

Contents

<i>Preface</i>	<i>xiii</i>
1 Introduction	1
1.1 Abstract Views of an Operating System	1
1.2 Goals of an OS	5
1.3 Operation of an OS	7
1.4 Preview of the Book	16
<i>Exercises</i>	23
<i>Bibliography</i>	25

Part I Fundamental Concepts

2 Overview of Operating Systems	31
2.1 OS and the Computer System	31
2.2 Efficiency, System Performance and User Convenience	46
2.3 Classes of Operating Systems	49
2.4 Batch Processing Systems	52
2.5 Multiprogramming Systems	56
2.6 Time Sharing Systems	65
2.7 Real Time Operating Systems	70
2.8 Distributed Operating Systems	73
2.9 Modern Operating Systems	75
<i>Exercises</i>	76
<i>Bibliography</i>	80
3 Processes and Threads	84
3.1 Processes and Programs	84
3.2 Programmer View of Processes	87
3.3 OS View of Processes	92
3.4 Threads	106
3.5 Case Studies of Processes and Threads	115
3.6 Interacting Processes—An Advanced Programmer View of Processes	130
<i>Exercises</i>	139
<i>Bibliography</i>	141
4 Scheduling	143
4.1 Preliminaries	143
4.2 Non-preemptive Scheduling Policies	148
4.3 Preemptive Scheduling Policies	151
4.4 Scheduling in Practice	156
4.5 Real Time Scheduling	168
4.6 Scheduling in Unix	175

viii Contents

4.7	Scheduling in Linux	178
4.8	Scheduling in Windows	181
4.9	Performance Analysis of Scheduling Policies	181
	<i>Exercises</i>	186
	<i>Bibliography</i>	188
5	Memory Management	191
5.1	Managing the Memory Hierarchy	191
5.2	Static and Dynamic Memory Allocation	193
5.3	Memory Allocation to a Process	196
5.4	Reuse of Memory	202
5.5	Contiguous Memory Allocation	211
5.6	Noncontiguous Memory Allocation	213
5.7	Paging	217
5.8	Segmentation	219
5.9	Segmentation with Paging	220
5.10	Kernel Memory Allocation	221
5.11	A Review of Relocation, Linking and Program Forms	226
	<i>Exercises</i>	233
	<i>Bibliography</i>	235
6	Virtual Memory	237
6.1	Virtual Memory Basics	237
6.2	Demand Paging	240
6.3	Page Replacement Policies	266
6.4	Memory Allocation to a Process	275
6.5	Shared Pages	279
6.6	Memory Mapped Files	282
6.7	Unix Virtual Memory	284
6.8	Linux Virtual Memory	289
6.9	Virtual Memory in Windows	290
6.10	Virtual Memory Using Segmentation	291
	<i>Exercises</i>	295
	<i>Bibliography</i>	299
7	File Systems	302
7.1	File System and IOCS	302
7.2	Files and File Operations	305
7.3	Fundamental File Organizations	307
7.4	Directory Structures	311
7.5	File Protection	318
7.6	Interface between File System and IOCS	319
7.7	Allocation of Disk Space	322
7.8	Implementing File Access	326
7.9	File Sharing Semantics	331
7.10	File System Reliability	333
7.11	Virtual File System	341
7.12	Unix File System	343

7.13	Linux File System	348
7.14	Windows File System	348
7.15	Performance of File Systems	350
	<i>Exercises</i>	352
	<i>Bibliography</i>	355
8	Security and Protection	357
8.1	Overview of Security and Protection	357
8.2	Goals of Security and Protection	360
8.3	Security Attacks	362
8.4	Formal and Practical Aspects of Security	364
8.5	Encryption	366
8.6	Authentication and Password Security	373
8.7	Access Descriptors and the Access Control Matrix	375
8.8	Protection Structures	377
8.9	Capabilities	382
8.10	Unix Security	390
8.11	Linux Security	391
8.12	Windows Security	392
	<i>Exercises</i>	394
	<i>Bibliography</i>	395

Part II Advanced Topics

9	Process Synchronization	403
9.1	Data Access Synchronization and Control Synchronization	403
9.2	Critical Sections	404
9.3	Race Conditions in Control Synchronization	409
9.4	Implementing Critical Sections and Indivisible Operations	411
9.5	Classic Process Synchronization Problems	414
9.6	Structure of Concurrent Systems	421
9.7	Algorithmic Approach Implementing Critical Sections	423
9.8	Semaphores	431
9.9	Conditional Critical Regions	443
9.10	Monitors	447
9.11	Process Synchronization in Unix	456
9.12	Process Synchronization in Linux	457
9.13	Process Synchronization in Windows	458
	<i>Exercises</i>	459
	<i>Bibliography</i>	463
10	Message Passing	466
10.1	Overview of Message Passing	466
10.2	Implementing Message Passing	470
10.3	Mailboxes	474
10.4	Message Passing in Unix	476

x Contents

10.5	Message Passing in Windows	480
	<i>Exercises</i>	481
	<i>Bibliography</i>	481
11	Deadlocks	483
11.1	Definition of deadlock	483
11.2	Deadlocks in Resource Allocation	484
11.3	Handling Deadlocks	490
11.4	Deadlock Detection and Resolution	491
11.5	Deadlock Prevention	496
11.6	Deadlock Avoidance	500
11.7	Formal Characterization of Resource Deadlocks	506
11.8	Deadlock Handling in Practice	513
	<i>Exercises</i>	515
	<i>Bibliography</i>	518
12	Implementation of File Operations	520
12.1	Layers of the Input Output Control System	520
12.2	Overview of I/O Organization	523
12.3	I/O Devices	525
12.4	Device Level I/O	537
12.5	Disk Scheduling	551
12.6	Buffering of Records	554
12.7	Blocking of Records	559
12.8	Access Methods	562
12.9	Unified Disk Cache	564
12.10	File Processing in Unix	565
12.11	File Processing in Linux	568
12.12	File Processing in Windows	570
	<i>Exercises</i>	571
	<i>Bibliography</i>	573
13	Synchronization and Scheduling in Multiprocessor Operating Systems	576
13.1	Architecture of Multiprocessor Systems	576
13.2	Multiprocessor Operating Systems	583
13.3	Kernel Structure	584
13.4	Process Synchronization	587
13.5	Process Scheduling	591
13.6	Case Studies	593
	<i>Exercises</i>	596
	<i>Bibliography</i>	597
14	Structure of Operating Systems	599
14.1	Operation of an OS	600
14.2	Structure of an Operating System	601
14.3	Operating Systems with Monolithic Structure	603
14.4	Layered Design of Operating Systems	604
14.5	Virtual Machine Operating Systems	607

14.6	Kernel Based Operating Systems	610
14.7	Microkernel Based Operating Systems	612
14.8	Configuring and Installing the Kernel	614
14.9	Architecture of Unix	615
14.10	The Kernel of Linux	617
14.11	Architecture of Windows	618
	<i>Exercises</i>	619
	<i>Bibliography</i>	619

Part III

Distributed Operating Systems

15	Distributed Operating Systems	625
15.1	Features of Distributed Systems	625
15.2	Nodes of a Distributed System	627
15.3	Network Operating Systems	627
15.4	Distributed Operating Systems	628
15.5	Reliable Interprocess Communication	631
15.6	Distributed Computation Paradigms	637
15.7	Networking	644
15.8	Model of a Distributed System	657
15.9	Design Issues in Distributed Operating Systems	659
	<i>Exercises</i>	661
	<i>Bibliography</i>	662
16	Theoretical Issues in Distributed Systems	664
16.1	Notions of Time and State	665
16.2	States and Events in a Distributed System	665
16.3	Time, Clocks and Event Precedences	666
16.4	Recording the State of a Distributed System	672
	<i>Exercises</i>	680
	<i>Bibliography</i>	682
17	Distributed Control Algorithms	684
17.1	Operation of Distributed Control Algorithms	684
17.2	Correctness of Distributed Control Algorithms	686
17.3	Distributed Mutual Exclusion	688
17.4	Distributed Deadlock Handling	693
17.5	Distributed Scheduling Algorithms	698
17.6	Distributed Termination Detection	701
17.7	Election Algorithms	703
17.8	Practical Issues in Using Distributed Control Algorithms	705
	<i>Exercises</i>	707
	<i>Bibliography</i>	708
18	Recovery and Fault Tolerance	710
18.1	Faults, Failures and Recovery	711
18.2	Byzantine Faults and Agreement Protocols	714

xii Contents

18.3	Recovery	715
18.4	Fault Tolerance Techniques	717
18.5	Resiliency	721
	<i>Exercises</i>	722
	<i>Bibliography</i>	723
19	Distributed File Systems	725
19.1	Design Issues in Distributed File Systems	725
19.2	Transparency	728
19.3	Semantics of File Sharing	729
19.4	Fault Tolerance	730
19.5	DFS Performance	735
19.6	Case Studies	740
	<i>Exercises</i>	747
	<i>Bibliography</i>	748
20	Distributed System Security	750
20.1	Issues in Distributed System Security	750
20.2	Message Security	753
20.3	Authentication of Data and Messages	760
20.4	Third-Party Authentication	762
	<i>Exercises</i>	768
	<i>Bibliography</i>	769
	Index	770