

国外计算机科学教材系列

Pearson


★ William Stallings

操作系统

—— 精髓与设计原理 (第八版)

Operating Systems

Internals and Design Principles, Eighth Edition



英文版

[美] William Stallings 著




中国工信出版集团



电子工业出版社
PUBLISHING HOUSE OF ELECTRONICS INDUSTRY
<http://www.phei.com.cn>

国外计算机科学教材系列



操作系统

——精髓与设计原理

(第八版)(英文版)

Operating Systems: Internals and Design Principles, Eighth Edition

[美] William Stallings 著

电子工业出版社
Publishing House of Electronics Industry
北京·BEIJING

内 容 简 介

本书既是关于操作系统概念、结构和机制的教材,目的是尽可能清楚和全面地展示现代操作系统的本质和特点;也是讲解操作系统的经典教材,不仅系统地讲述了操作系统的基本概念、原理和方法,而且以当代最流行的操作系统 Windows 8、UNIX、Android、Linux 为例,展现了当代操作系统的本质和特点。全书共分背景知识、进程、内存、调度、输入/输出和文件、嵌入式系统六部分,内容包括计算机系统概述、操作系统概述、进程描述和控制、线程、并发性:互斥和同步、并发:死锁和饥饿、内存管理、虚拟内存、单处理器调度、多处理器和实时调度、I/O 管理和磁盘调度、文件管理、嵌入式操作系统、虚拟机、计算机安全技术、分布式处理、客户/服务器和集群等。此外,本书配套网站提供了及时、生动的材料。

本书可作为高等学校计算机科学与技术专业操作系统课程的教材,也可供其他专业学生参考。

Original edition, entitled Operating Systems: Internals and Design Principles, Eighth Edition, ISBN: 9780133805918 by William Stallings. Published by Pearson Education, Inc. Copyright © 2015 Pearson Education, Inc.

All rights Reserved. No part of this book may be reproduced or transmitted in any forms or by any means, electronic or mechanical, including photocopying recording or by any information storage retrieval systems, without permission from Pearson Education, Inc.

English reprint edition published by PEARSON EDUCATION ASIA LTD, and PUBLISHING HOUSE OF ELECTRONICS INDUSTRY, Copyright © 2017.

This edition is authorized for sale only in the People's Republic of China (excluding the Special Administrative Region of Hong Kong and Macau).

本书英文影印版专有版权由 Pearson Education(培生教育出版集团)授予电子工业出版社,未经出版者预先书面许可,不得以任何方式复制或抄袭本书的任何部分。

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

版权贸易合同登记号 图字:01-2017-4085

图书在版编目(CIP)数据

操作系统:精髓与设计原理:英文/(美)威廉·斯托林斯(William Stallings)著.—8版.

北京:电子工业出版社,2017.6

书名原文:Operating Systems: Internals and Design Principles, Eighth Edition

国外计算机科学教材系列

ISBN 978-7-121-31781-1

I. ①操… II. ①威… III. ①操作系统-高等学校-教材-英文 IV. ①TP316

中国版本图书馆 CIP 数据核字(2017)第 124167 号

策划编辑:谭海平

责任编辑:谭海平

印 刷:北京京师印务有限公司

装 订:北京京师印务有限公司

出版发行:电子工业出版社

北京市海淀区万寿路 173 信箱 邮编 100036

开 本:787×980 1/16 印张:49 字数:1164 千字

版 次:2017 年 6 月第 1 版(原著第 8 版)

印 次:2017 年 6 月第 1 次印刷

定 价:98.00 元

凡所购买电子工业出版社图书有缺损问题,请向购买书店调换。若书店售缺,请与本社发行部联系,联系及邮购电话:(010)88254888, 88258888。

质量投诉请发邮件至 zlts@phei.com.cn, 盗版侵权举报请发邮件至 dbqq@phei.com.cn。

本书咨询联系方式:(010)88254552, tan02@phei.com.cn。

PREFACE

WHAT'S NEW IN THE EIGHTH EDITION

Since the seventh edition of this book was published, the field has seen continued innovations and improvements. In this new edition, I try to capture these changes while maintaining a broad and comprehensive coverage of the entire field. To begin the process of revision, the seventh edition of this book was extensively reviewed by a number of professors who teach the subject and by professionals working in the field. The result is that, in many places, the narrative has been clarified and tightened, and illustrations have been improved.

Beyond these refinements to improve pedagogy and user friendliness, the technical content of the book has been updated throughout, to reflect the ongoing changes in this exciting field, and the instructor and student support has been expanded. The most noteworthy changes are as follows:

- **Windows 8:** Windows 8 is Microsoft's latest OS offering for PCs, workstations, and servers, which includes a number of changes to the internal architecture. The new edition provides details on Windows 8 internals in all of the key technology areas covered in this book, including process/thread management, scheduling, memory management, security, file systems, and I/O.
- **Android operating system:** Android is the fastest growing mobile platform. The real-world constraints and operating environment of mobile devices are quite different from traditional desktop or server computers. It is important for students to learn this new environment.
- **Embedded Linux:** The use of a minimal version of Linux for embedded systems has grown in popularity. This new edition provides an overview of the key elements of the embedded Linux approach.
- **Virtual machines:** Server virtualization and other forms of virtual machines are becoming increasingly widespread. A new chapter deals with the operating system design issues for virtual machines.
- **Multicore design issues:** The dominant computer architecture is now multicore. This raises new OS design issues that are addressed in this new edition.
- **I/O standards:** The book has been updated to reflect the latest developments, including Thunderbolt.
- **Storage hardware:** The discussion of storage hardware has been updated and now includes discussion of solid-state drives.
- **Fault tolerance:** The ACM/IEEE Computer Science Curricula 2013 lists fault tolerance as one of the core topics for an OS course. A new section provides an overview of fault tolerance.

OBJECTIVES

This book is about the concepts, structure, and mechanisms of operating systems. Its purpose is to present, as clearly and completely as possible, the nature and characteristics of modern-day operating systems.

This task is challenging for several reasons. First, there is a tremendous range and variety of computer systems for which operating systems are designed. These include embedded systems, smart phones, single-user workstations and personal computers, medium-sized shared systems, large mainframe and supercomputers, and specialized machines such as real-time systems. The variety is not just in the capacity and speed of machines, but in applications and system support requirements. Second, the rapid pace of change that has always characterized computer systems continues with no letup. A number of key areas in operating system design are of recent origin, and research into these and other new areas continues.

In spite of this variety and pace of change, certain fundamental concepts apply consistently throughout. To be sure, the application of these concepts depends on the current state of technology and the particular application requirements. The intent of this book is to provide a thorough discussion of the fundamentals of operating system design and to relate these to contemporary design issues and to current directions in the development of operating systems.

EXAMPLE SYSTEMS

This text is intended to acquaint the reader with the design principles and implementation issues of contemporary operating systems. Accordingly, a purely conceptual or theoretical treatment would be inadequate. To illustrate the concepts and to tie them to real-world design choices that must be made, four operating systems have been chosen as running examples:

- **Windows 8:** A multitasking operating system for personal computers, workstations, servers, and mobile devices. This operating system incorporates many of the latest developments in operating system technology. In addition, Windows is one of the first important commercial operating systems to rely heavily on object-oriented design principles. This book covers the technology used in the most recent versions of Windows, known as Windows 8.
- **Android:** Android tailored for embedded devices, especially mobile phones. Focusing on the unique requirements of the embedded environment, the book provides details of Android internals.
- **UNIX:** A multiuser operating system, originally intended for minicomputers, but implemented on a wide range of machines from powerful microcomputers to supercomputers. Several flavors of UNIX are included as examples. FreeBSD is a widely used system that incorporates many state-of-the-art features. Solaris is a widely used commercial version of UNIX.
- **Linux:** An open-source version of UNIX that is now widely used.

These systems were chosen because of their relevance and representativeness. The discussion of the example systems is distributed throughout the text rather than assembled as a single chapter or appendix. Thus, during the discussion of concurrency, the concurrency mechanisms of each example system are described, and the motivation for the individual design choices is discussed. With this approach, the design concepts discussed in a given chapter are immediately reinforced with real-world examples. For convenience all of the material for each of the example systems is also available as an online document.

SUPPORT OF ACM/IEEE COMPUTER SCIENCE CURRICULA 2013

The book is intended for both an academic and a professional audience. As a textbook, it is intended as a one-semester or two-semester undergraduate course in operating systems for computer science, computer engineering, and electrical engineering majors. This edition is designed to support the recommendations of the current (February 2013) draft version of the ACM/IEEE Computer Science Curricula 2013 (CS2013). The CS2013 curriculum recommendation includes Operating Systems (OSs) as one of the Knowledge Areas in the Computer Science Body of Knowledge. CS2013 divides all course work into three categories: Core-Tier 1 (all topics should be included in the curriculum), Core-Tier-2 (all or almost all topics should be included), and Elective (desirable to provide breadth and depth). In the OS area, CS2013 includes two Tier 1 topics, four Tier 2 topics, and 6 Elective topics, each of which has a number of subtopics. This text covers all of the topics and subtopics listed by CS2013 in these three categories.

Table P.1 shows the support for the OS Knowledge Area provided in this textbook. A detailed list of subtopics for each topic is available as the file CS2013-OS.pdf at box.com/OS8e.

PLAN OF THE TEXT

The book is divided into five parts (see Chapter 0 for an overview):

- Background
- Processes
- Memory
- Scheduling
- Advanced topics (embedded OSs, virtual machines, OS security, and distributed systems)

The book includes a number of pedagogic features, including the use of animations and video notes and numerous figures and tables to clarify the discussion. Each chapter includes a list of key words, review questions, homework problems, and suggestions for further reading. The book also includes an extensive glossary, a list of frequently used acronyms, and a bibliography. In addition, a test bank is available to instructors.

Table P.1 Coverage of CS2013 Operating Systems (OSs) Knowledge Area

Topic	Coverage in Book
Overview of Operating Systems (Tier 1)	Chapter 2: Operating System Overview
Operating System Principles (Tier 1)	Chapter 1: Computer System Overview Chapter 2: Operating System Overview
Concurrency (Tier 2)	Chapter 5: Mutual Exclusion and Synchronization Chapter 6: Deadlock and Starvation Appendix A: Topics in Concurrency Chapter 18: Distributed Process Management
Scheduling and Dispatch (Tier 2)	Chapter 9: Uniprocessor Scheduling Chapter 10: Multiprocessor and Real-Time Scheduling
Memory Management (Tier 2)	Chapter 7: Memory Management Chapter 8: Virtual Memory
Security and Protection (Tier 2)	Chapter 15: Operating System Security
Virtual Machines (Elective)	Chapter 14: Virtual Machines
Device Management (Elective)	Chapter 11: I/O Management and Disk Scheduling
File System (Elective)	Chapter 12: File Management
Real Time and Embedded Systems (Elective)	Chapter 10: Multiprocessor and Real-Time Scheduling Chapter 13: Embedded Operating Systems Material on Android throughout the text
Fault Tolerance (Elective)	Section 2.5: Fault Tolerance
System Performance Evaluation (Elective)	Performance issues related to memory management, scheduling, and other areas addressed throughout the text

INSTRUCTOR SUPPORT MATERIALS

The major goal of this text is to make it as effective a teaching tool for this fundamental yet evolving subject as possible. This goal is reflected both in the structure of the book and in the supporting material. The text is accompanied by the following supplementary material to aid the instructor:

- **Solutions manual:** Solutions to end-of-chapter Review Questions and Problems.
- **Projects manual:** Suggested project assignments for all of the project categories listed in this Preface.
- **PowerPoint slides:** A set of slides covering all chapters, suitable for use in lecturing.
- **PDF files:** Reproductions of all figures and tables from the book.
- **Test bank:** A chapter-by-chapter set of questions with a separate file of answers.

- **VideoNotes on concurrency:** Professors perennially cite concurrency as perhaps the most difficult concept in the field of operating systems for students to grasp. The edition is accompanied by a number of VideoNotes lectures discussing the various concurrency algorithms defined in the book.
- **Sample syllabuses:** The text contains more material than can be conveniently covered in one semester. Accordingly, instructors are provided with several sample syllabuses that guide the use of the text within limited time. These samples are based on real-world experience by professors with the seventh edition.

All of these support materials are available at the **Instructor Resource Center (IRC)** for this textbook, which can be reached through the publisher's Web site www.pearsonhighered.com/stallings or by clicking on the link labeled *Pearson Resources for Instructors* at this book's Companion Web site at WilliamStallings.com/OperatingSystems. To gain access to the IRC, please contact your local Pearson sales representative via pearsonhighered.com/educator/relocator/requestSalesRep, page or call Pearson Faculty Services at 1-800-526-0485.

The **Companion Web site**, at WilliamStallings.com/OperatingSystems (click on *Instructor Resources* link), includes the following:

- Links to Web sites for other courses being taught using this book.
- Sign-up information for an Internet mailing list for instructors using this book to exchange information, suggestions, and questions with each other and with the author.

PROJECTS AND OTHER STUDENT EXERCISES

For many instructors, an important component of an OS course is a project or set of projects by which the student gets hands-on experience to reinforce concepts from the text. This book provides an unparalleled degree of support for including a projects component in the course. In the online portion of the text, two major programming projects are defined. In addition, the instructor's support materials available through Prentice Hall not only includes guidance on how to assign and structure the various projects but also includes a set of user's manuals for various project types plus specific assignments, all written especially for this book. Instructors can assign work in the following areas:

- **OS/161 projects:** Described below.
- **Simulation projects:** Described below.
- **Programming projects:** Described below.
- **Research projects:** A series of research assignments that instruct the student to research a particular topic on the Internet and write a report.
- **Reading/report assignments:** A list of papers that can be assigned for reading and writing a report, plus suggested assignment wording.

- **Writing assignments:** A list of writing assignments to facilitate learning the material.
- **Discussion topics:** These topics can be used in a classroom, chat room, or message board environment to explore certain areas in greater depth and to foster student collaboration.

In addition, information is provided on a software package known as BACI that serves as a framework for studying concurrency mechanisms.

This diverse set of projects and other student exercises enables the instructor to use the book as one component in a rich and varied learning experience and to tailor a course plan to meet the specific needs of the instructor and students. See Appendix B in this book for details.

OS/161

This edition provides support for an active learning component based on OS/161. OS/161 is an educational operating system that is becoming increasingly recognized as the preferred teaching platform for OS internals. It aims to strike a balance between giving students experience in working on a real operating system and potentially overwhelming students with the complexity that exists in a fully-fledged operating system, such as Linux. Compared to most deployed operating systems, OS/161 is quite small (approximately 20,000 lines of code and comments), and therefore it is much easier to develop an understanding of the entire code base.

The IRC includes:

1. A packaged set of html files that the instructor can upload to a course server for student access.
2. A getting-started manual to be handed out to students to help them begin using OS/161.
3. A set of exercises using OS/161, to be handed out to students.
4. Model solutions to each exercise for the instructor's use.
5. All of this will be cross-referenced with appropriate sections in the book, so that the student can read the textbook material and then do the corresponding OS/161 project.

SIMULATIONS

The IRC provides support for assigning projects based on a set of seven **simulations** that cover key areas of OS design. The student can use a set of simulation packages to analyze OS design features. The simulators are all written in Java and can be run either locally as a Java application or online through a browser. The IRC includes specific assignments to give to students, telling them specifically what they are to do and what results are expected.

ANIMATIONS

This edition also incorporates animations. Animations provide a powerful tool for understanding the complex mechanisms of a modern OS. A total of 53 animations are used to illustrate key functions and algorithms in OS design. The animations are used for Chapters 3, 5, 6, 7, 8, 9, and 11. The animations are available at the Premium Web site for this book.

PROGRAMMING PROJECTS

This edition provides support for programming projects. Two major programming projects, one to build a shell, or command line interpreter, and one to build a process dispatcher are described in the online portion of this textbook. The IRC provides further information and step-by-step exercises for developing the programs.

As an alternative, the instructor can assign a more extensive series of projects that cover many of the principles in the book. The student is provided with detailed instructions for doing each of the projects. In addition, there is a set of homework problems, which involve questions related to each project for the student to answer.

Finally, the project manual provided at the IRC includes a series of programming projects that cover a broad range of topics and that can be implemented in any suitable language on any platform.

ONLINE DOCUMENTS, ANIMATIONS, AND VIDEONOTES FOR STUDENTS

For this new edition, a substantial amount of original supporting material for students has been made available online, at two Web locations. The **Companion Web site**, at WilliamStallings.com/OperatingSystems (click on *Student Resources* link), includes a list of relevant links organized by chapter and an errata sheet for the book.



Purchasing this textbook new also grants the reader six months of access to the **Premium Content site**, which includes the following materials:

- **Online chapters:** To limit the size and cost of the book, two chapters of the book, covering security, are provided in PDF format. The chapters are listed in this book's table of contents.
- **Online appendices:** There are numerous interesting topics that support material found in the text but whose inclusion is not warranted in the printed text. A total of 18 online appendices cover these topics for the interested student. The appendices are listed in this book's table of contents.

- **Homework problems and solutions:** To aid the student in understanding the material, a separate set of homework problems with solutions is available.
- **Animations:** Animations provide a powerful tool for understanding the complex mechanisms of a modern OS. A total of 53 animations are used to illustrate key functions and algorithms in OS design. The animations are used for Chapters 3, 5, 6, 7, 8, 9, and 11.
- **Videonotes:** VideoNotes are step-by-step video tutorials specifically designed to enhance the programming concepts presented in this textbook. The book is accompanied by a number of VideoNotes lectures discussing the various concurrency algorithms defined in the book.



To access the Premium Content site, click on the *Premium Content* link at the Companion Web site or at pearsonhighered.com/stallings and enter the student access code found on the card in the front of the book.

ACKNOWLEDGMENTS

Through its multiple editions, this book has benefited from review by hundreds of instructors and professionals, who gave generously of their time and expertise. Here I acknowledge those whose help contributed to this latest edition.

The following instructors reviewed all or a large part of the manuscript: Bernard Ku (Austin Community College, Northridge), Hemant Pendharker (Worcester State University), Nancy Birkenheuer (Regis University, Denver), Greg Ozbirn (University of Texas, Dallas).

Thanks also to the people who provided detailed reviews of the example systems. Reviews of the Android material were provided by Kristopher Micinski, Ron Munitz, Atte Peltomaki, Durgadoss Ramanathan, Manish Shakya, Samuel Simon, Wei Wang, and Chen Yang. The Linux reviewers were Tigran Aivazian, Kaiwan Billimoria, Peter Huewe, Manmohan Manoharan, Rami Rosen, Neha Naik, and Hualing Yu. The Windows material was reviewed by Francisco Cotrina, Sam Haidar, Christopher Kuleci, Benny Olsson, and Dave Probert. Nick Garnett of eCosCentric, for the review of the material on eCos, and Philip Levis, one of the developers of TinyOS, for the review of the material on TinyOS.

Andrew Peterson of the University of Toronto prepared the OS/161 supplements for the IRC. James Craig Burley authored and recorded the VideoNotes.

Adam Critchley (University of Texas at San Antonio) developed the simulation exercises. Matt Sparks (University of Illinois at Urbana-Champaign) adapted a set of programming problems for use with this textbook.

Lawrie Brown of the Australian Defence Force Academy produced the material on buffer overflow attacks. Ching-Kuang Shene (Michigan Tech University) provided the examples used in the section on race conditions and reviewed the section. Tracy Camp and Keith Hellman, both at the Colorado School of Mines, developed a new set of homework problems. In addition, Fernando Ariel Gont

contributed a number of homework problems; he also provided detailed reviews of all of the chapters.

I would also like to thank Bill Bynum (College of William and Mary) and Tracy Camp (Colorado School of Mines) for contributing Appendix O; Steve Taylor (Worcester Polytechnic Institute) for contributing the programming projects and reading/report assignments in the instructor's manual; and Professor Tan N. Nguyen (George Mason University) for contributing the research projects in the instruction manual. Ian G. Graham (Griffith University) contributed the two programming projects in the textbook. Oskars Rieksts (Kutztown University) generously allowed me to make use of his lecture notes, quizzes, and projects.

Finally, I thank the many people responsible for the publication of this book, all of whom did their usual excellent job. This includes the staff at Pearson, particularly my editor Tracy Johnson, her assistant Jenah Blitz-Stoehr, program manager Carole Snyder, and project manager Bob Engelhardt. I also thank Shiny Rajesh and the production staff at Integra for another excellent and rapid job. Thanks also to the marketing and sales staffs at Pearson, without whose efforts this book would not be in front of you.

ABOUT THE AUTHOR

Dr. William Stallings has authored 17 titles, and counting revised editions, over 40 books on computer security, computer networking, and computer architecture. His writings have appeared in numerous publications, including the *Proceedings of the IEEE*, *ACM Computing Reviews* and *Cryptologia*.

He has 12 times received the award for the best Computer Science textbook of the year from the Text and Academic Authors Association.

In over 30 years in the field, he has been a technical contributor, technical manager, and an executive with several high-technology firms. He has designed and implemented both TCP/IP-based and OSI-based protocol suites on a variety of computers and operating systems, ranging from microcomputers to mainframes. As a consultant, he has advised government agencies, computer and software vendors, and major users on the design, selection, and use of networking software and products.

He created and maintains the *Computer Science Student Resource Site* at ComputerScienceStudent.com. This site provides documents and links on a variety of subjects of general interest to computer science students (and professionals). He is a member of the editorial board of *Cryptologia*, a scholarly journal devoted to all aspects of cryptology. He is Networking Category Expert Writer at networking.answers.com.

Dr. Stallings holds a Ph.D. from M.I.T. in Computer Science and a B.S. from Notre Dame in electrical engineering.

CONTENTS

Online Resources viii

VideoNotes x

Preface xii

About the Author xxi

Chapter 0 Guide for Readers and Instructors 1

- 0.1 Outline of this Book 2
- 0.2 Example Systems 2
- 0.3 A Roadmap for Readers and Instructors 3
- 0.4 Internet and Web Resources 4

PART 1 BACKGROUND 7

Chapter 1 Computer System Overview 7

- 1.1 Basic Elements 8
- 1.2 Evolution of the Microprocessor 10
- 1.3 Instruction Execution 10
- 1.4 Interrupts 13
- 1.5 The Memory Hierarchy 24
- 1.6 Cache Memory 27
- 1.7 Direct Memory Access 31
- 1.8 Multiprocessor and Multicore Organization 32
- 1.9 Recommended Reading 36
- 1.10 Key Terms, Review Questions, and Problems 36
- 1A Performance Characteristics of Two-Level Memories 39

Chapter 2 Operating System Overview 46

- 2.1 Operating System Objectives and Functions 47
- 2.2 The Evolution of Operating Systems 52
- 2.3 Major Achievements 61
- 2.4 Developments Leading to Modern Operating Systems 70
- 2.5 Fault Tolerance 73
- 2.6 OS Design Considerations for Multiprocessor and Multicore 76
- 2.7 Microsoft Windows Overview 79
- 2.8 Traditional UNIX Systems 86
- 2.9 Modern UNIX Systems 89
- 2.10 Linux 91
- 2.11 Android 96
- 2.12 Recommended Reading and Animations 102
- 2.13 Key Terms, Review Questions, and Problems 103

PART 2 PROCESSES 105

Chapter 3 Process Description and Control 105

- 3.1 What Is a Process? 107
- 3.2 Process States 109
- 3.3 Process Description 124
- 3.4 Process Control 133
- 3.5 Execution of the Operating System 139
- 3.6 UNIX SVR4 Process Management 142
- 3.7 Summary 147
- 3.8 Recommended Reading and Animations 147
- 3.9 Key Terms, Review Questions, and Problems 148

Chapter 4 Threads 152

- 4.1 Processes and Threads 153
- 4.2 Types of Threads 159
- 4.3 Multicore and Multithreading 166
- 4.4 Windows 8 Process and Thread Management 171
- 4.5 Solaris Thread and SMP Management 178
- 4.6 Linux Process and Thread Management 182
- 4.7 Android Process and Thread Management 186
- 4.8 Mac OS X Grand Central Dispatch 190
- 4.9 Summary 193
- 4.10 Recommended Reading 193
- 4.11 Key Terms, Review Questions, and Problems 194

Chapter 5 Concurrency: Mutual Exclusion and Synchronization 199

- 5.1 Principles of Concurrency 202
- 5.2 Mutual Exclusion: Hardware Support 210
- 5.3 Semaphores 214
- 5.4 Monitors 227
- 5.5 Message Passing 233
- 5.6 Readers/Writers Problem 240
- 5.7 Summary 244
- 5.8 Recommended Reading and Animations 245
- 5.9 Key Terms, Review Questions, and Problems 246

Chapter 6 Concurrency: Deadlock and Starvation 259

- 6.1 Principles of Deadlock 260
- 6.2 Deadlock Prevention 269
- 6.3 Deadlock Avoidance 271
- 6.4 Deadlock Detection 277
- 6.5 An Integrated Deadlock Strategy 279
- 6.6 Dining Philosophers Problem 280
- 6.7 UNIX Concurrency Mechanisms 282
- 6.8 Linux Kernel Concurrency Mechanisms 286

- 6.9 Solaris Thread Synchronization Primitives 293
- 6.10 Windows 7 Concurrency Mechanisms 295
- 6.11 Android Interprocess Communication 299
- 6.12 Summary 300
- 6.13 Recommended Reading and Animations 301
- 6.14 Key Terms, Review Questions, and Problems 302

PART 3 MEMORY 309

Chapter 7 Memory Management 309

- 7.1 Memory Management Requirements 310
- 7.2 Memory Partitioning 314
- 7.3 Paging 325
- 7.4 Segmentation 328
- 7.5 Summary 330
- 7.6 Recommended Reading and Animations 330
- 7.7 Key Terms, Review Questions, and Problems 330
- 7A Loading and Linking 333

Chapter 8 Virtual Memory 340

- 8.1 Hardware and Control Structures 341
- 8.2 Operating System Software 358
- 8.3 UNIX and Solaris Memory Management 377
- 8.4 Linux Memory Management 383
- 8.5 Windows Memory Management 386
- 8.6 Android Memory Management 389
- 8.7 Summary 389
- 8.8 Recommended Reading and Animations 390
- 8.9 Key Terms, Review Questions, and Problems 391

PART 4 SCHEDULING 397

Chapter 9 Uniprocessor Scheduling 397

- 9.1 Types of Processor Scheduling 398
- 9.2 Scheduling Algorithms 402
- 9.3 Traditional UNIX Scheduling 424
- 9.4 Summary 426
- 9.5 Recommended Reading and Animations 427
- 9.6 Key Terms, Review Questions, and Problems 428

Chapter 10 Multiprocessor, Multicore, and Real-Time Scheduling 432

- 10.1 Multiprocessor and Multicore Scheduling 433
- 10.2 Real-Time Scheduling 446
- 10.3 Linux Scheduling 461
- 10.4 UNIX SVR4 Scheduling 465
- 10.5 UNIX FreeBSD Scheduling 466

vi CONTENTS

- 10.6 Windows Scheduling 470
- 10.7 Summary 472
- 10.8 Recommended Reading 473
- 10.9 Key Terms, Review Questions, and Problems 473

PART 5 INPUT/OUTPUT AND FILES 477

Chapter 11 I/O Management and Disk Scheduling 477

- 11.1 I/O Devices 478
- 11.2 Organization of the I/O Function 480
- 11.3 Operating System Design Issues 483
- 11.4 I/O Buffering 486
- 11.5 Disk Scheduling 489
- 11.6 RAID 496
- 11.7 Disk Cache 505
- 11.8 UNIX SVR4 I/O 509
- 11.9 Linux I/O 512
- 11.10 Windows I/O 515
- 11.11 Summary 517
- 11.12 Recommended Reading and Animations 518
- 11.13 Key Terms, Review Questions, and Problems 519

Chapter 12 File Management 522

- 12.1 Overview 523
- 12.2 File Organization and Access 529
- 12.3 B-Trees 533
- 12.4 File Directories 536
- 12.5 File Sharing 541
- 12.6 Record Blocking 542
- 12.7 Secondary Storage Management 544
- 12.8 UNIX File Management 552
- 12.9 Linux Virtual File System 557
- 12.10 Windows File System 561
- 12.11 Android File Management 566
- 12.12 Summary 567
- 12.13 Recommended Reading 568
- 12.14 Key Terms, Review Questions, and Problems 568

PART 6 EMBEDDED SYSTEMS 571

Chapter 13 Embedded Operating Systems 571

- 13.1 Embedded Systems 572
- 13.2 Characteristics of Embedded Operating Systems 573
- 13.3 Embedded Linux 576
- 13.4 TinyOS 577
- 13.5 Recommended Reading 587
- 13.6 Key Terms, Review Questions, and Problems 587