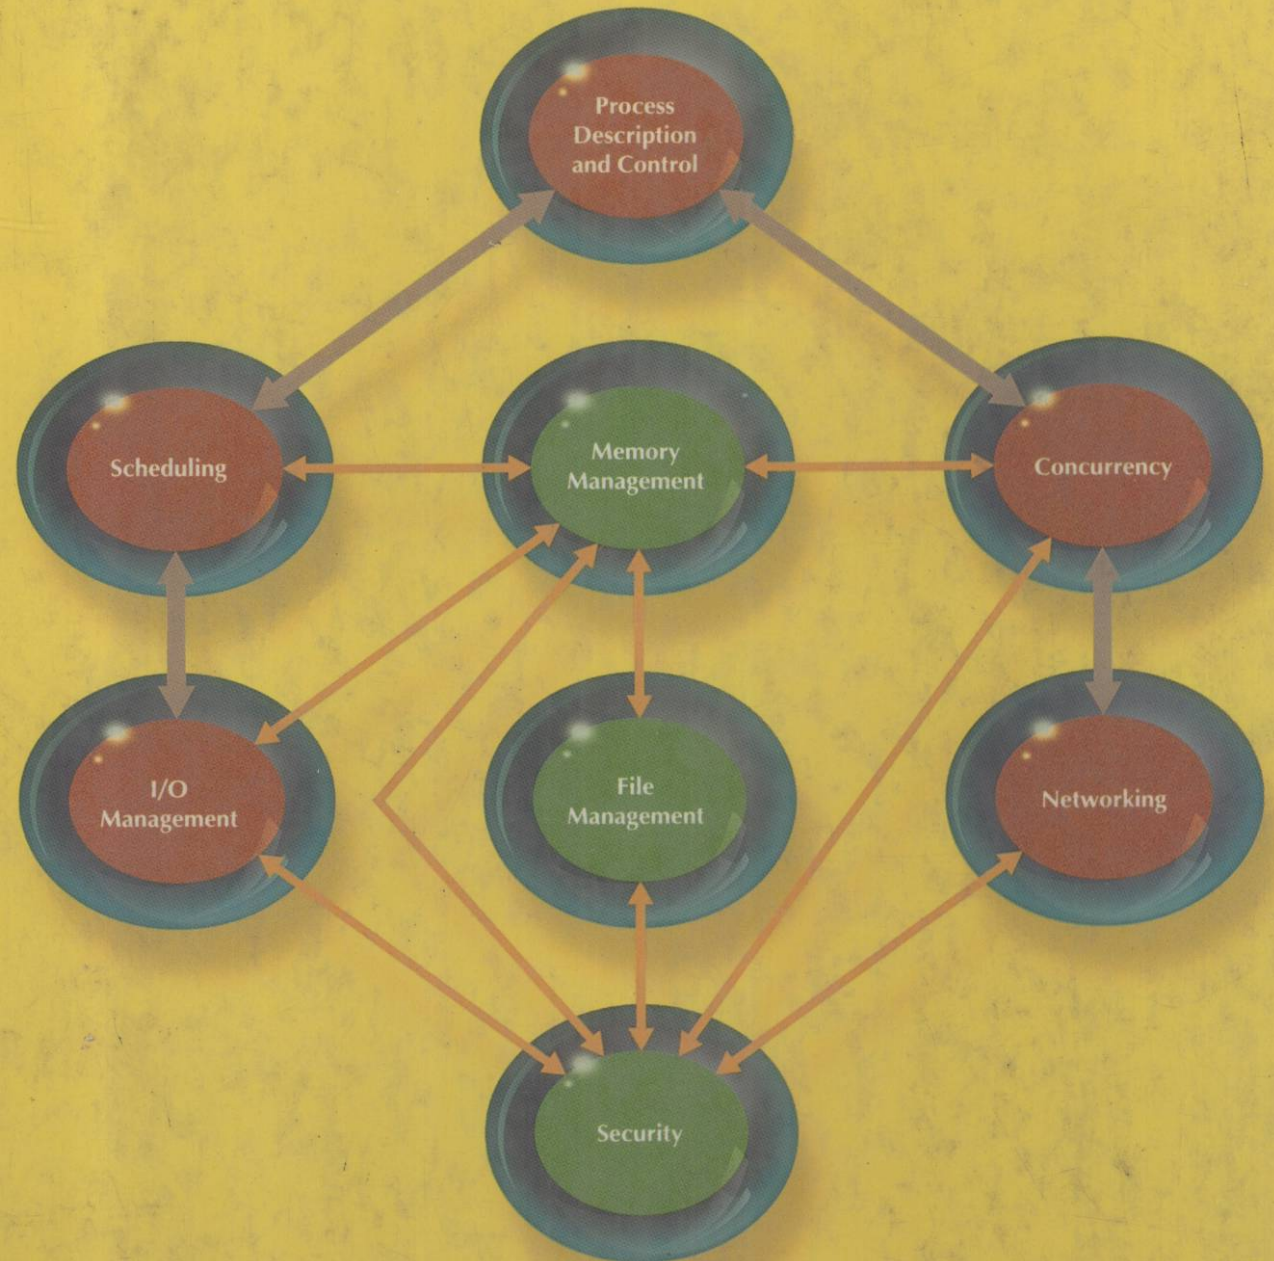


OPERATING SYSTEMS

WILLIAM STALLINGS





OPERATING SYSTEMS

William Stallings

MACMILLAN PUBLISHING COMPANY
New York

MAXWELL MACMILLAN CANADA
Toronto

MAXWELL MACMILLAN INTERNATIONAL
New York Oxford Singapore Sydney

Editor: John Griffin
Production Supervisor: John Travis
Production Manager: Paul Smolenski
Text Designer: Jane Edelstein
Cover Designer: Jane Edelstein
Cover Illustration: Brian Sheridan
Illustrations by Precision Graphics

This book was set in 10/12 Palatino by Waldman Graphics, and was printed and bound by R. R. Donnelley & Sons—Crawfordsville. The cover was printed by Lehigh Press Lithographers.

Manufacturers often claim as trademarks forms of their own name and the names they give to their products and product lines. Some of these trademark names are used in this book. Where Macmillan Publishing Company was aware of a manufacturer's trademark claim before publication of this book, the trademarked product name has been capitalized herein, either initially or wholly. For example, AT&T is a registered trademark of American Telephone & Telegraph Company; DEC, PDP, VAX, and VMS are registered trademarks of Digital Equipment Corporation; IBM is a registered trademark of International Business Machines Corporation; and UNIX is a registered trademark of Bell Laboratories. All trademarks are the property of their respective owners.

Copyright © 1992 by Macmillan Publishing Company, a division of Macmillan, Inc.

Printed in the United States of America

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without permission in writing from the publisher.

Macmillan Publishing Company
866 Third Avenue, New York, New York 10022

Macmillan Publishing Company is
part of the Maxwell Communication
Group of Companies.

Maxwell Macmillan Canada, Inc.
1200 Eglinton Avenue East
Suite 200
Don Mills, Ontario M3C 3N1

Library of Congress Cataloging in Publication Data

Stallings, William.

Operating systems / William Stallings.

p. cm.

Includes bibliographical references and index.

ISBN 0-02-415481-4

1. Operating systems, (Computers) I. Title.

QA76.76.063S733 1992

005.4'3--dc20

91-10247

CIP

Printing: 2 3 4 5 6 7 8 Year: 2 3 4 5 6 7 8 9 0 1

Acknowledgments are on p. 683, which constitutes an extension of the copyright page.

THE WILLIAM STALLINGS BOOKS ON COMPUTER AND DATA COMMUNICATIONS TECHNOLOGY

DATA AND COMPUTER COMMUNICATIONS, THIRD EDITION

A comprehensive survey that has become the standard in the field, covering four main areas: (1) data communications, including transmission, media, signal encoding, link control, and multiplexing; (2) communication networks, including circuit- and packet-switched, local, packet radio, and satellite; (3) communications architecture, including the OSI model and related protocols; and (4) ISDN and broadband ISDN.

LOCAL NETWORKS: AN INTRODUCTION

An in-depth presentation of the technology of local area networks. Covers topology, transmission media, multiple access, networking, and interface issues. Provides a survey of current standards.

ISDN AND BROADBAND ISDN: SERVICES AND STANDARDS

An in-depth presentation of the technology of integrated services digital networks (ISDN). Covers ISDN services, architecture, signaling systems, and standards of the 1988 and 1990 CCITT standards. Also covers frame relay and broadband ISDN technology.

COMPUTER ORGANIZATION AND ARCHITECTURE SECOND EDITION

A unified view of this broad field. Covers computer architecture, microprogramming, instruction set design, and advanced topics such as RISC and parallel processing.

BUSINESS DATA COMMUNICATIONS

A comprehensive presentation of data communications from a business perspective. Covers data communications and applications technology.

THE WILLIAM STALLINGS ON COMPUTER AND DATA COMMUNICATIONS TECHNOLOGY

OPERATING SYSTEMS: CONCEPTS AND EXAMPLES

A state-of-the art survey of operating system principles. Covers fundamental technology as well as contemporary design issues, such as threads, real-time systems, multiprocessor scheduling, distributed systems, and security.

HANDBOOK OF COMPUTER-COMMUNICATIONS STANDARDS VOLUME 1 THE OPEN SYSTEMS INTERCONNECTION (OSI) MODEL AND OSI-RELATED STANDARDS, SECOND EDITION

A description of the master plan for all computer-communications standards: the OSI model. The book also provides a detailed presentation of OSI-related standards at all 7 layers, including HDLC, X.25, ISO internet, ISO transport, ISO session, ISO presentation, Abstract Syntax ONE (ASN.1), and common application service elements (CASE).

HANDBOOK OF COMPUTER-COMMUNICATIONS STANDARDS VOLUME 2 LOCAL AREA NETWORK STANDARDS, SECOND EDITION

A detailed examination of all current local network standards, including logical link control (LLC, IEEE 802.2), CSMA/CD (IEEE 802.3), token bus (IEEE 802.4), token ring (IEEE 802.5), and fiber distributed data interface (FDDI, ANS X3T9.5).

HANDBOOK OF COMPUTER-COMMUNICATIONS STANDARDS VOLUME 3 THE TCP/IP PROTOCOL SUITE, SECOND EDITION

A description of the protocol standards that are mandated on all DOD computer procurements and are becoming increasingly popular on commercial local network products, including TCP, IP, FTP, SMTP, and TELNET. The network management standards, SNMP and CMOT, are also presented.

OPERATING SYSTEMS

**To Tricia,
the bravest person in the world**

Preface

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 a challenging one for several reasons.

First, there is a tremendous range and variety of computer systems for which operating systems are designed. These include 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.

The object is to provide the reader with a solid understanding of the key mechanisms of modern operating systems, the types of design trade-offs and decisions involved in OS design, and the context within which the operating system functions (hardware, other system programs, application programs, interactive users).

The Process

The concept of the process is central to the study of operating systems. While all books on the subject deal with this topic, no other leading text devotes a major section to introducing and explaining the basic principles of processes.

In this book, Chapter 3 is devoted to this task. The result is a solid foundation for examining the many issues addressed in later chapters.

Recent Developments in Operating System Design

In addition to providing coverage of the fundamentals of operating systems, this book examines the most important recent developments in OS design. Among the topics covered:

- *Threads*: The concept of process is more complex and subtle than usually presented and in fact embodies two separate and potentially independent concepts: one relating to resource ownership and one relating to execution. This distinction has led to the development in some operating systems of a construct known as the thread.
- *Real time systems*: In recent years, real-time computing has come to be viewed as an important emerging discipline in computer science and engineering. The operating system, and in particular the scheduler, is perhaps the most important component of a real-time system.
- *Multiprocessor scheduling*: Traditionally, there has been little difference between the scheduling approaches taken on multiprogrammed uniprocessors and multiprocessor systems. With the increasing interest in the use of threads and in parallel programming, multiprocessor scheduling has become an area of intense study and development.
- *Distributed systems*: With the increasing availability of inexpensive yet powerful personal computers and minicomputers, there has been an increasing trend toward distributed data processing (DDP). With DDP, processors, data, and other aspects of a data processing system may be dispersed within an organization. Many of the OS design issues covered in this book take on added complexity in a distributed environment.
- *Process migration*: Process migration is the ability to move an active process from one machine to another; it has become an increasingly "hot" topic in distributed operating systems. Interest in this concept grew out of research into ways to do load balancing across multiple networked systems, although the application of the concept now extends beyond that one area. Until recently, a number of observers felt that process migration was impractical. Such assessments have proved too pessimistic. New implementations, including those in commercial products, have fueled a continuing interest and new developments in this area.
- *Security*: Security has long been a concern in the design of operating systems. However, the design approaches to security have continued to evolve as the threats to security have evolved. Examples of areas of threat that present new and more complex difficulties include viruses and attacks on distributed operating systems. An example of a new approach to coping with these threats is the concept of a trusted system.

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. In order to illustrate the concepts and to tie them to real-world design choices that must be made, three operating systems have been chosen as running examples:

- *OS/2*: A single-user, multitasking operating system for the IBM PS/2 and compatible systems. It is one of the few truly new operating systems that has been essentially designed from scratch (although it includes an MS-DOS compatibility mode). As such, it is in a position to incorporate in a clean fashion the latest developments in operating system technology.
- *UNIX*: A multiuser operating system, originally intended for minicomputers, but implemented on a wide range of machines from powerful microcomputers to supercomputers.
- *MVS*: The top-of-the-line operating system for IBM mainframes, and perhaps the most complex operating system ever developed. It provides both batch and timesharing capabilities.

These three systems were chosen because of their relevance and representativeness. Most personal computer operating systems on new machines are single-user multitasking systems, with OS/2 becoming the market leader. UNIX has become the dominant operating system on a wide variety of workstations and multiuser systems. MVS is the most widely used mainframe operating system. Thus, most readers will be exposed to one or more of these operating systems by the time they use this book or within a few years of doing so.

As with the technique used in the author's *Computer Organization and Architecture*, 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 virtual memory, the page replacement algorithm of each example system is 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.

Intended Audience

The book is intended for both an academic and a professional audience. As a textbook, it is intended as a one-semester undergraduate course in operating systems for computer science, computer engineering, and electrical engineering majors. It covers the topics in Subject Area 5 of the IEEE Model Program in Computer Science and Engineering and also covers the OS-related topics in CS6 and CS10 of the ACM Recommendations for the Undergraduate Program in Computer Science. It also covers all of the topics in the operating systems topic area of the ACM/IEEE Joint Computing Curricula 1991.

The book also serves as a basic reference volume and is suitable for self-study.

Plan of the Text

The organization of the chapters is as follows:

1. *Computer system overview*: Provides an overview of computer architecture and organization, with emphasis on topics that relate to operating system design.
 2. *Operating system overview*: Provides an overview of the remainder of the book.
 3. *Process description and control*: Presents the concept of a process and examines the data structures used by the operating system to control processes. The related concepts of threads and sessions are also discussed.
 4. *Concurrency*: Examines the key aspects of concurrency on a single system, with emphasis on issues of mutual exclusion and deadlock.
 5. *Memory management*: Provides a comprehensive survey of techniques for memory management.
 6. *Scheduling*: Provides a comparative discussion of various approaches to process scheduling. The unique requirements of multiprocessors and real-time systems are examined.
 7. *I/O management and disk scheduling*: Examines the issues involved in OS control of the I/O function. Special attention is devoted to disk I/O, which is the key to system performance.
 8. *File management*: Provides an overview of file management, with an emphasis on those aspects that are usually implemented as part of the operating system or are closely tied to the operating system.
 9. *Networking and distributed processing*: Examines the major trends in this area, including the OSI model, the use of network servers, and distributed operating systems.
 10. *Security*: Provides a survey of threats and mechanisms for providing computer and network security.
- A. *Queuing analysis*: An appendix provides a practical guide to the use of queuing analysis to model performance.

In addition, the book includes an extensive glossary, a list of frequently-used acronyms, and a bibliography. Each chapter includes problems and suggestions for further reading.

Acknowledgments

Many people have reviewed all or a portion of the manuscript for this book. I would like to especially thank the following: Umakishore Ramachandran, Georgia Institute of Technology; Margaret Reek, Rochester Institute of Technology; Anand Tripathi, University of Minnesota; Sol Shatz, University of Illinois–Chicago; Philip Krueger, The Ohio State University; Samuel Gulden, Lehigh University; E. K. Parti, U.S. Naval Academy; Armin Roeseler, AT&T Bell Labs; Charles Shub, University of Colorado–Colorado Springs; Eric Cooper,

Carnegie Mellon University; Evelyn Obaid, San Jose State University; Evan Ivie, Brigham Young University; Mark Smotherman, Clemson University; Gary Harkin, Montana State University; Rance Cleaveland, North Carolina State University; James Silver, Indiana University–Purdue University at Fort Wayne.

I would like to express my deep appreciation to Margaret Reek and Sol Shatz who contributed some of the end-of-chapter homework problems.

W.S.

-
- 1.1 BASIC ELEMENTS
 - 1.2 PROCESSOR REGISTERS
 - User-Visible Registers
 - Control and Status Registers
 - 1.3 INSTRUCTION EXECUTION
 - Instruction Fetch and Execute
 - I/O Function
 - 1.4 INTERRUPTS
 - Interrupts and the Instruction Cycle
 - Interrupt Processing
 - Multiple Interrupts
 - Multiprogramming
 - 1.5 INTERCONNECTION STRUCTURES
 - 1.6 THE MEMORY HIERARCHY
 - 1.7 CACHE MEMORY
 - Motivation
 - Cache Principles
 - Cache Design
 - 1.8 INPUT/OUTPUT ORGANIZATION
 - External Devices
 - I/O Module Function
 - I/O Module Structure
 - 1.9 I/O COMMUNICATION TECHNIQUES
 - Interrupt-Driven I/O
 - Direct Memory Access
 - 1.10 RECOMMENDED READING
 - 1.11 PROBLEMS
 - APPENDIX 1A PERFORMANCE CHARACTERISTICS OF TWO-LEVEL MEMORIES
 - 1A.1 Locality
 - 1A.2 Operation of Two-Level Memory
 - 1A.3 Performance
 - APPENDIX 1B PROCEDURE CONTROL
 - 1B.1 Stack Implementation
 - 1B.2 Procedure Calls and Returns
 - 1B.3 Reentrant Procedures

ACKNOWLEDGMENTS

Table 1.2: © 1990 Morgan Kaufmann Publishers. Reprinted with permission from David Patterson and John Hennessy, *Computer Architecture: A Quantitative Approach*. Figure 2.3: From A. Silberschatz, J. Peterson, and P. Galvin, *Operating System Concepts*. © 1991, by Addison-Wesley Publishing Company. Reprinted with permission of the publisher. Figures 2.7 and 2.8: Reprinted with permission of Macmillan Publishing Company from *Operating Systems: Design and Implementation* by Raymond W. Turner. Copyright © 1986 by Macmillan Publishing Company, a Division of Macmillan, Inc. Figure 2.14: Reprinted with permission from *Version 4.4 VAX/VMS Internals and Data Structures* by Lawrence J. Kenah, Ruth E. Goldenberg, and Simon F. Bate. Copyright © 1988 Digital Press/Digital Equipment Corporation, 12 Crosby Drive, Bedford, MA. 01730. Figure 2.18: From Maurice J. Bach, *The Design of the UNIX® Operating System*, © 1986, pp. 20. Reprinted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Figure 2.19: Reprinted, by permission, from *Architecture and Implementation of Large Scale IBM Computer Systems*, Second Edition, 1981, QED Information Sciences, Inc., Wellesley, MA. Table 3.2: From James R. Pinkert and Larry L. Wear, *Operating Systems: Concepts, Policies, and Mechanisms*, © 1989, p. 69. Reprinted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Table 3.11: From S. Krakowiak and D. Beeson, *Principles of Operating Systems*. Copyright © Massachusetts Institute of Technology 1988. Published by the MIT Press. Figure 3.14: From Maurice J. Bach, *The Design of the UNIX® Operating System*, © 1986, p. 148. Reprinted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Table 4.4: From S. Isloor and T. Marsland, "The Deadlock Problem: An Overview," *Computer*, September 1980. © 1980 IEEE. Tables 4.7, 4.8, 4.9, and 4.16: From M. Ben-Ari, *Principles of Concurrent Programming*, © 1982, pp. 29, 32, 39, 52. Reprinted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Figure 4.18: From J. Hayes, *Computer Architecture and Organization*, Second Edition. © 1988. Reprinted by permission of McGraw-Hill, Inc. Figure 4.24: From Lubomir Bic and Alan C. Shaw, *The Logical Design of Operating Systems*, 2e., © 1988, p. 78. Reprinted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Figure 4.27: From M. Milenkovic, *Operating Systems: Concepts and Design*. © 1987. Reprinted by permission of McGraw-Hill, Inc. Figure 5.24: From J. Hayes, *Computer Architecture and Organization*, Second Edition. © 1988. Reprinted by permission of McGraw-Hill, Inc. Figure 5.27: From Maurice J. Bach, *The Design of the UNIX® Operating System*, © 1986, p. 287. Reprinted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Figure 5.28: From M. Maekawa, A. Oldehoeft, and R. Oldehoeft, *Operating Systems: Advanced Concepts*. Copyright © 1987. Reprinted by permission of Benjamin/Cummings Publishing Co. Table 6.7: From James Martin, *Principles of Data Communication*. © 1988, p. 290. Reprinted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Figures 6.15 and 6.14: From Raphael A. Finkel, *An Operating Systems Vade Mecum*, 2e, © 1988, pp. 35 and 36. Adapted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Figure 6.16: From Maurice J. Bach, *The Design of the UNIX® Operating System*, © 1986, p. 257. Reprinted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Figure 6.17: From Charles H. Sauer and K. Mani Chandy, *Computer Systems Performance Modeling*, © 1981, p. 25. Reprinted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Figure 6.23: From Maurice J. Bach, *The Design of the UNIX® Operating System*, © 1986, p. 253. Reprinted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Figure 6.24: From E. Iacobucci, *OS/2 Programmer's Guide*. © 1988. Reprinted by permission of McGraw-Hill, Inc. Tables 7.4 and 8.1 and Figure 8.6: From G. Wiederhold, *File Organization for Database Design*. © 1987, pp. 75, 225, and 420. Reprinted by permission of McGraw-Hill, Inc. Table 7.6: From J. Krantz, A. Mizell, and R. Williams, *OS/2: Features, Functions, and Applications*. © 1988, p. 32. John Wiley & Sons, Inc. Figure 8.1: From Daniel Grosshans, *File Systems: Design & Implementation*, © 1986, p. 395. Adapted by permission of Prentice Hall, Englewood Cliffs, New Jersey. Figure 8.6: From G. Wiederhold, *File Organization for Database Design*. © 1987. Reprinted by permission of McGraw-Hill, Inc. Figure 9.19: From S. Krakowiak and D. Beeson, *Principles of Operating Systems*. Copyright © Massachusetts Institute of Technology 1988. Published by the MIT Press.

Contents

PREFACE vii

CHAPTER 1 COMPUTER SYSTEM OVERVIEW 1

1-1	Basic Elements	1
1-2	Processor Registers	3
1-3	Instruction Execution	5
1-4	Interrupts	10
1-5	Interconnection Structures	22
1-6	The Memory Hierarchy	25
1-7	Cache Memory	29
1-8	Input/Output Organization	34
1-9	I/O Communication Techniques	40
1-10	Recommended Reading	44
1-11	Problems	45
	Appendix 1A Performance Characteristics of Two-Level Memories	47
	Appendix 1B Procedure Control	53

CHAPTER 2 OPERATING SYSTEM OVERVIEW 59

2-1	Operating System Objectives and Functions	59
2-2	The Evolution of Operating Systems	64
2-3	Major Achievements	75
2-4	Example Systems	91
2-5	Outline of the Remainder of the Book	104
2-6	Recommended Reading	108
2-7	Problems	109

CHAPTER 3 PROCESS DESCRIPTION AND CONTROL 113

3-1	Process States	114
3-2	Process Description	133

3-3	Process Control	143
3-4	Processes and Threads	152
3-5	Examples of Process Description and Control	158
3-6	Summary	172
3-7	Recommended Reading	172
3-8	Problems	173

CHAPTER 4 CONURRENCY 179

4-1	Principles of Concurrency	180
4-2	Mutual Exclusion—Software Approaches	193
4-3	Mutual Exclusion—Hardware Support	200
4-4	Mutual Exclusion—Operating System Support	204
4-5	Deadlock	221
4-6	Classic Problems in Concurrency	236
4-7	Example Systems	243
4-8	Summary	250
4-9	Recommended Reading	253
4-10	Problems	254

CHAPTER 5 MEMORY MANAGEMENT 265

5-1	Memory Management Requirements	265
5-2	Loading Programs into Main Memory	269
5-3	Virtual Memory—Hardware and Control Structures	286
5-4	Virtual Memory—Operating System Software	304
5-5	Examples of Memory Management	323
5-6	Summary	335
5-7	Recommended Reading	336
5-8	Problems	337
	Appendix 5A Loading and Linking	339
	Appendix 5B Hash Tables	346

CHAPTER 6 SCHEDULING 351

6-1	Types of Scheduling	352
6-2	Scheduling Algorithms	356
6-3	Multiprocessor Scheduling	380
6-4	Real-Time Scheduling	393
6-5	Example Systems	399
6-6	Summary	404
6-7	Recommended Reading	406
6-8	Problems	406
	Appendix 6A Response Time	409

CHAPTER 7 I/O MANAGEMENT AND DISK SCHEDULING 413

7-1	I/O Devices	413
7-2	Organization of the I/O Function	415
7-3	Operating System Design Issues	421
7-4	I/O Buffering	424
7-5	Disk I/O	428
7-6	Example Systems	439
7-7	Summary	449
7-8	Recommended Reading	449
7-9	Problems	450
	Appendix 7A Secondary Storage Devices	452

CHAPTER 8 FILE MANAGEMENT 465

8-1	Overview	466
8-2	File Organization and Access	471
8-3	File Directories	478
8-4	File Sharing	483
8-5	Record Blocking	484
8-6	Secondary Storage Management	485
8-7	Example System—UNIX System V	496
8-8	Summary	498
8-9	Recommended Reading	500
8-10	Problems	501

CHAPTER 9 NETWORKING AND DISTRIBUTED PROCESSING 504

9-1	Communications Architecture	506
9-2	Servers	523
9-3	Process Migration	530
9-4	Distributed Process Communication	536
9-5	Distributed Process Management—Mutual Exclusion	547
9-6	Distributed Process Management—Deadlock	559
9-7	Summary	566
9-8	Recommended Reading	567
9-9	Problems	568

CHAPTER 10 SECURITY 571

10-1	Security Threats	573
10-2	Computer Security	580
10-3	Trusted Systems	587
10-4	Network Security	589