

Contents

About the Author

xxvi

PART 1 OVERVIEW

CHAPTER 1 INTRODUCTION

Outline	1
1.1 Opening Remarks	2
1.2 What is an Operating System?	3
1.3 Early History: The 1940s and The 1950s	3
1.4 The 1960s	4
1.5 The Emergence of a New Field: Software Engineering	4
1.6 The 1970s	6
1.7 The 1980s	6
1.8 The 1990s and Beyond	6
1.9 Distributed Computing	7
1.10 The Key Architectural Trend: Parallel Computation	8
1.11 Input-Output Trends	9
1.12 Open Systems	10
1.13 UNIX	11
1.14 Ethical Issues	12
1.15 Application Bases	13
1.16 The Key Operating Systems for the 1990s	13
Appendix: Open System Organizations	14
	20

CHAPTER 2 HARDWARE, SOFTWARE, FIRMWARE

Outline

- 2.1 Introduction
- 2.2 Hardware
- 2.3 Software
- 2.4 Firmware
- Summary

PART 2 PROCESS MANAGEMENT

CHAPTER 3 PROCESS CONCEPTS

Outline

- 3.1 Introduction
- 3.2 Definitions of "Process"
- 3.3 Process States
- 3.4 Process State Transitions
- 3.5 The Process Control Block
- 3.6 Operations on Processes
- 3.7 Suspend and Resume
- 3.8 Interrupt Processing
- 3.9 The Nucleus of the Operating System
- Summary

CHAPTER 4 ASYNCHRONOUS CONCURRENT PROCESSES

Outline

- 4.1 Introduction
- 4.2 Parallel Processing
- 4.3 A Control Structure for Indicating Parallelism: Parbegin/Parend
- 4.4 Mutual Exclusion
- 4.5 Critical Sections
- 4.6 Mutual Exclusion Primitives
- 4.7 Implementing Mutual Exclusion Primitives
- 4.8 Dekker's Algorithm
- 4.9 Peterson's Algorithm
- 4.10 N-Process Mutual Exclusion
- 4.11 A Hardware Solution to Mutual Exclusion: The Testandset Instruction
- 4.12 Semaphores
- 4.13 Process Synchronization with Semaphores
- 4.14 The Producer-consumer Relationship

4.15 Counting Semaphores	92
4.16 Implementing Semaphores, P and V	92
4.17 Eventcounts and Sequenceers	94
Summary	96

CHAPTER 5 CONCURRENT PROGRAMMING 111

Outline	112
5.1 Introduction	113
5.2 Critical Regions and Conditional Critical Regions	114
5.3 Monitors	115
5.4 Path Expressions	121
5.5 Message Passing	121
5.6 Ada: A Key Concurrent Programming Language for the 1990s	124
Summary	135

CHAPTER 6 DEADLOCK AND INDEFINITE POSTPONEMENT 153

Outline	154
6.1 Introduction	155
6.2 Examples of Deadlock	155
6.3 A Related Problem: Indefinite Postponement	158
6.4 Resource Concepts	158
6.5 Four Necessary Conditions for Deadlock	159
6.6 Major Areas of Deadlock Research	160
6.7 Deadlock Prevention	160
6.8 Deadlock Avoidance and the Banker's Algorithm	164
6.9 Deadlock Detection	168
6.10 Deadlock Recovery	170
6.11 Deadlock Considerations in Future Systems	172
Summary	173

PART 3 STORAGE MANAGEMENT

CHAPTER 7 REAL STORAGE 285

Outline	186
7.1 Introduction	187
7.2 Storage Organization	187
7.3 Storage Management	187
7.4 Storage Hierarchy	188
7.5 Storage Management Strategies	188

7.6	Contiguous vs. Noncontiguous Storage Allocation	1
7.7	Single User Contiguous Storage Allocation	1
7.8	Fixed Partition Multiprogramming	1
7.9	Variable Partition Multiprogramming	1
7.10	Multiprogramming with Storage Swapping	1
	Summary	2

CHAPTER 8 VIRTUAL STORAGE ORGANIZATION

	Outline	2
8.1	Introduction	2
8.2	Evolution of Storage Organizations	2
8.3	Virtual Storage: Basic Concepts	2
8.4	Multilevel Storage Organization	2
8.5	Block Mapping	2
8.6	Paging: Basic Concepts	2
8.7	Segmentation	2
8.8	Paging/Segmentation Systems	2
	Summary	2

CHAPTER 9 VIRTUAL STORAGE MANAGEMENT

	Outline	2
9.1	Introduction	2
9.2	Virtual Storage Management Strategies	2
9.3	Page Replacement Strategies	2
9.4	Locality	2
9.5	Working Sets	2
9.6	Page Fault Frequency Page Replacement	2
9.7	Demand Paging	2
9.8	Antic	2
9.9	Page Release	2
9.10	Page Size	2
9.11	Program Behavior under Paging	2
	Summary	2

PART 4 PROCESSOR MANAGEMENT

CHAPTER 10 JOB AND PROCESSOR SCHEDULING

	Outline	2
10.1	Introduction	2
10.2	Scheduling Levels	2

10.3 Scheduling Objectives	287
10.4 Scheduling Criteria	289
10.5 Preemptive vs. Nonpreemptive Scheduling	290
10.6 The Interval Timer or Interrupting Clock	291
10.7 Priorities	291
10.8 Deadline Scheduling	292
10.9 First-In-First-Out (FIFO) Scheduling	292
10.10 Round Robin (RR) Scheduling	293
10.11 Quantum Size	294
10.12 Shortest-Job-First (SJF) Scheduling	295
10.13 Shortest-Remaining-Time (SRT) Scheduling	295
10.14 Highest-Response-Ratio-Next (HRRN) Scheduling	296
10.15 Multilevel Feedback Queues	296
10.16 Fair Share Scheduling	299
10.17 Sun UNIX Process Scheduling	301
10.18 VAX/VMS Process Scheduling	301
Summary	301

CHAPTER 11

DISTRIBUTED COMPUTING: THE PARALLEL COMPUTATION VIEW

313

Outline	314
11.1 Introduction	316
11.2 The Fifth Generation Project and the Strategic Computing Initiative	316
11.3 Supercomputer Projects	316
11.4 Classification of Sequential and Parallel Architectures	317
11.5 Pipelining	318
11.6 Vector Processing	318
11.7 Array Processors	318
11.8 Data Flow Computers	319
11.9 Multiprocessors	319
11.10 Fault Tolerance	320
11.11 Computational Complexity Issues	321
11.12 Detecting Parallelism	322
11.13 The "Never Wait" Rule	324
11.14 Processor Interconnection Schemes	325
11.15 Loosely-Coupled Vs. Tightly-Coupled Systems	329
11.16 Fetch-and-Add	330
11.17 Multiprocessor Operating System Organizations	332
11.18 Case Study: C.mmp and Cm*	334
11.19 Case Study: The Butterfly	334
11.20 Case Study: The Connection Machine	340
Summary	346
Appendix: Organizations Researching and Developing Highly-parallel Computers	357

PART 5 AUXILIARY STORAGE MANAGEMENT

CHAPTER 12 DISK PERFORMANCE OPTIMIZATION

Outline

- 12.1 Introduction
- 12.2 Operation of Moving-Head Disk Storage
- 12.3 Why Disk Scheduling Is Necessary
- 12.4 Desirable Characteristics of Disk Scheduling Policies
- 12.5 Seek Optimization
- 12.6 Rotational Optimization
- 12.7 Systems Considerations
- 12.8 Disk Caching
- 12.9 Other Performance Enhancement Techniques
- 12.10 RAM Disks
- 12.11 Optical Disks

Summary

CHAPTER 13 FILE AND DATABASE SYSTEMS

Outline

- 13.1 Introduction
- 13.2 The File System
- 13.3 File System Functions
- 13.4 The Data Hierarchy
- 13.5 Blocking and Buffering
- 13.6 File Organization
- 13.7 Queued and Basic Access Methods
- 13.8 Allocating and Freeing Space
- 13.9 File Descriptor
- 13.10 Access Control Matrix
- 13.11 Access Control by User Classes
- 13.12 Backup and Recovery
- 13.13 File Servers
- 13.14 Distributed File Systems
- 13.15 CD-ROMs, WORMs, and Magneto-Laser Disks
- 13.16 Database Systems
- 13.17 Database Models

Summary

PART 6 PERFORMANCE

CHAPTER 14 PERFORMANCE COPROCESSORS, RISC, AND DATA FLOW 419

Outline	420
14.1 Introduction	421
14.2 Important Trends Affecting Performance Issues	421
14.3 Why Performance Monitoring and Evaluation Are Needed	421
14.4 Performance Measures	422
14.5 Performance Evaluation Techniques	424
14.6 Bottlenecks and Saturation	429
14.7 Feedback Loops	429
14.8 Coprocessors	430
14.9 Reduced Instruction Set Computing (RISC)	431
14.10 Case Study: Data Flow, RISC, and Pipelining	434
Summary	438
Appendix: Companies with Commitments to RISC-based Architectures	452

CHAPTER 15 ANALYTIC MODELING 453

Outline	454
15.1 Introduction	455
15.2 Queuing Theory	455
15.3 Markov Processes	471
Summary	479

PART 7 NETWORKS AND SECURITY

CHAPTER 16 DISTRIBUTED COMPUTING: THE OPEN SYSTEMS INTERCONNECTION VIEW 489

Outline	490
16.1 Introduction	492
16.2 Network Architecture: OSI and Other Reference Models	496
16.3 The Application Layer (Layer 7)	502

- 16.4 The Presentation Layer (Layer 6)
- 16.5 The Session Layer (Layer 5)
- 16.6 The Transport Layer (Layer 4)
- 16.7 The Network Layer (Layer 3)
- 16.8 The Data Link Layer (Layer 2)
- 16.9 The Physical Layer (Layer 1)
- 16.10 Transmission Media
- 16.11 OSI Network Management
- 16.12 OSI Security
- 16.13 Integrated Digital Services Digital Network (ISDN)
- 16.14 MAP, TOP, and GOSIP
- 16.15 OSI Testing
- 16.16 Transmission Control Protocol/Internet Protocol (TCP/IP)
- 16.17 The Future of OSI
- Appendix: OSI-oriented Standards

CHAPTER 17 OPERATING SYSTEMS SECURITY

Outline

- 17.1 Introduction
- 17.2 Security Requirements
- 17.3 A Total Approach to Security
- 17.4 External Security
- 17.5 Operational Security
- 17.6 Surveillance
- 17.7 Threat Monitoring
- 17.8 Amplification
- 17.9 Password Protection
- 17.10 Auditing
- 17.11 Access Controls
- 17.12 Security Kernels
- 17.13 Hardware Security
- 17.14 Fault-Tolerant Systems
- 17.15 Capabilities and Object-Oriented Systems
- 17.16 Case Study: The Object-Oriented Architecture of the IBM System/38
- 17.17 Cryptography
- 17.18 Operating System Penetration
- 17.19 Case Study: Penetrating an Operating System
- 17.20 UNIX Operating System Security
- 17.21 Worms and Viruses
- Summary
- Appendix: Information Sources

PART 8 CASE STUDIES

CHAPTER 18 CASE STUDY: UNIX SYSTEMS 567

Outline	568
18.1 Introduction	570
18.2 History	571
18.3 The Shell	573
18.4 The Kernel	575
18.5 The File System	576
18.6 Process Management	582
18.7 Memory Management	587
18.8 The Input/Output System	590
18.9 The SunOS Operating System	593
8.10 Distributed UNIX Systems	596
8.11 UNIX Systems Standardization and Open Systems	602
8.12 The Future of UNIX Systems	605
Appendix A: Information Sources	618
Appendix B: A Tour of the AT&T System V Interface Definition (SVID)	619

CHAPTER 19 CASE STUDY: MS-DOS, DOS 629

Outline	630
19.1 Introduction	632
19.2 Early History of MS-DOS	633
19.3 MS-DOS Version 1.0	634
19.4 MS-DOS Version 2.0	635
19.5 MS-DOS Version 3.0	637
19.6 The IBM Personal System/2	638
19.7 The Micro Channel	639
19.8 MS-DOS 3.3	639
19.9 Operating System/2 Standard Edition	640
9.10 Operating System/2 Extended Edition	641
9.11 Systems Application Architecture	641
9.12 The User's View of MS-DOS	642
9.13 The Systems's View of MS-DOS	646
9.14 The Programmer's View of MS-DOS System Calls	655
9.15 DOS 4.0	660
9.16 The Future of MS-DOS and OS/2	661
Appendix: Information Sources	667

CHAPTER 20

CASE STUDY: MVS

Outline

- 20.1 History of MVS
- 20.2 MVS Design Objectives and Capabilities
- 20.3 System/370 Hardware
- 20.4 Storage Management
- 20.5 MVS/XA Address Spaces
- 20.6 Multiprocessing
- 20.7 Managing System Resources
- 20.8 Supervising the Execution of Work
- 20.9 Input/Output
- 20.10 Job Entry Subsystem
- 20.11 Monitoring System Activity
- 20.12 Enterprise Systems Architecture and MVS/ESA
- 20.13 The Future of MVS

CHAPTER 21

CASE STUDY: VM: A VIRTUAL MACHINE OPERATING SYSTEM

Outline

- 21.1 Introduction
 - 21.2 History
 - 21.3 The Control Program (CP)
 - 21.4 The Conversational Monitor System (CMS)
 - 21.5 Remote Spooling and Communications Subsystem
 - 21.6 Strengths of VM
 - 21.7 VM/370 Evolution
 - 21.8 Performance Considerations
 - 21.9 Reliability, Availability, and Serviceability
 - 21.10 The Versions of VM
 - 21.11 AIX/370
 - 21.12 UNIX vs. VM
 - 21.13 VM/SP Extensions
 - 21.14 VM: IBM's Large-Scale Operating System for the 1990s?
- Appendix: Other IBM Publications

CHAPTER 22

CASE STUDY: INSIDE THE APPLE MACINTOSH

Outline

- 22.1 Introduction
- 22.2 History
- 22.3 Hardware
- 22.4 Software
- 22.5 Resource Manager

22.6	QuickDraw	744
22.7	Font Manager	744
22.8	Toolbox Event Manager	745
22.9	Window Manager	747
22.10	Control Manager	747
22.11	Menu Manager	748
22.12	TextEdit	749
22.13	Dialog Manager	749
22.14	Desk Manager	750
22.15	Package Manager	750
22.16	Memory Manager	751
22.17	Segment Loader	753
22.18	Operating System Event Manager	754
22.19	File Manager	754
22.20	Printing Manager	759
22.21	Device Manager	761
22.22	System Error Handler	770
22.23	Operating System Utilities	770
22.24	The Evolving Macintosh Family	771
22.25	Macintosh Plus Operating System Enhancements	773
22.26	Macintosh SE and Macintosh II Operating System Enhancements	773
22.27	The Future of the Macintosh	774
	Appendix: Information Sources	779

CHAPTER 23

CASE STUDY: OS/2

781

	Outline	782
23.1	Introduction	784
23.2	History	784
23.3	Design Overview	785
23.4	Multitasking	788
23.5	Dynamic Linking	793
23.6	The File System	794
23.7	Memory Management	795
23.8	Interprocess Communication	799
23.9	The Presentation Manager and VIO	803
23.10	Device Monitors	804
23.11	Data Integrity	807
23.12	Timer Services	807
23.13	Device Drivers	808
23.14	The Compatibility Box	814
23.15	System Configuration	815
23.16	OS/2 Software Developer's Kit	816
23.17	Handling Errors and Exceptions	816
23.18	The Future	816
	Appendix: Summary of OS/2 Commands	827

Index

833