

CONTENTS

PREFACE	v
1 CHARACTERIZATION OF DISTRIBUTED SYSTEMS	1
1.1 Introduction	2
1.2 Examples of distributed systems	3
1.3 Resource sharing and the Web	7
1.4 Challenges	16
1.5 Summary	25
2 SYSTEM MODELS	29
2.1 Introduction	30
2.2 Architectural models	31
2.3 Fundamental models	47
2.4 Summary	62
3 NETWORKING AND INTERNETWORKING	65
3.1 Introduction	66
3.2 Types of network	70
3.3 Network principles	73
3.4 Internet protocols	90
3.5 Network case studies: Ethernet, wireless LAN and ATM	110
3.6 Summary	121

4	INTERPROCESS COMMUNICATION	125
4.1	Introduction	126
4.2	The API for the Internet protocols	127
4.3	External data representation and marshalling	138
4.4	Client-server communication	145
4.5	Group communication	153
4.6	Case study: interprocess communication in UNIX	158
4.7	Summary	161
5	DISTRIBUTED OBJECTS AND REMOTE INVOCATION	165
5.1	Introduction	166
5.2	Communication between distributed objects	169
5.3	Remote procedure call	183
5.4	Events and notifications	187
5.5	Java RMI case study	194
5.6	Summary	202
6	OPERATING SYSTEM SUPPORT	207
6.1	Introduction	208
6.2	The operating system layer	209
6.3	Protection	212
6.4	Processes and threads	214
6.5	Communication and invocation	231
6.6	Operating system architecture	242
6.7	Summary	246
7	SECURITY	251
7.1	Introduction	252
7.2	Overview of security techniques	261
7.3	Cryptographic algorithms	272
7.4	Digital signatures	282
7.5	Cryptography pragmatics	289
7.6	Case studies: Needham–Schroeder, Kerberos, SSL & Millicent	291
7.7	Summary	307

8	DISTRIBUTED FILE SYSTEMS	309
8.1	Introduction	310
8.2	File service architecture	318
8.3	Sun Network File System	323
8.4	The Andrew File System	335
8.5	Recent advances	344
8.6	Summary	350
9	NAME SERVICES	353
9.1	Introduction	354
9.2	Name services and the Domain Name System	357
9.3	Directory and discovery services	371
9.4	Case study of the Global Name Service	374
9.5	Case study of the X.500 Directory Service	378
9.6	Summary	382
10	TIME AND GLOBAL STATES	385
10.1	Introduction	386
10.2	Clocks, events and process states	387
10.3	Synchronizing physical clocks	389
10.4	Logical time and logical clocks	397
10.5	Global states	400
10.6	Distributed debugging	409
10.7	Summary	416
11	COORDINATION AND AGREEMENT	419
11.1	Introduction	420
11.2	Distributed mutual exclusion	423
11.3	Elections	431
11.4	Multicast communication	436
11.5	Consensus and related problems	451
11.6	Summary	462

12	TRANSACTIONS AND CONCURRENCY CONTROL	465
12.1	Introduction	466
12.2	Transactions	469
12.3	Nested transactions	480
12.4	Locks	482
12.5	Optimistic concurrency control	497
12.6	Timestamp ordering	501
12.7	Comparison of methods for concurrency control	508
12.8	Summary	509
13	DISTRIBUTED TRANSACTIONS	515
13.1	Introduction	516
13.2	Flat and nested distributed transactions	516
13.3	Atomic commit protocols	519
13.4	Concurrency control in distributed transactions	528
13.5	Distributed deadlocks	531
13.6	Transaction recovery	539
13.7	Summary	549
14	REPLICATION	553
14.1	Introduction	554
14.2	System model and group communication	556
14.3	Fault-tolerant services	565
14.4	Highly available services	572
14.5	Transactions with replicated data	591
14.6	Summary	603
15	DISTRIBUTED MULTIMEDIA SYSTEMS	607
15.1	Introduction	608
15.2	Characteristics of multimedia data	612
15.3	Quality of service management	614
15.4	Resource management	623
15.5	Stream adaptation	625
15.6	Case study: the Tiger video file server	627
15.7	Summary	632

16	DISTRIBUTED SHARED MEMORY	635
16.1	Introduction	636
16.2	Design and implementation issues	640
16.3	Sequential consistency and Ivy	649
16.4	Release consistency and Munin	657
16.5	Other consistency models	663
16.6	Summary	664
17	CORBA CASE STUDY	669
17.1	Introduction	670
17.2	CORBA RMI	671
17.3	CORBA services	686
17.4	Summary	695
18	MACH CASE STUDY	699
18.1	Introduction	700
18.2	Ports, naming and protection	703
18.3	Tasks and threads	705
18.4	Communication model	707
18.5	Communication implementation	710
18.6	Memory management	713
18.7	Summary	719
	REFERENCES	723
	INDEX	757