

大学计算机教育丛书（影印版）

Distributed Operating Systems

Andrew S. Tanenbaum

分布式操作系统



清华大学出版社 • PRENTICE HALL

Distributed Operating Systems

分布式操作系统

Andrew S. Tanenbaum

清华大学出版社
Prentice-Hall International, Inc.

(京)新登字 158 号

Distributed operation systems/Andrew S. Tanenbaum

© 1995 by Prentice-Hall, Inc.

Original Edition Published by Prentice Hall, Inc.

All Rights Reserved.

For sale in Mainland China Only.

Prentice Hall 公司授权清华大学出版社在中国境内(不包括中国香港、澳门特别行政区和台湾地区)独家出版发行本书影印本。

本书任何部分内容,未经出版者书面许可,不得以任何方式抄袭、节录或翻印。

本书封面贴有清华大学出版社激光防伪标签,无标签者不得销售。

北京市版权局著作权合同登记号: 01-97-0167

图书在版编目(CIP)数据

分布式操作系统: 英文/(美)特南鲍姆(Tanenbaum, A. S.)著. - 北京:清华大学出版社, 1996. 12

(大学计算机教育丛书: 影印版)

ISBN 7-302-02411-1

I. 分… II. 特… III. 分布式操作系统-高等学校-教材-英文
IV. TP316

中国版本图书馆 CIP 数据核字(96)第 25162 号

出版者: 清华大学出版社(北京清华大学学研楼, 邮编 100084)

[http:// www.tup.tsinghua.edu.cn](http://www.tup.tsinghua.edu.cn)

印刷者: 清华大学印刷厂

发行者: 新华书店总店北京发行所

开 本: 850×1168 1/32 印张: 19.75

版 次: 1997 年 2 月第 1 版 2000 年 2 月第 6 次印刷

书 号: ISBN 7-302-02411-1/TP·1213

印 数: 16001~19000

定 价: 30.00 元

出版者的话

今天,我们的大学生、研究生和教学、科研工作者,面临的是一个国际化的信息时代。他们将需要随时查阅大量的外文资料;会有更多的机会参加国际性学术交流活动;接待外国学者;走上国际会议的讲坛。作为科技工作者,他们不仅应有与国外同行进行口头和书面交流的能力,更为重要的是,他们必须具备极强的查阅外文资料获取信息的能力。有鉴于此,在国家教委所颁布的“大学英语教学大纲”中有一条规定:专业阅读应作为必修课程开设。同时,在大纲中还规定了这门课程的学时和教学要求。有些高校除开设“专业阅读”课之外,还在某些专业课拟进行英语授课。但教、学双方都苦于没有一定数量的合适的英文原版教材作为教学参考书。为满足这方面的需要,我们陆续精选了一批国外计算机科学方面最新版本的著名教材,进行影印出版。我社获得国外著名出版公司和原著作者的授权将国际先进水平的教材引入我国高等学校,为师生们提供了教学用书,相信会对高校教材改革产生积极的影响。

我们欢迎高校师生将使用影印版教材的效果、意见反馈给我们,更欢迎国内专家、教授积极向我社推荐国外优秀计算机教育教材,以利我们将《大学计算机教育丛书(影印版)》做得更好,更适合高校师生的需要。

清华大学出版社
《大学计算机教育丛书(影印版)》项目组
1999.6

Preface

With the publication of *Distributed Operating Systems* I have now completed my trilogy on operating systems. The three volumes of this trilogy are:

- *Operating Systems: Design and Implementation*
- *Distributed Operating Systems*
- *Modern Operating Systems*

The three volumes are not completely independent, however. For schools having a two-course sequence in operating systems (or an undergraduate course plus a graduate course), one possible choice is to use *Operating Systems: Design and Implementation* in the first course and *Distributed Operating Systems* in the second one.

The former book treats the standard principles of single-processor systems, including processes, synchronization, I/O, deadlocks, memory management, file systems, security, and so on. It also illustrates these principles in great detail through the use of MINIX, a UNIX-clone whose source listing is given in an appendix. MINIX is available on diskette from Prentice Hall for the IBM PC (8088 and up), Atari, Amiga, Macintosh, and SPARC processors.

The latter book (this one), covers distributed operating systems in detail, including communication, synchronization, processes, file systems, and memory management, but this time in the context of distributed systems. Four examples of distributed systems are given in great detail: Amoeba, Mach, Chorus, and DCE. Amoeba is available for free to universities for educational use. It runs

on the Intel 386/486, SPARC, and Sun 3 processors. For information on how to obtain Amoeba please FTP the file *amoebaIntro.ps.Z* from *ftp.cs.vu.nl* or contact the author by electronic mail at *ast@cs.vu.nl*. Potential users should be forewarned that Amoeba is considerably more complex than MINIX: the documentation alone (available by FTP), runs to well over 1000 pages and the system requires at least five large machines and an Ethernet to run well.

By studying these two books in sequence and using both MINIX and Amoeba, students will obtain a thorough knowledge of the principles and practice of both single-processor and distributed operating systems. Now that the trilogy is completed, I plan to revise MINIX and the book describing it.

For universities or computer professionals with less time available, *Modern Operating Systems* can be thought of as a condensed version of the other two books. It provides an introduction to the principles of both single-processor and distributed operating systems, but without the detailed example of MINIX. It also omits many of the advanced topics present in this book, including an introduction to ATM, fault-tolerant distributed systems, real time distributed systems, distributed shared memory, Chorus, DCE, and other topics. In all, about 230 pages of material on distributed systems present in this book have been omitted from *Modern Operating Systems*.

Many people have helped me with this book. I would especially like to thank the following people for reading portions of the manuscript and giving me many useful suggestions for improvement: Irina Athanasii, Henri Bal, Saniya Ben Hassen, David Black, John Carter, Randall Dean, Wiebren de Jonge, John Dugas, Dick Grune, Anoop Gupta, Frans Kaashoek, Marcus Koebler, Hermann Kopetz, Ed Lazowska, Dan Lenoski, Kai Li, Marc Maathuis, David Mosberger, Douglas Orr, Craig Partridge, Carlton Pu, Marc Rozier, Rich Salz, Mike Schroeder, Karsten Schwan, Greg Sharp, Dennis Shasha, Sol Shatz, Jennifer Steiner, Chuck Thacker, John Turek, Walt Tuvell, Leendert van Doorn, Robbert van Renesse, Kees Verstoep, Ellen Zegura, Willy Zwaenpoel, and the anonymous reviewers. My editor, Bill Zobrist, put up with my attempts to get everything perfect with nary a whimper.

Despite all this help, no doubt some errors remain. That seems to be inevitable, no matter how many people read the manuscript. People who wish to report errors should contact me by electronic mail.

Finally, I would like to thank Suzanne again. After eight books, she knows the implications of another one, but her patience and love are boundless. I also want to thank Barbara and Marvin for using their computers and leaving mine alone (except for the printer). Teaching them how to use PC word processing programs has made me appreciate *troff* more than ever. Finally, I would like to thank Little Bram for being quiet while I was writing.

Andrew S. Tanenbaum

Contents

PREFACE	xvi
1 INTRODUCTION TO DISTRIBUTED SYSTEMS	1
1.1 WHAT IS A DISTRIBUTED SYSTEM?	2
1.2 GOALS	3
1.2.1 Advantages of Distributed Systems over Centralized Systems	3
1.2.2 Advantages of Distributed Systems over Independent PCs	6
1.2.3 Disadvantages of Distributed Systems	6
1.3 HARDWARE CONCEPTS	8
1.3.1 Bus-Based Multiprocessors	10
1.3.2 Switched Multiprocessors	12
1.3.3 Bus-Based Multicomputers	13
1.3.4 Switched Multicomputers	14
1.4 SOFTWARE CONCEPTS	15
1.4.1 Network Operating Systems	16
1.4.2 True Distributed Systems	18
1.4.3 Multiprocessor Timesharing Systems	20
1.5 DESIGN ISSUES	22
1.5.1 Transparency	22
1.5.2 Flexibility	25
1.5.3 Reliability	27
1.5.4 Performance	28
1.5.5 Scalability	29
1.6 SUMMARY	31

2 COMMUNICATION IN DISTRIBUTED SYSTEMS**34**

- 2.1 LAYERED PROTOCOLS 35
 - 2.1.1 The Physical Layer 38
 - 2.1.2 The Data Link Layer 38
 - 2.1.3 The Network Layer 40
 - 2.1.4 The Transport Layer 40
 - 2.1.5 The Session Layer 41
 - 2.1.6 The Presentation Layer 41
 - 2.1.7 The Application Layer 42
- 2.2 ASYNCHRONOUS TRANSFER MODE NETWORKS 42
 - 2.2.1 What Is Asynchronous Transfer Mode? 42
 - 2.2.2 The ATM Physical Layer 44
 - 2.2.3 The ATM Layer 45
 - 2.2.4 The ATM Adaptation Layer 46
 - 2.2.5 ATM Switching 47
 - 2.2.6 Some Implications of ATM for Distributed Systems 49
- 2.3 THE CLIENT-SERVER MODEL 50
 - 2.3.1 Clients and Servers 51
 - 2.3.2 An Example Client and Server 52
 - 2.3.3 Addressing 56
 - 2.3.4 Blocking versus Nonblocking Primitives 58
 - 2.3.5 Buffered versus Unbuffered Primitives 61
 - 2.3.6 Reliable versus Unreliable Primitives 63
 - 2.3.7 Implementing the Client-Server Model 65
- 2.4 REMOTE PROCEDURE CALL 68
 - 2.4.1 Basic RPC Operation 68
 - 2.4.2 Parameter Passing 72
 - 2.4.3 Dynamic Binding 77
 - 2.4.4 RPC Semantics in the Presence of Failures 80
 - 2.4.5 Implementation Issues 84
 - 2.4.6 Problem Areas 95
- 2.5 GROUP COMMUNICATION 99
 - 2.5.1 Introduction to Group Communication 99
 - 2.5.2 Design Issues 101
 - 2.5.3 Group Communication in ISIS 110
- 2.6 SUMMARY 114

3 SYNCHRONIZATION IN DISTRIBUTED SYSTEMS**118**

- 3.1 CLOCK SYNCHRONIZATION 119
 - 3.1.1 Logical Clocks 120
 - 3.1.2 Physical Clocks 124
 - 3.1.3 Clock Synchronization Algorithms 127
 - 3.1.4 Use of Synchronized Clocks 132
- 3.2 MUTUAL EXCLUSION 134
 - 3.2.1 A Centralized Algorithm 134
 - 3.2.2 A Distributed Algorithm 135
 - 3.2.3 A Token Ring Algorithm 138
 - 3.2.4 A Comparison of the Three Algorithms 139
- 3.3 ELECTION ALGORITHMS 140
 - 3.3.1 The Bully Algorithm 141
 - 3.3.2 A Ring Algorithm 143
- 3.4 ATOMIC TRANSACTIONS 144
 - 3.4.1 Introduction to Atomic Transactions 144
 - 3.4.2 The Transaction Model 145
 - 3.4.3 Implementation 150
 - 3.4.4 Concurrency Control 154
- 3.5 DEADLOCKS IN DISTRIBUTED SYSTEMS 158
 - 3.5.1 Distributed Deadlock Detection 159
 - 3.5.2 Distributed Deadlock Prevention 163
- 3.6 SUMMARY 165

4 PROCESSES AND PROCESSORS IN DISTRIBUTED SYSTEMS**169**

- 4.1 THREADS 169
 - 4.1.1 Introduction to Threads 170
 - 4.1.2 Thread Usage 171
 - 4.1.3 Design Issues for Threads Packages 174
 - 4.1.4 Implementing a Threads Package 178
 - 4.1.5 Threads and RPC 184
- 4.2 SYSTEM MODELS 186
 - 4.2.1 The Workstation Model 186
 - 4.2.2 Using Idle Workstations 189
 - 4.2.3 The Processor Pool Model 193
 - 4.2.4 A Hybrid Model 197
- 4.3 PROCESSOR ALLOCATION 197
 - 4.3.1 Allocation Models 197

- 4.3.2 Design Issues for Processor Allocation Algorithms 199
- 4.3.3 Implementation Issues for Processor Allocation Algorithms 201
- 4.3.4 Example Processor Allocation Algorithms 203
- 4.4 SCHEDULING IN DISTRIBUTED SYSTEMS 210
- 4.5 FAULT TOLERANCE 212
 - 4.5.1 Component Faults 212
 - 4.5.2 System Failures 213
 - 4.5.3 Synchronous versus Asynchronous Systems 214
 - 4.5.4 Use of Redundancy 214
 - 4.5.5 Fault Tolerance Using Active Replication 215
 - 4.5.6 Fault Tolerance Using Primary-Backup 217
 - 4.5.7 Agreement in Faulty Systems 219
- 4.6 REAL-TIME DISTRIBUTED SYSTEMS 223
 - 4.6.1 What Is a Real-Time System? 223
 - 4.6.2 Design Issues 226
 - 4.6.3 Real-Time Communication 230
 - 4.6.4 Real-Time Scheduling 234
- 4.7 SUMMARY 240

5 DISTRIBUTED FILE SYSTEMS

245

- 5.1 DISTRIBUTED FILE SYSTEM DESIGN 246
 - 5.1.1 The File Service Interface 246
 - 5.1.2 The Directory Server Interface 248
 - 5.1.3 Semantics of File Sharing 253
- 5.2 DISTRIBUTED FILE SYSTEM IMPLEMENTATION 256
 - 5.2.1 File Usage 256
 - 5.2.2 System Structure 258
 - 5.2.3 Caching 262
 - 5.2.4 Replication 268
 - 5.2.5 An Example: Sun's Network File System 272
 - 5.2.6 Lessons Learned 278
- 5.3 TRENDS IN DISTRIBUTED FILE SYSTEMS 279
 - 5.3.1 New Hardware 280
 - 5.3.2 Scalability 282
 - 5.3.3 Wide Area Networking 283
 - 5.3.4 Mobile Users 284
 - 5.3.5 Fault Tolerance 284
 - 5.3.6 Multimedia 285
- 5.4 SUMMARY 285

6 DISTRIBUTED SHARED MEMORY**289**

- 6.1 INTRODUCTION 290
- 6.2 WHAT IS SHARED MEMORY? 292
 - 6.2.1 On-Chip Memory 293
 - 6.2.2 Bus-Based Multiprocessors 293
 - 6.2.3 Ring-Based Multiprocessors 298
 - 6.2.4 Switched Multiprocessors 301
 - 6.2.5 NUMA Multiprocessors 308
 - 6.2.6 Comparison of Shared Memory Systems 312
- 6.3 CONSISTENCY MODELS 315
 - 6.3.1 Strict Consistency 315
 - 6.3.2 Sequential Consistency 317
 - 6.3.3 Causal Consistency 321
 - 6.3.4 PRAM Consistency and Processor Consistency 322
 - 6.3.5 Weak Consistency 325
 - 6.3.6 Release Consistency 327
 - 6.3.7 Entry Consistency 330
 - 6.3.8 Summary of Consistency Models 331
- 6.4 PAGE-BASED DISTRIBUTED SHARED MEMORY 333
 - 6.4.1 Basic Design 334
 - 6.4.2 Replication 334
 - 6.4.3 Granularity 336
 - 6.4.4 Achieving Sequential Consistency 337
 - 6.4.5 Finding the Owner 339
 - 6.4.6 Finding the Copies 342
 - 6.4.7 Page Replacement 343
 - 6.4.8 Synchronization 344
- 6.5 SHARED-VARIABLE DISTRIBUTED SHARED MEMORY 345
 - 6.5.1 Munin 346
 - 6.5.2 Midway 353
- 6.6 OBJECT-BASED DISTRIBUTED SHARED MEMORY 356
 - 6.6.1 Objects 356
 - 6.6.2 Linda 358
 - 6.6.3 Orca 365
- 6.7 COMPARISON 371
- 6.8 SUMMARY 372

7	CASE STUDY 1: AMOEBA	376
7.1	INTRODUCTION TO AMOEBA	376
7.1.1	History of Amoeba	376
7.1.2	Research Goals	377
7.1.3	The Amoeba System Architecture	378
7.1.4	The Amoeba Microkernel	380
7.1.5	The Amoeba Servers	382
7.2	OBJECTS AND CAPABILITIES IN AMOEBA	384
7.2.1	Capabilities	384
7.2.2	Object Protection	385
7.2.3	Standard Operations	387
7.3	PROCESS MANAGEMENT IN AMOEBA	388
7.3.1	Processes	388
7.3.2	Threads	391
7.4	MEMORY MANAGEMENT IN AMOEBA	392
7.4.1	Segments	392
7.4.2	Mapped Segments	393
7.5	COMMUNICATION IN AMOEBA	393
7.5.1	Remote Procedure Call	394
7.5.2	Group Communication in Amoeba	398
7.5.3	The Fast Local Internet Protocol	407
7.6	THE AMOEBA SERVERS	415
7.6.1	The Bullet Server	415
7.6.2	The Directory Server	420
7.6.3	The Replication Server	425
7.6.4	The Run Server	425
7.6.5	The Boot Server	427
7.6.6	The TCP/IP Server	427
7.6.7	Other Servers	428
7.7	SUMMARY	428
8	CASE STUDY 2: MACH	431
8.1	INTRODUCTION TO MACH	431
8.1.1	History of Mach	431
8.1.2	Goals of Mach	433
8.1.3	The Mach Microkernel	433
8.1.4	The Mach BSD UNIX Server	435

8.2	PROCESS MANAGEMENT IN MACH	436
8.2.1	Processes	436
8.2.2	Threads	439
8.2.3	Scheduling	442
8.3	MEMORY MANAGEMENT IN MACH	445
8.3.1	Virtual Memory	446
8.3.2	Memory Sharing	449
8.3.3	External Memory Managers	452
8.3.4	Distributed Shared Memory in Mach	456
8.4	COMMUNICATION IN MACH	457
8.4.1	Ports	457
8.4.2	Sending and Receiving Messages	464
8.4.3	The Network Message Server	469
8.5	UNIX EMULATION IN MACH	471
8.6	SUMMARY	472

9 CASE STUDY 3: CHORUS

475

9.1	INTRODUCTION TO CHORUS	475
9.1.1	History of Chorus	476
9.1.2	Goals of Chorus	477
9.1.3	System Structure	478
9.1.4	Kernel Abstractions	479
9.1.5	Kernel Structure	481
9.1.6	The UNIX Subsystem	483
9.1.7	The Object-Oriented Subsystem	483
9.2	PROCESS MANAGEMENT IN CHORUS	483
9.2.1	Processes	484
9.2.2	Threads	485
9.2.3	Scheduling	486
9.2.4	Traps, Exceptions, and Interrupts	487
9.2.5	Kernel Calls for Process Management	488
9.3	MEMORY MANAGEMENT IN CHORUS	490
9.3.1	Regions and Segments	490
9.3.2	Mappers	491
9.3.3	Distributed Shared Memory	492
9.3.4	Kernel Calls for Memory Management	493

- 9.4 COMMUNICATION IN CHORUS 495
 - 9.4.1 Messages 495
 - 9.4.2 Ports 495
 - 9.4.3 Communication Operations 496
 - 9.4.4 Kernel Calls for Communication 498
- 9.5 UNIX EMULATION IN CHORUS 499
 - 9.5.1 Structure of a UNIX Process 500
 - 9.5.2 Extensions to UNIX 500
 - 9.5.3 Implementation of UNIX on Chorus 501
- 9.6 COOL: AN OBJECT-ORIENTED SUBSYSTEM 507
 - 9.6.1 The COOL Architecture 507
 - 9.6.2 The COOL Base Layer 507
 - 9.6.3 The COOL Generic Runtime System 509
 - 9.6.4 The Language Runtime System 509
 - 9.6.5 Implementation of COOL 510
- 9.7 COMPARISON OF AMOEBA, MACH, AND CHORUS 510
 - 9.7.1 Philosophy 511
 - 9.7.2 Objects 512
 - 9.7.3 Processes 513
 - 9.7.4 Memory Model 514
 - 9.7.5 Communication 515
 - 9.7.6 Servers 516
- 9.8 SUMMARY 517

10 CASE STUDY 4: DCE

520

- 10.1 INTRODUCTION TO DCE 520
 - 10.1.1 History of DCE 520
 - 10.1.2 Goals of DCE 521
 - 10.1.3 DCE Components 522
 - 10.1.4 Cells 525
- 10.2 THREADS 527
 - 10.2.1 Introduction to DCE Threads 527
 - 10.2.2 Scheduling 529
 - 10.2.3 Synchronization 530
 - 10.2.4 Thread Calls 531
- 10.3 REMOTE PROCEDURE CALL 535
 - 10.3.1 Goals of DCE RPC 535
 - 10.3.2 Writing a Client and a Server 536
 - 10.3.3 Binding a Client to a Server 538
 - 10.3.4 Performing an RPC 539

CONTENTS

xiii

10.4	TIME SERVICE	540
10.4.1	DTS Time Model	541
10.4.2	DTS Implementation	543
10.5	DIRECTORY SERVICE	544
10.5.1	Names	546
10.5.2	The Cell Directory Service	547
10.5.3	The Global Directory Service	549
10.6	SECURITY SERVICE	554
10.6.1	Security Model	555
10.6.2	Security Components	557
10.6.3	Tickets and Authenticators	558
10.6.4	Authenticated RPC	559
10.6.5	ACLs	562
10.7	DISTRIBUTED FILE SYSTEM	564
10.7.1	DFS Interface	565
10.7.2	DFS Components in the Server Kernel	566
10.7.3	DFS Components in the Client Kernel	569
10.7.4	DFS Components in User Space	571
10.8	SUMMARY	573

11 BIBLIOGRAPHY AND SUGGESTED READINGS

577

11.1	SUGGESTED READINGS	577
11.2	ALPHABETICAL BIBLIOGRAPHY	584

INDEX

603

Introduction to Distributed Systems

Computer systems are undergoing a revolution. From 1945, when the modern computer era began, until about 1985, computers were large and expensive. Even minicomputers normally cost tens of thousands of dollars each. As a result, most organizations had only a handful of computers, and for lack of a way to connect them, these operated independently from one another.

Starting in the mid-1980s, however, two advances in technology began to change that situation. The first was the development of powerful microprocessors. Initially, these were 8-bit machines, but soon 16-, 32-, and even 64-bit CPUs became common. Many of these had the computing power of a decent-sized mainframe (i.e., large) computer, but for a fraction of the price.

The amount of improvement that has occurred in computer technology in the past half century is truly staggering and totally unprecedented in other industries. From a machine that cost 10 million dollars and executed 1 instruction per second, we have come to machines that cost 1000 dollars and execute 10 million instructions per second, a price/performance gain of 10^{11} . If cars had improved at this rate in the same time period, a Rolls Royce would now cost 10 dollars and get a billion miles per gallon. (Unfortunately, it would probably also have a 200-page manual telling how to open the door.)

The second development was the invention of high-speed computer networks. The **local area networks** or LANs allow dozens, or even hundreds, of machines within a building to be connected in such a way that small amounts of

information can be transferred between machines in a millisecond or so. Larger amounts of data can be moved between machines at rates of 10 to 100 million bits/sec and sometimes more. The **wide area networks** or **WANs** allow millions of machines all over the earth to be connected at speeds varying from 64 Kbps (kilobits per second) to gigabits per second for some advanced experimental networks.

The result of these technologies is that it is now not only feasible, but easy, to put together computing systems composed of large numbers of CPUs connected by a high-speed network. They are usually called **distributed systems**, in contrast to the previous **centralized systems** (or **single-processor systems**) consisting of a single CPU, its memory, peripherals, and some terminals.

There is only one fly in the ointment: software. Distributed systems need radically different software than centralized systems do. In particular, the necessary operating systems are only beginning to emerge. The first few steps have been taken, but there is still a long way to go. Nevertheless, enough is already known about these distributed operating systems that we can present the basic ideas. The rest of this book is devoted to studying concepts, implementation, and examples of distributed operating systems.

1.1. WHAT IS A DISTRIBUTED SYSTEM?

Various definitions of distributed systems have been given in the literature, none of them satisfactory and none of them in agreement with any of the others. For our purposes it is sufficient to give a loose characterization:

A distributed system is a collection of independent computers that appear to the users of the system as a single computer.

This definition has two aspects. The first one deals with hardware: the machines are autonomous. The second one deals with software: the users think of the system as a single computer. Both are essential. We will come back to these points later in this chapter, after going over some background material on both the hardware and the software.

Rather than going further with definitions, it is probably more helpful to give several examples of distributed systems. As a first example, consider a network of workstations in a university or company department. In addition to each user's personal workstation, there might be a pool of processors in the machine room that are not assigned to specific users but are allocated dynamically as needed. Such a system might have a single file system, with all files accessible from all machines in the same way and using the same path name. Furthermore, when a user typed a command, the system could look for the best place to execute that command, possibly on the user's own workstation, possibly on an idle