, which is a finite sequence of instructions, each of which has a clear meaning and can be performed with a finite amount of effort in a finite length of time. An integer assignment statement such as $x := y + z$ is an example of an instruction that can be executed in a finite amount of effort. In an algorithm instructions can be executed any number of times, provided the instructions themselves indicate the repetition. However, we require that, no matter what the input values may be, an algorithm terminate after executing a finite number of instructions. Thus, a program is an algorithm as long as it never enters an infinite loop on any input.

There is one aspect of this definition of an algorithm that needs some clarification. We said each instruction of an algorithm must have a “clear meaning” and must be executable with a “finite amount of effort.” Now what is clear to one person may not be clear to another, and it is often difficult to prove rigorously that an instruction can be carried out in a finite amount of time. It is often difficult as well to prove that on any input, a sequence of instructions terminates, even if we understand clearly what each instruction means. By argument and counterargument, however, agreement can usually be reached as to whether a sequence of instructions constitutes an algorithm. The burden of proof lies with the person claiming to have an algorithm. In Section 1.5 we discuss how to estimate the running time of common programming language constructs that can be shown to require a finite amount of time for their execution.

In addition to using Pascal programs as algorithms, we shall often present algorithms using a *pseudo-language* that is a combination of the constructs of a programming language together with informal English statements. We shall use Pascal as the programming language, but almost any common programming language could be used in place of Pascal for the algorithms we shall discuss. The following example illustrates many of the steps in our approach to writing a computer program.

Example 1.1. A mathematical model can be used to help design a traffic light for a complicated intersection of roads. To construct the pattern of lights, we shall create a program that takes as input a set of permitted turns at an intersection (continuing straight on a road is a “turn”) and partitions this set into as few groups as possible such that all turns in a group are simultaneously permissible without collisions. We shall then associate a phase of the traffic light with each group in the partition. By finding a partition with the smallest number of groups, we can construct a traffic light with the smallest number of phases.

For example, the intersection shown in Fig. 1.1 occurs by a watering hole called JoJo’s near Princeton University, and it has been known to cause some navigational difficulty, especially on the return trip. Roads *C* and *E* are one-way, the others two way. There are 13 turns one might make at this

intersection. Some pairs of turns, like AB (from A to B) and EC , can be carried out simultaneously, while others, like AD and EB , cause lines of traffic to cross and therefore cannot be carried out simultaneously. The light at the intersection must permit turns in such an order that AD and EB are never permitted at the same time, while the light might permit AB and EC to be made simultaneously.




Fig. 1.1. An intersection.

We can model this problem with a mathematical structure known as a graph. A *graph* consists of a set of points called *vertices*, and lines connecting the points, called *edges*. For the traffic intersection problem we can draw a graph whose vertices represent turns and whose edges connect pairs of vertices whose turns cannot be performed simultaneously. For the intersection of Fig. 1.1, this graph is shown in Fig. 1.2, and in Fig. 1.3 we see another representation of this graph as a table with a 1 in row i and column j whenever there is an edge between vertices i and j .

The graph can aid us in solving the traffic light design problem. A *coloring* of a graph is an assignment of a color to each vertex of the graph so that no two vertices connected by an edge have the same color. It is not hard to see that our problem is one of coloring the graph of incompatible turns using as few colors as possible.

The problem of coloring graphs has been studied for many decades, and the theory of algorithms tells us a lot about this problem. Unfortunately, coloring an arbitrary graph with as few colors as possible is one of a large class of problems called “NP-complete problems,” for which all known solutions are essentially of the type “try all possibilities.” In the case of the coloring problem, “try all possibilities” means to try all assignments of colors to vertices using at first one color, then two colors, then three, and so on, until a legal coloring is found. With care, we can be a little speedier than this, but it is generally believed that no algorithm to solve this problem can be substantially more efficient than this most obvious approach.

We are now confronted with the possibility that finding an optimal solution for the problem at hand is computationally very expensive. We can adopt




Fig. 1.2. Graph showing incompatible turns.

	<i>AB</i>	<i>AC</i>	<i>AD</i>	<i>BA</i>	<i>BC</i>	<i>BD</i>	<i>DA</i>	<i>DB</i>	<i>DC</i>	<i>EA</i>	<i>EB</i>	<i>EC</i>	<i>ED</i>
<i>AB</i>					1	1	1			1			
<i>AC</i>						1	1	1		1	1		
<i>AD</i>										1	1	1	
<i>BA</i>													
<i>BC</i>	1							1			1		
<i>BD</i>	1	1					1				1	1	
<i>DA</i>	1	1				1					1	1	
<i>DB</i>		1			1							1	
<i>DC</i>													
<i>EA</i>	1	1	1										
<i>EB</i>		1	1		1	1	1						
<i>EC</i>			1			1	1	1					
<i>ED</i>													

Fig. 1.3. Table of incompatible turns.

one of three approaches. If the graph is small, we might attempt to find an optimal solution exhaustively, trying all possibilities. This approach, however, becomes prohibitively expensive for large graphs, no matter how efficient we try to make the program. A second approach would be to look for additional information about the problem at hand. It may turn out that the graph has some special properties, which make it unnecessary to try all possibilities in finding an optimal solution. The third approach is to change the problem a little and look for a good but not necessarily optimal solution. We might be happy with a solution that gets close to the minimum number of colors on small graphs, and works quickly, since most intersections are not even as complex as Fig. 1.1. An algorithm that quickly produces good but not necessarily optimal solutions is called a *heuristic*.

One reasonable heuristic for graph coloring is the following “greedy” algorithm. Initially we try to color as many vertices as possible with the first color, then as many as possible of the uncolored vertices with the second color, and so on. To color vertices with a new color, we perform the following steps.

1. Select some uncolored vertex and color it with the new color.
2. Scan the list of uncolored vertices. For each uncolored vertex, determine whether it has an edge to any vertex already colored with the new color. If there is no such edge, color the present vertex with the new color.

This approach is called “greedy” because it colors a vertex whenever it can, without considering the potential drawbacks inherent in making such a move. There are situations where we could color more vertices with one color if we were less “greedy” and skipped some vertex we could legally color. For example, consider the graph of Fig. 1.4, where having colored vertex 1 red, we can color vertices 3 and 4 red also, provided we do not color 2 first. The greedy algorithm would tell us to color 1 and 2 red, assuming we considered vertices in numerical order.




Fig. 1.4. A graph.

As an example of the greedy approach applied to Fig. 1.2, suppose we start by coloring AB blue. We can color AC , AD , and BA blue, because none of these four vertices has an edge in common. We cannot color BC blue because there is an edge between AB and BC . Similarly, we cannot color BD ,

DA , or DB blue because each of these vertices is connected by an edge to one or more vertices already colored blue. However, we can color DC blue. Then EA , EB , and EC cannot be colored blue, but ED can.

Now we start a second color, say by coloring BC red. BD can be colored red, but DA cannot, because of the edge between BD and DA . Similarly, DB cannot be colored red, and DC is already blue, but EA can be colored red. Each other uncolored vertex has an edge to a red vertex, so no other vertex can be colored red.

The remaining uncolored vertices are DA , DB , EB , and EC . If we color DA green, then DB can be colored green, but EB and EC cannot. These two may be colored with a fourth color, say yellow. The colors are summarized in Fig. 1.5. The "extra" turns are determined by the greedy approach to be compatible with the turns already given that color, as well as with each other. When the traffic light allows turns of one color, it can also allow the extra turns safely.

color	turns	extras
blue	AB, AC, AD, BA, DC, ED	—
red	BC, BD, EA	BA, DC, ED
green	DA, DB	AD, BA, DC, ED
yellow	EB, EC	BA, DC, EA, ED

Fig. 1.5. A coloring of the graph of Fig. 1.2.

The greedy approach does not always use the minimum possible number of colors. We can use the theory of algorithms again to evaluate the goodness of the solution produced. In graph theory, a k -clique is a set of k vertices, every pair of which is connected by an edge. Obviously, k colors are needed to color a k -clique, since no two vertices in a clique may be given the same color.

In the graph of Fig. 1.2 the set of four vertices AC, DA, BD, EB is a 4-clique. Therefore, no coloring with three or fewer colors exists, and the solution of Fig. 1.5 is optimal in the sense that it uses the fewest colors possible. In terms of our original problem, no traffic light for the intersection of Fig. 1.1 can have fewer than four phases.

Therefore, consider a traffic light controller based on Fig. 1.5, where each phase of the controller corresponds to a color. At each phase the turns indicated by the row of the table corresponding to that color are permitted, and the other turns are forbidden. This pattern uses as few phases as possible. $\square$

Pseudo-Language and Stepwise Refinement

Once we have an appropriate mathematical model for a problem, we can formulate an algorithm in terms of that model. The initial versions of the algorithm are often couched in general statements that will have to be refined subsequently into smaller, more definite instructions. For example, we described the greedy graph coloring algorithm in terms such as "select some uncolored vertex." These instructions are, we hope, sufficiently clear that the reader grasps our intent. To convert such an informal algorithm to a program, however, we must go through several stages of formalization (called *stepwise refinement*) until we arrive at a program the meaning of whose steps are formally defined by a language manual.

Example 1.2. Let us take the greedy algorithm for graph coloring part of the way towards a Pascal program. In what follows, we assume there is a graph G , some of whose vertices may be colored. The following program *greedy* determines a set of vertices called *newclr*, all of which can be colored with a new color. The program is called repeatedly, until all vertices are colored. At a coarse level, we might specify *greedy* in pseudo-language as in Fig. 1.6.

```

procedure greedy ( var  $G$ : GRAPH; var newclr: SET );
    { greedy assigns to newclr a set of vertices of  $G$  that may be
      given the same color }
    begin
(1)      newclr :=  $\emptyset$ ; †
(2)      for each uncolored vertex  $v$  of  $G$  do
(3)          if  $v$  is not adjacent to any vertex in newclr then begin
(4)              mark  $v$  colored;
(5)              add  $v$  to newclr
            end
    end; { greedy }

```

Fig. 1.6. First refinement of *greedy* algorithm.

We notice from Fig. 1.6 certain salient features of our pseudo-language. First, we use boldface lower case keywords corresponding to Pascal reserved words, with the same meaning as in standard Pascal. Upper case types such as GRAPH and SET‡ are the names of "abstract data types." They will be defined by Pascal type definitions and the operations associated with these abstract data types will be defined by Pascal procedures when we create the final program. We shall discuss abstract data types in more detail in the next two sections.

The flow-of-control constructs of Pascal, like **if**, **for**, and **while**, are

† The symbol $\emptyset$ stands for the empty set.

‡ We distinguish the abstract data type SET from the built-in set type of Pascal.

available for pseudo-language statements, but conditionals, as in line (3), may be informal statements rather than Pascal conditional expressions. Note that the assignment at line (1) uses an informal expression on the right. Also, the **for**-loop at line (2) iterates over a set.

To be executed, the pseudo-language program of Fig. 1.6 must be refined into a conventional Pascal program. We shall not proceed all the way to such a program in this example, but let us give one example of refinement, transforming the **if**-statement in line (3) of Fig. 1.6 into more conventional code.

To test whether vertex v is adjacent to some vertex in *newclr*, we consider each member w of *newclr* and examine the graph G to see whether there is an edge between v and w . An organized way to make this test is to use *found*, a boolean variable to indicate whether an edge has been found. We can replace lines (3)–(5) of Fig. 1.6 by the code in Fig. 1.7.

```

procedure greedy ( var  $G$ : GRAPH; var newclr: SET );
begin
(1)      newclr :=  $\emptyset$ ;
(2)      for each uncolored vertex  $v$  of  $G$  do begin
(3.1)    found := false;
(3.2)    for each vertex  $w$  in newclr do
(3.3)    if there is an edge between  $v$  and  $w$  in  $G$  then
(3.4)    found := true;
(3.5)    if found = false then begin
           {  $v$  is adjacent to no vertex in newclr }
(4)      mark  $v$  colored;
(5)      add  $v$  to newclr
           end
        end
end; { greedy }

```

Fig. 1.7. Refinement of part of Fig. 1.6.

We have now reduced our algorithm to a collection of operations on two sets of vertices. The outer loop, lines (2)–(5), iterates over the set of uncolored vertices of G . The inner loop, lines (3.2)–(3.4), iterates over the vertices currently in the set *newclr*. Line (5) adds newly colored vertices to *newclr*.

There are a variety of ways to represent sets in a programming language like Pascal. In Chapters 4 and 5 we shall study several such representations. In this example we can simply represent each set of vertices by another abstract data type LIST, which here can be implemented by a list of integers terminated by a special value *null* (for which we might use the value 0). These integers might, for example, be stored in an array, but there are many other ways to represent LIST's, as we shall see in Chapter 2.