



CONCEPTS IN DATA STRUCTURES & SOFTWARE DEVELOPMENT

▼

A TEXT FOR
THE SECOND
COURSE IN
COMPUTER
SCIENCE

▲

G. MICHAEL SCHNEIDER
STEVEN C. BRUELL

CONCEPTS IN DATA STRUCTURES AND SOFTWARE DEVELOPMENT

A Text for the Second Course
in Computer Science

G. Michael Schneider
Macalester College

Steven C. Bruell
University of Iowa

WEST PUBLISHING COMPANY

St. Paul New York Los Angeles San Francisco

COPYEDITING: Pam McMurry
TEXT DESIGN: Lucy Lesiak Design
TECHNICAL ILLUSTRATIONS: G & S Typesetters, Inc.
COMPOSITION: G & S Typesetters, Inc.
COVER DESIGN: David J. Farr, Imagesmythe, Inc.

Turbo Pascal is a registered trademark of Borland International, Inc.
Ada is a registered trademark of the United States Department of Defense.
Cray-2 is a registered trademark of Cray Research, Inc.
VAX is a registered trademark of Digital Equipment Corporation
Macintosh is a registered trademark of Apple Computer, Inc.

COPYRIGHT © 1992 By WEST PUBLISHING COMPANY
610 Opperman Drive
P.O. Box 64526
St. Paul, MN 55164-1003

All rights reserved

Printed in the United States of America

98 97 96

8 7 6 5 4 3 2

Library of Congress Cataloging-in-Publication Data

Schneider, G. Michael.

Concepts in data structures and software development : a text for the second course in
computer science / G. Michael Schneider, Steven C. Bruell.

p. cm.

Includes index.

ISBN 0-314-77460-2 (hard)

1. Data structures (Computer science) 2. Computer software—Development.

I. Bruell, Steven C. II. Title.

QA76.9.D35S36 1991

005.1—dc20

Printed with **Printwise**
Environmentally Advanced Water Washable Ink



90-42783
CIP

CONCEPTS IN DATA STRUCTURES AND SOFTWARE DEVELOPMENT

A Text for the Second Course
in Computer Science

I dedicate this book to my wife Ruthann, my children Benjamin and Rebecca,
and my sister Karen.

G. Michael Schneider

To Sandy

Steven C. Bruell

ABOUT THE AUTHORS

G. Michael Schneider

G. Michael Schneider is currently Professor of computer science and Director of the computer science program at Macalester College in St. Paul, Minnesota, where he has been since 1982. Prior to that he was on the faculty of computer science at the University of Minnesota for eight years.

In addition to his research interests, he is the author of a number of best-selling college textbooks, including *Introduction to Programming and Problem Solving with Pascal*, *Advanced Programming and Problem Solving with Pascal* (with Steven C. Bruell), and *Principles of Computer Organization*. For the past ten years he has been extremely active in computer science education and curriculum development through membership in the ACM Special Interest Group in Computer Science Education (SIGCSE), the national Computer Science Accreditation Board (CSAB), and the Consortium of Liberal Arts Computer Science Programs (LACS). He has written and presented papers on a range of issues related to the undergraduate curriculum and the design and use of computer laboratories. He is the recipient of two recent National Science Foundation grants to create formal teaching and research laboratories and to study how they can be integrated into the undergraduate computer science curriculum.

Professor Schneider received his Ph.D. in computer science from the University of Wisconsin with a research specialization in computer networks, distributed systems, and parallel processing.

Steven C. Bruell

Steven C. Bruell received the B.A. degree in mathematics (with honors) from the University of Texas at Austin in 1973, and the M.S. and Ph.D. degrees in computer science from Purdue University, West Lafayette, Indiana, in 1975 and 1978, respectively.

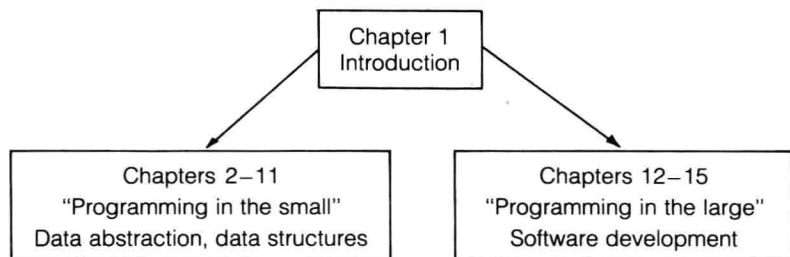
He is currently an Associate Professor of computer science at the University of Iowa, Iowa City. His current research interests include distributed systems, parallel simulation, and computer performance evaluation. He has coauthored many papers and three books. He has been an ACM National Lecturer and is currently serving as the Associate Chair of the Computer Science Department.

PREFACE

This book is intended for the second course in computer science, the course called CS 2 by the ACM in its curriculum recommendations. It assumes that the reader has completed a traditional first course that introduced programming in a high-level language, most likely Pascal. In this text we investigate the topics of data structures and software development.

The material is divided into two parts. Part I (chapters 2 through 11) covers abstract data types, data structures, recursion, and the analysis of algorithms. This section can be viewed as a study of “programming in the small” because it treats issues related to the construction of correct, efficient, and maintainable program units. Part II (chapters 12 through 15) addresses topics related to the software life cycle—specification, design, implementation, validation, and maintenance. This part can be viewed as a study of “programming in the large,” since it looks at issues related to the construction of software systems made up of hundreds or thousands of individual units of the type discussed in Part I.

Pictorially, the organization of this book is:



The placement of the discussion on software development following the data abstraction and data structures chapters does not imply that the material must be presented in that order. We have tried to keep these two sections independent from a teaching perspective, and the instructor who wishes to reverse the order of presentation and treat the material on software development first can do so without difficulty.

The material in Part I of this text is a continuation of ideas begun in CS 1. In that first course, students focused on the construction of correct, robust, and reliable program units through a study of algorithm design, structured programming, data types, elementary data structures, procedural abstraction, empirical testing, and programming style. In Part I we continue that investigation and present concepts that permit the design of program units with increased robustness, efficiency, shareability, and maintainability. The topics covered include the specification, design, and implementation of abstract data types (chapter 2) and advanced data structures such as stacks, queues, strings (chapters 3 and 4), trees (chapters 7 and 8), sets (chapter 9), and graphs (chapter 10). The topics of sorting and searching are woven throughout the chapters of Part I, but in chapter 11, we present a separate discussion on the important topic of external sorting and searching. In chapter 5, we introduce recursion so that the student can understand and utilize recursively defined data structures and feel comfortable using recursion as a general problem-solving paradigm. Finally, in chapter 6, we provide the mathematical tools needed for analyzing the time and space efficiency of algorithms, including sequential, recursive, and parallel algorithms.

When presenting the topic of data structures in Part I, we have been guided by and strongly influenced by the following four principles:

1. A student should first study a data structure in terms of its high-level behavior before becoming enmeshed in the details of its implementation. It is not that these implementation details are unimportant. However, the separation of the external and internal properties of a data structure allows us to look first at the formal properties of a data type, such as the nature of operations on objects of that type. Only after these fundamental properties are fully understood do we move on and look at how this structure can be implemented in a typical high-level language. The tool that allows us to create this conceptual separation is the abstract data type (abbreviated ADT), and it is the central theme of Part I. It is the method used to present the data structures studied in chapters 2 through 11. We first study a data structure as an abstract data type, concentrating only on its formal properties and its operations. Later we look at ways in which this structure can be implemented, and we investigate efficiency issues related to the choice of implementation.
2. One of the shortcomings of computer science has been a resistance to the widespread adoption of formal notation to represent key concepts. For example, in CS 1 it is common to describe the semantics of Pascal statements using natural language. However, natural language can be imprecise and highly ambiguous, and its use can lead to inconsistencies, misinterpretations, and incorrect implementations. Just as the freshman physics major learns to use calculus to express fundamental laws of physics, so must the CS student learn to use precise notation to express the “laws” of computer science. Therefore, to describe the behavior of operations on an abstract data type, we use the method of *axiomatic semantics*, which is a powerful, albeit complex, technique for describing the actions of operations on data objects. While this text contains only an introduction (with additional coverage coming in advanced courses in compilers and formal languages), it should make the student aware of the importance of precise and accurate notation. In addition to axiomatic semantics, we also use Big-O notation for the time and space analysis of algorithms (chapter 6), decision tables and finite state automata for the specification

of software (chapter 13), and a directed graph representation of program design documents (chapter 14).

3. There are many different data structures that can be covered in a CS 2 course, from arrays to AVL trees, from heaps to hash tables. Because of the large number of topics that can be studied, any treatment of this subject should initially provide an overall structure or classification scheme. If this classification is done well, the student will not view the subsequent discussion as a large and apparently unrelated collection of facts but as an organized and related body of knowledge. (The problem is identical to the one faced by a biologist teaching a first course. He or she first presents the well-known taxonomy of living organisms (kingdom, phylum, class, order, . . .) rather than immediately presenting a “Noah’s Ark” of animal types. This latter approach would quickly collapse into chaos.) In chapter 3, before beginning our discussion of data structures, we present a taxonomy of data structures based on the following four classes: sets, linear types, hierarchical types, and graphs (Figure 3.1). We then show why all data structures, regardless of what properties they may possess, fall into one of these four classes, thus giving the student a much better idea of how these ideas fit together (the “big picture”).

4. Finally, and most controversially, we made the decision not to use a specific high-level language to present our programming examples. Instead, we have chosen to use the pseudocode shown in Figure 2–16. While this pseudocode is closely related to a number of existing languages (especially Modula-2 and Ada), it provides these capabilities in a much more simplified manner. There will no doubt be those who ask why we did not select —, where the missing word is the name of the individual’s favorite language, such as Modula-2, Ada, Turbo Pascal, or C++. To completely and accurately describe the abstract data type capabilities of a “real” language, such as Ada packages, would take many pages of language- and system-specific information. This material would dwell on highly technical syntactic and semantic details such as semicolons, indentation, reserved words, and programming environments. But this approach is exactly opposite from what we stress in this book—that the ideas and concepts we are presenting are language-, system-, and vendor-independent, and that one can study them at a high level of abstraction without worrying about the specific language into which they will be translated (the very definition of abstraction!). This is not a book on advanced concepts in Modula-2 or Ada or — but a text on advanced concepts in computer science.

However, computer science is an applied as well as a theoretical discipline, and it is important for a student to implement the abstract ideas presented in a course. It is through this hands-on experience that a student can gain a deeper understanding of these concepts. To support this empirical experience, we have provided a series of language-specific supplements that are coordinated with the material in the text. These supplements describe an actual high-level programming language that can then be used to implement the examples, case studies, and homework assignments. Alternatively, an instructor could choose to provide this language- and system-related information through reference manuals, other texts, or locally prepared handouts. Regardless of the approach, however, it is important to send a message to the student that this is *not* just a language course or a programming course, but a course in computer science. The language that will ultimately be used, so important in CS 1, is more properly viewed in CS 2 as material supplementary to the central discussion.

Part II of the text, chapters 12 to 15, is an overview of the software life cycle. It focuses on those life cycle phases not stressed in the student's previous class work. Four chapters cannot possibly cover software development in great detail, and this text does not purport to replace a full semester course on the subject of software engineering. However, many students graduate in computer science without completing such a course and with little or no introduction to the life cycle of a software project. Part II of the text provides that needed introduction.

In chapter 12 we overview the eight stages in the software life cycle, from the initial rough problem statement to finished documentation and ongoing maintenance. Some of these eight stages may be inappropriate for extended discussion in a CS 2 course, for example, the feasibility study can require a knowledge of economics, accounting, and marketing as well as computer science. Some software development stages have been studied extensively in CS 1, namely coding and debugging. However, the remaining steps in the software life cycle are both highly appropriate to discuss and generally unfamiliar to students in a second course. They are presented in the next three chapters.

Chapter 13 discusses the topics of problem specification and the content of the problem specification document, including functional, error, and performance specifications. As we mentioned earlier, natural language is not the most precise way to express user needs. Therefore, in this chapter we look at two formal alternatives to natural language—decision tables and finite-state automata. These methods can provide the precision needed to develop high-quality specifications.

Chapter 14 treats the program design phase. Here we use the top-down problem-solving strategy called *divide and conquer* to select and specify the program units and data structures needed in our solution. The relationships between these program units will be expressed as a structure chart using a formal notation based on directed graphs. This chapter also describes the information contained in the program design document and presents guidelines for evaluating the quality of a proposed design.

Finally, chapter 15 looks at the two validation and verification (V & V) techniques called *empirical testing* and *formal verification*. We review the strengths and weaknesses of each of these approaches and conclude by describing a third method called *structured walkthroughs*. Regardless of which method is used, the goal of all V & V techniques is the same—to increase the confidence with which we can assert the correctness of each individual program unit as well as the correctness of the software system as a whole.

This is the outline of the text. As the preface clearly shows, it covers a good deal of material and ranges over many subject areas. However, its central themes can be clearly and simply stated: To improve the students' ability to specify, design, and solve real-world problems and to deepen their understanding and appreciation of the discipline of computer science.

Acknowledgments

I would like to thank the people who assisted me during the writing of this text. This certainly includes all the reviewers who read the initial drafts so carefully and made

many helpful suggestions and comments. Their ideas were invaluable to me during the rewriting and revising of the manuscript. Alphabetically, these reviewers were

Douglas Bickerstaff	C. Ray Russell
Eastern Washington University	Virginia Commonwealth University
Thomas E. Byther	Paul S. Schnare
University of Maine	Eastern Kentucky University
Maurice Eggen	Greg Scragg
Trinity University	State University of New York
Henry Etlinger	at Geneseo
Rochester Institute of Technology	Theodore Sjoerdsma
Charles E. Frank	Washington and Lee University
Northern Kentucky University	Jeffrey Slomka
Michael Hennessy	Southwest Texas State University
University of Oregon	Greg Starling
Lawrence A. Jehn	University of Arkansas, Fayetteville
University of Dayton	Stan Thomas
Kenneth Modesitt	Wake Forest University
Western Kentucky University	Susan R. Wallace
William R. Nico	University of North Florida
California State University at Hayward	
Alex T. Pregel, Jr.	
Brandeis University	

A special thanks to Mr. David Sielaff who helped prepare and debug all the algorithms and code in the text and I am deeply grateful. Thanks also to Ms. Michelle Villinski who typed the original manuscript and did an outstanding job. Finally, I would like to thank the people at West Publishing who worked with me during the preparation and production of the book. They are professionals in every sense of the word, and they were a pleasure to work with: Mark Jacobsen, Denis Ralling, and, most of all, Jerry Westby, whom I consider to be not only my editor but also a good friend.

G. Michael Schneider
Macalester College

In addition to those named above, I would like to thank Pam Betzel, Jean Bartley, and Wen Zhang.

S.C. Bruell
University of Iowa
August, 1990

CONTENTS

Preface xiii

Chapter 1 An Introduction to Advanced Programming Concepts 1

- 1.1 Introduction 2
- 1.2 Part I: Abstract Data Types and Advanced Data Structures 3
 - 1.2.1 Abstract Data Types 3
 - 1.2.2 Advanced Data Structures 5
 - 1.2.3 Recursion 6
 - 1.2.4 The Analysis of Algorithms 7
- 1.3 Part II: Software Development and the Software Life Cycle 8
 - 1.3.1 Specification and Design 8
 - 1.3.2 Testing and Formal Verification 10
- 1.4 Conclusion 11

PART I. ABSTRACT DATA TYPES, ADVANCED DATA STRUCTURES, RECURSION, AND THE ANALYSIS OF ALGORITHMS 13

Chapter 2 Abstract Data Types 15

- 2.1 Introduction to Abstract Data Types 16
- 2.2 Examples of Abstract Data Types 19
 - 2.2.1 Water Tank: An Atomic Data Type 19
 - 2.2.2 Train: A Structured Data Type 22
- 2.3 The Syntax and Semantics of Abstract Data Types 24
 - 2.3.1 A Formal Syntax for Abstract Data Types 25
 - 2.3.2 Formal Semantics for Abstract Data Types 27
 - 2.3.3 Pre- and Postconditions 33
- 2.4 Language Support for Abstract Data Types 35
 - 2.4.1 Required Language Constructs 35
 - 2.4.2 High-Level Language Implementation of the Water Tank Abstract Data Type 43
 - 2.4.3 Data Abstraction Facilities in Existing Languages 43

- 2.5 Advantages of Data Abstraction 50
- 2.6 Conclusion 52

Chapter 3 Linear Data Structures and Their Array Implementation 57

- 3.1 The Classification of Data Structures 58
- 3.2 Arrays 59
 - 3.2.1 The Abstract Data Type Array 60
 - 3.2.2 Implementation of Arrays 62
 - 3.2.3 Strings 67
- 3.3 Stacks 72
 - 3.3.1 Formal Description 72
 - 3.3.2 Applications of Stacks 79
- 3.4 Queues 84
- 3.5 Array Implementation of a Stack 91
- 3.6 Array Implementation of a Queue 93
- 3.7 Case Study: Disk Request Queues 96
- 3.8 Conclusion 103

Chapter 4 Linear Data Structures and Their Linked List Implementation 110

- 4.1 The Linear List Abstract Data Type 111
 - 4.1.1 Definition of a Linear List 111
 - 4.1.2 Operations on Linear Lists 112
- 4.2 The Implementation of Linear Lists 115
 - 4.2.1 Pointers in High-Level Programming Languages 118
 - 4.2.2 The Linked List Implementation of Linear List Operations 121
 - 4.2.3 The Array Implementation of Linear List Operations 128
- 4.3 Linked List Implementation of a Stack 131
- 4.4 Linked List Implementation of a Queue 134
- 4.5 Variations on the Linked List Structure 137
 - 4.5.1 Circular Lists 137
 - 4.5.2 Doubly Linked Lists 138
- 4.6 Case Study: A Memory Manager 141
- 4.7 Conclusion 145

Chapter 5 Recursion 157

- 5.1 Introduction 158
- 5.2 The Basic Concepts of Recursive Algorithms 160
- 5.3 How Recursion Is Implemented 164
- 5.4 Examples 167
 - 5.4.1 Sequential Search of a Linked List 167
 - 5.4.2 Reversing a Linked List 169
 - 5.4.3 Recursive Binary Search 171
 - 5.4.4 MergeSort 173

- 5.4.5 QuickSort 175
- 5.4.6 Finding a Path Through a Graph 179
- 5.5 Conclusion 185

Chapter 6 The Analysis of Algorithms 189

- 6.1 Introduction 190
- 6.2 The Asymptotic Analysis of Algorithms 190
- 6.3 Examples of Time Complexity Analysis 193
- 6.4 Examples of Space Complexity Analysis 204
- 6.5 The Analysis of Recursive Algorithms 206
- 6.6 The Analysis of Parallel Algorithms 210
- 6.7 Conclusion 215

Chapter 7 Trees and Binary Trees 221

- 7.1 Introduction 222
- 7.2 General Trees 224
- 7.3 Binary Trees 228
 - 7.3.1 Introduction 228
 - 7.3.2 Operations on Binary Trees 229
 - 7.3.3 The Binary Tree Representation of General Trees 236
 - 7.3.4 The Linked List Implementation of Binary Trees 238
 - 7.3.5 The Array Implementation of Binary Trees 242
- 7.4 Case Study: Expression Trees 245
- 7.5 Conclusion 254

Chapter 8 Special Purpose Tree Structures 260

- 8.1 Introduction 261
- 8.2 Binary Search Trees 261
 - 8.2.1 Definition 261
 - 8.2.2 Using Binary Search Trees in Searching Operations 262
 - 8.2.3 TreeSort 268
- 8.3 Indexed Search Trees 269
- 8.4 Heaps 274
 - 8.4.1 Definition 274
 - 8.4.2 Implementation of Heaps Using Arrays 277
 - 8.4.3 Application of Heaps 281
- 8.5 Case Study: A Spelling Checker 285
- 8.6 Conclusion 294

Chapter 9 Sets and Search Tables 299

- 9.1 Sets 300
 - 9.1.1 Definition and Operations 300
 - 9.1.2 Implementation 302
- 9.2 Search Tables 314
 - 9.2.1 Definition and Operations 314

- 9.2.2 Implementation Using Arrays and Linked Lists 316
- 9.2.3 Hashing 318
- 9.3 Case Study: Symbol Tables for Block-Structured Programming Languages 329
- 9.4 Conclusion 341

Chapter 10 Graph Structures 346

- 10.1 Introduction and Definitions 347
- 10.2 Operations on Graphs 350
 - 10.2.1 Basic Creation, Insertion, and Deletion Operations 350
 - 10.2.2 Traversal 352
 - 10.2.3 Minimal Spanning Trees 357
 - 10.2.4 Shortest Paths 361
- 10.3 Implementation 368
 - 10.3.1 The Adjacency Matrix Representation 368
 - 10.3.2 The Adjacency List Representation 370
- 10.4 Case Study: Computer Networks 372
- 10.5 Conclusion 380

Chapter 11 External Sorting and Searching 385

- 11.1 Introduction 386
- 11.2 External Storage Media 387
- 11.3 External Sorting Using MergeSort 391
- 11.4 External Searching 400
 - 11.4.1 Formal Definition of m -Way Search Trees 400
 - 11.4.2 Complexity Analysis of m -Way Search Trees 404
 - 11.4.3 Other Approaches to External Searching 408
- 11.5 Conclusion 413

PART II. LARGE SCALE SOFTWARE DEVELOPMENT AND THE SOFTWARE LIFE CYCLE 417

Chapter 12 An Overview of the Software Life Cycle 419

- 12.1 Introduction 420
- 12.2 The Steps Involved in Software Development 422
 - 12.2.1 The Feasibility Study 422
 - 12.2.2 Problem Specification 424
 - 12.2.3 High-Level Program Design 426
 - 12.2.4 Module Design and ADT Implementation (Low-Level Design) 427
 - 12.2.5 Implementation 427
 - 12.2.6 Testing, Verification, and Benchmarking 430
 - 12.2.7 Documentation 430
 - 12.2.8 Program Maintenance 431
- 12.3 Conclusion 433

Chapter 13 The Problem Specification Phase 436

- 13.1 Introduction 437
- 13.2 Contents of the Problem Specification Document 437
 - 13.2.1 Input/Output Specifications 437
 - 13.2.2 Exception Handling 442
 - 13.2.3 Performance Specifications 446
- 13.3 Formal Specification Methods 447
 - 13.3.1 Decision Tables 448
 - 13.3.2 Finite State Machines 452
- 13.4 Conclusion 453

Chapter 14 The Program Design and Implementation Phases 460

- 14.1 Introduction 461
- 14.2 Top-Down Program Design 462
- 14.3 The Program Design Document 464
 - 14.3.1 The Structure Chart 464
 - 14.3.2 Procedure Specifications 469
 - 14.3.3 The Data Dictionary 471
- 14.4 The Advantages of Top-Down Program Design 472
- 14.5 Design Evaluation 475
 - 14.5.1 Logical Coherence 475
 - 14.5.2 Independence 479
 - 14.5.3 Connectivity 483
 - 14.5.4 Size 484
- 14.6 Implementation Issues 485
 - 14.6.1 Phased Implementation/Testing 485
 - 14.6.2 Top-Down versus Bottom-Up Implementation 489
- 14.7 Conclusion 490

Chapter 15 The Testing and Verification of Software 497

- 15.1 Introduction 498
- 15.2 Empirical Testing 498
 - 15.2.1 Unit Testing 498
 - 15.2.2 System Integration 505
 - 15.2.3 Acceptance Testing 506
- 15.3 Formal Verification 507
- 15.4 Informal Verification 514
- 15.5 Conclusion 516

AN INTRODUCTION TO ADVANCED PROGRAMMING CONCEPTS

CHAPTER OUTLINE

- 1.1 Introduction
- 1.2 Part I: Abstract Data Types and Advanced Data Structures
 - 1.2.1 Abstract Data Types
 - 1.2.2 Advanced Data Structures
 - 1.2.3 Recursion
 - 1.2.4 The Analysis of Algorithms
- 1.3 Part II: Software Development and the Software Life Cycle
 - 1.3.1 Specification and Design
 - 1.3.2 Testing and Formal Verification
- 1.4 Conclusion