

大学计算机教育丛书(影印版)

网络互连技术系列

THIRD EDITION

Internetworking with TCP/IP

VOLUME I

Principles, Protocols, and Architecture

TCP/IP 网络互连技术

卷 I (第 3 版)

原理, 协议和体系结构

DOUGLAS E. COMER

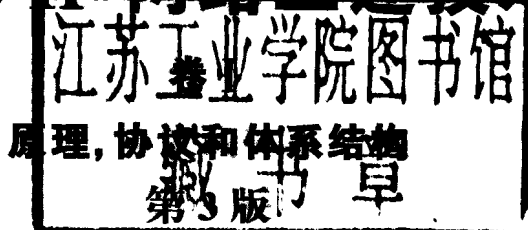


清华大学出版社 · PRENTICE HALL

<http://www.tup.tsinghua.edu.cn>

Internetworking With TCP/IP
Vol I :
Principles, Protocols, and Architecture
Third Edition

TCP/IP 网络互连技术



DOUGLAS E. COMER

Department of Computer Sciences

(京) 新登字 158 号

Internetworking with TCP/IP Vol I: Principles, Protocols, and architecture
3rd ed. /Douglas E. Comer

© 1995 by Prentice Hall, Inc.

Original English Language Edition published by Prentice Hall, Inc., a Simon & Schuster Company.

All Rights Reserved.

For sale in Mainland China only.

本书影印版由西蒙与舒斯特国际出版公司授权清华大学出版社在中国境内(不包括中国香港特别行政区、澳门地区和台湾地区)独家出版发行。

未经出版者书面许可,不得用任何方式复制或抄袭本书的任何部分。

本书封面贴有 Prentice Hall 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: 01-98-0957

图书在版编目(CIP)数据

TCP/IP 网络互连技术 卷 I:第3版:英文(美)科默(Comer, D. E.)著. - 影印版. - 北京:清华大学出版社,1998.7

(大学计算机教育丛书)

ISBN 7-302-02946-6

I. T… II. 科… III. 计算机网络-连接技术-英文 IV. TP393

中国版本图书馆 CIP 数据核字(98)第 09294 号

出版者:清华大学出版社(北京清华大学校内,邮编 100084)

<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:20

版次:1998年9月第1版 1999年2月第2次印刷

书号:ISBN 7-302-02946-6/TP·1557

印数:5001~10000

定价:30.00 元

出版前言

清华大学出版社与 Prentice Hall 出版公司合作推出的“大学计算机教育丛书(影印版)”和“ATM 与 B-ISDN 技术丛书(影印版)”受到了广大读者的欢迎。很多读者通过电话、信函、电子函件给我们的工作以积极的评价,并提出了不少中肯的建议。其中,很多读者希望我们能够出版一些网络方面较深层次的书籍,这也就成为我们出版这套“网络互连技术系列”的最初动机。

众所周知,网络协议是网络与通信技术的关键组成部分。而今,因特网技术、移动通信技术的飞速发展,为网络协议注入了新内容。本套丛书以 Douglas Comer 教授的网络协议的经典名著 TCP/IP 网络互连技术系列为主干,并补充以论述新协议如 IPV6 和移动 IP 等国外最新专著,力求为从事网络互连技术研究与开发的人员以及大专院校师生提供充分的技术支持。

衷心希望所有阅读这套丛书的读者能从中受益。

清华大学出版社
Prentice Hall 公司

1998.9

Additional Enthusiastic Comments About Internetworking With TCP/IP Volume 1

"Unquestionably THE reference for TCP/IP; both informative and easy to read, this book is liked by both novice and experienced."

– Raj Yavatkar
University of Kentucky
US Editor, Computer Communications

"The third edition maintains Comer's Internetworking with TCP/IP as the acknowledged leader in TCP/IP books by adding up-to-the-minute material on ATM, CIDR, firewalls, DHCP and the next version of IP, IPng."

– Ralph Droms
Bucknell University
IETF Working Group Chair

"Doug Comer remains the first and best voice of Internet technology. Despite the legion of 'Internet carpetbaggers' (the current crop of 'authors' who can barely spell F-T-P) which contributes noise – but no knowledge – on the Internet and its infrastructure, Dr. Comer shines through as the premiere source for lucid explanations and accurate information. He sets a standard for which many strive, but precious few attain."

– Marshall Rose
Dover Beach Consulting
IETF Area Director

"Comer's Volume 1 drastically changed the course of networking history."

– Dan Lynch
Interop Company
IAB Member

"When you need to teach the details of TCP/IP, you need the latest information. Once again, Comer separates the chaff from the wheat with his latest edition of the TCP/IP book that a generation of networkers grew up with."

– Shawn Ostermann
Ohio University

Foreword

Professor Douglas Comer's book has become *the* classic text for an introduction to TCP/IP. Writing an introduction to TCP/IP for the uninitiated is a very difficult task. While combining the explanation of the general principles of computer communication with the specific examples from the TCP/IP protocol suite, Doug Comer has provided a very readable book.

While this book is specifically about the TCP/IP protocol suite, it is a good book for learning about computer communications protocols in general. The principles of architecture, layering, multiplexing, encapsulation, addressing and address mapping, routing, and naming are quite similar in any protocol suite, though, of course, different in detail.

Computer communication protocols do not do anything themselves. Like operating systems, they are in the service of application processes. Processes are the active elements that request communication and are the ultimate senders and receivers of the data transmitted. The various layers of protocols are like the various layers in a computer operating system, especially the file system. Understanding protocol architecture is like understanding operating system architecture. In this book Doug Comer has taken the "bottom up" approach – starting with the physical networks and moving up in levels of abstraction to the applications.

Since application processes are the active elements using the communication supported by the protocols, TCP/IP is an "interprocess communication" (IPC) mechanism. While there are several experiments in progress with operating system style message passing and procedure call types of IPC based on IP, the focus in this book is on more traditional applications that use the UDP datagram or TCP logical connection forms of IPC. Typically in operating systems there is a set of functions provided by the operating system to the application processes. This system call interface usually includes calls for opening, reading, writing, and closing files, among other things. In many systems there are similar system calls for IPC functions including network communication. As an example of such an interface Doug Comer presents an overview of the socket interface.

One of the key ideas inherent in TCP/IP and in the title of this book is "internet-working." The power of a communication system is directly related to the number of entities in that system. The telephone network is very useful because (nearly) all the telephones are connected to one network (as it appears to the users). Computer communication systems and networks are currently separated and fragmented. As more users and enterprises adopt TCP/IP as their network communication technology and are joining the Internet this is becoming less of a problem, but there is still a long way to

go. The goal of interconnection and internetworking, to have a single powerful computer communication network, is fundamental to the design of TCP/IP.

Essential to internetworking is addressing, and a universal protocol – the Internet Protocol. Of course, the individual networks have their own protocols which are used to carry the IP datagrams, and there must be a mapping between the individual network address and the IP address. Over the lifetime of TCP/IP, the nature of these individual networks have changed from the early days of the ARPANET to the recently developed ATM networks. A new chapter in this edition discusses IP over ATM networks. This book now includes recent developments in Dynamic Host Configuration (DHCP) that will ease the administration of networks and the installation of new computers.

To have an internetwork, the individual networks must be connected. The connecting devices are called routers. Further, these routers must have some procedures for forwarding data from one network to the next. The data is in the form of IP datagrams and the destination is specified by an IP address, but the router must make a routing decision based on the IP address and what it knows about the connectivity of the networks making up the Internet. The procedures for distributing the current connectivity information to the routers are called routing algorithms, and these are currently the subject of much study and development. In particular, the recent development of the Classless InterDomain Routing (CIDR) technique to reduce the amount of routing information exchanged is important.

Like all communication systems, the TCP/IP protocol suite is an unfinished system. It is evolving to meet changing requirements and new opportunities. Thus, this book is, in a sense, a snapshot of TCP/IP. And, as Doug Comer points out, there are many loose ends. With the recent rapid growth of the Internet there is concern about it outgrowing the capabilities of the TCP/IP protocols, particularly the address space. In response the research and engineering community has developed a “next generation” version of the Internet Protocol called IPng. Many of the enterprises now joining the Internet have concerns about security. A new chapter in this edition discusses the security and firewalls.

Most chapters end with a few pointers to material “for further study.” Many of these refer to memos of the RFC series of notes. This series of notes is the result of a policy of making the working ideas and the protocol specifications developed by the TCP/IP research and development community widely available. This availability of the basic and detailed information about these protocols, and the availability of the early implementations of them, has had much to do with their current widespread use. This commitment to public documentation at this level of detail is unusual for a research effort, and has had significant benefits for the development of computer communication.

This book brings together information about the various parts of the TCP/IP architecture and protocols and makes it accessible. Its publication is a very significant milestone in the evolution of computer communications.

Jon Postel,
Associate Director for Networking
Information Sciences Institute
University of Southern California

Preface

The world has changed dramatically since the second edition of this book was published. It hardly seems possible only four years have elapsed. When I began the second edition in the summer of 1990, the Internet had grown to nearly 300,000 host computers, up from 5,000 hosts when the book was first written. At the time, we marveled at how large an obscure research project had become. Cynics predicted that continued growth would lead to a complete collapse by 1993. Instead of collapsing, the Internet has continued its explosive expansion; the “large” Internet of 1990 is only 7% of the current Internet.

TCP/IP and the Internet have accommodated change well. The basic technology has survived over a decade of exponential growth and the associated increases in traffic. The protocols have worked over new high-speed network technologies, and the design has handled applications that could not be imagined a decade ago. Of course, the entire protocol suite has not remained static. New protocols have been deployed, and new techniques have been developed to adapt existing protocols to new network technologies. Changes are documented in RFCs, which have increased by over 50 percent.

This edition contains updated information throughout the text (including use of the commercially popular term *IP router* in place of the traditional scientific term *IP gateway*) as well as new material that describes technical advances and changes. The chapter on subnet addressing now describes supernetting as well as subnetting, and shows how the two techniques are motivated by the same goal. The chapter on bootstrapping explains a significant advance that will eliminate the need for manual configuration of host computers and allow a computer to obtain an IP address automatically: the Dynamic Host Configuration Protocol (DHCP). The chapter on TCP includes a description of Silly Window Syndrome and an explanation of the heuristics TCP uses to prevent the problem. The chapter on electronic mail includes a description of the Multipurpose Internet Mail Extensions (MIME), which permit non-ASCII data to be sent in a standard e-mail message.

Three new chapters contain detailed information about significant developments. Chapter 18 explains how TCP/IP is being used over ATM networks. The chapter discusses the organization of ATM hardware, the purpose of adaptation layer protocols, IP encapsulation, address binding, routing, and virtual circuit management. The chapter illustrates how a connectionless protocol like IP can use the connection-oriented interface that ATM provides. Chapter 28 covers a topic that is crucial to many organizations as they contemplate connecting to the global Internet – security. The chapter describes the Internet firewall concept, and shows how a firewall architecture can be

used to protect networks and computers inside an organization from unwanted access. The chapter also discusses the principles underlying a two-level firewall design, and considers outside access from a secure computer. Finally, a new chapter is devoted to what may be the most significant change in TCP/IP since its inception: the imminent adoption of a next generation Internet Protocol (IPng). Chapter 29 describes the protocol that the IETF has developed to serve as IPng. Although it has not been thoroughly tested or approved as a permanent standard, the new design appears to be the consensus choice. The chapter presents the proposed design and address assignment scheme.

The third edition retains the same general contents and overall organization as the second edition. The entire text focuses on the concept of internetworking in general and the TCP/IP internet technology in particular. Internetworking is a powerful abstraction that allows us to deal with the complexity of multiple underlying communication technologies. It hides the details of network hardware and provides a high level communication environment. The text reviews both the architecture of network interconnections and the principles underlying protocols that make such interconnected networks function as a single, unified communication system. It also shows how an internet communication system can be used for distributed computation.

After reading this book, you will understand how it is possible to interconnect multiple physical networks into a coordinated system, how internet protocols operate in that environment, and how application programs use the resulting system. As a specific example, you will learn the details of the global TCP/IP Internet, including the architecture of its router system and the application protocols it supports. In addition, you will understand some of the limitations of the internet approach.

Designed as both a college text and as a professional reference, the book is written at an advanced undergraduate or graduate level. For professionals, the book provides a comprehensive introduction to the TCP/IP technology and the architecture of the Internet. Although it is not intended to replace protocol standards, the book is an excellent starting point for learning about internetworking because it provides a uniform overview that emphasizes principles. Moreover, it gives the reader perspective that can be extremely difficult to obtain from individual protocol documents.

When used in the classroom, the text provides more than sufficient material for a single semester network course at either the undergraduate or graduate level. Such a course can be extended to a two-semester sequence if accompanied by programming projects and readings from the literature. For undergraduate courses, many of the details are unnecessary. Students should be expected to grasp the basic concepts described in the text, and they should be able to describe or use them. At the graduate level, students should be expected to use the material here as a basis for further exploration. They should understand the details well enough to answer exercises or solve problems that require them to explore extensions and subtleties. Many of the exercises suggest such subtleties; solving them often requires students to read protocol standards and apply creative energy to comprehend consequences.

At all levels, hands-on experience sharpens the concepts and helps students gain intuition. Thus, I encourage instructors to invent projects that force students to use Internet services and protocols. The semester project in my graduate Internetworking

course at Purdue requires students to build an IP router. We supply hardware and the source code for an operating system, including device drivers for network interfaces; students build a working router that interconnects three networks with different MTUs. The course is extremely rigorous, students work in teams, and the results have been impressive (many industries recruit graduates from the course). Although such experimentation is safest when the instructional laboratory network is isolated from production computing facilities, we have found that students exhibit the most enthusiasm, and benefit the most, when they have access to a functional TCP/IP internet.

The book is organized into four main parts. Chapters 1 and 2 form an introduction that provides an overview and discusses existing network technologies. In particular, Chapter 2 reviews physical network hardware. The intention is to provide basic intuition about what is possible, not to spend inordinate time on hardware details. Chapters 3-13 describe the TCP/IP Internet from the viewpoint of a single host, showing the protocols a host contains and how they operate. They cover the basics of Internet addressing and routing as well as the notion of protocol layering. Chapters 14-18 and 28 describe the architecture of an internet when viewed globally. They explore routing architecture and the protocols routers use to exchange routing information. Finally, Chapters 19-27 discuss application level services available in the Internet. They present the client-server model of interaction, and give several examples of client and server software.

The chapters have been organized bottom up. They begin with an overview of hardware and continue to build new functionality on top of it. This view will appeal to anyone who has developed Internet software because it follows the same pattern one uses in implementation. The concept of layering does not appear until Chapter 11. The discussion of layering emphasizes the distinction between conceptual layers of functionality and the reality of layered protocol software in which multiple objects appear at each layer.

A modest background is required to understand the material. The reader is expected to have a basic understanding of computer systems, and to be familiar with data structures like stacks, queues, and trees. Readers need basic intuition about the organization of computer software into an operating system that supports concurrent programming and application programs that users invoke to perform computation. Readers do not need sophisticated mathematics, nor do they need to know information theory or theorems from data communications; the book describes the physical network as a black box around which an internetwork can be built. It states design principles in English and discusses motivations and consequences.

I thank all the people who have contributed to versions of this book. John Lin provided extensive assistance with this edition, including classifying RFCs. Ralph Droms reviewed the chapter on bootstrapping, and Sandeep Kumar, Steve Lodin, and Christoph Schuba, from the COAST security project at Purdue, commented on the security chapter. Special thanks go to my wife, Chris, whose careful editing made many improvements in wording.

Contents

Foreword	xix
-----------------	------------

Preface	xxi
----------------	------------

Chapter 1 Introduction And Overview	1
--	----------

1.1	<i>The Motivation For Internetworking</i>	1
1.2	<i>The TCP/IP Internet</i>	2
1.3	<i>Internet Services</i>	3
1.4	<i>History And Scope Of The Internet</i>	6
1.5	<i>The Internet Architecture Board</i>	8
1.6	<i>The IAB Reorganization</i>	9
1.7	<i>The Internet Society</i>	11
1.8	<i>Internet Request For Comments</i>	11
1.9	<i>Internet Protocols And Standardization</i>	12
1.10	<i>Future Growth And Technology</i>	12
1.11	<i>Organization Of The Text</i>	13
1.12	<i>Summary</i>	14

Chapter 2 Review Of Underlying Network Technologies	17
--	-----------

2.1	<i>Introduction</i>	17
2.2	<i>Two Approaches To Network Communication</i>	18
2.3	<i>Wide Area And Local Area Networks</i>	19
2.4	<i>Ethernet Technology</i>	20
2.5	<i>Fiber Distributed Data Interconnect (FDDI)</i>	32
2.6	<i>Asynchronous Transfer Mode</i>	36
2.7	<i>ARPANET Technology</i>	37
2.8	<i>National Science Foundation Networking</i>	39
2.9	<i>ANSNET</i>	44

2.10	<i>A Planned Wide Area Backbone</i>	44
2.11	<i>Other Technologies Over Which TCP/IP Has Been Used</i>	44
2.12	<i>Summary And Conclusion</i>	47

Chapter 3 Internetworking Concept And Architectural Model 49

3.1	<i>Introduction</i>	49
3.2	<i>Application-Level Interconnection</i>	49
3.3	<i>Network-Level Interconnection</i>	50
3.4	<i>Properties Of The Internet</i>	51
3.5	<i>Internet Architecture</i>	52
3.6	<i>Interconnection Through IP Routers</i>	52
3.7	<i>The User's View</i>	54
3.8	<i>All Networks Are Equal</i>	54
3.9	<i>The Unanswered Questions</i>	55
3.10	<i>Summary</i>	56

Chapter 4 Internet Addresses 59

4.1	<i>Introduction</i>	59
4.2	<i>Universal Identifiers</i>	59
4.3	<i>Three Primary Classes Of IP Addresses</i>	60
4.4	<i>Addresses Specify Network Connections</i>	61
4.5	<i>Network And Broadcast Addresses</i>	61
4.6	<i>Limited Broadcast</i>	62
4.7	<i>Interpreting Zero To Mean "This"</i>	62
4.8	<i>Weaknesses In Internet Addressing</i>	63
4.9	<i>Dotted Decimal Notation</i>	65
4.10	<i>Loopback Address</i>	65
4.11	<i>Summary Of Special Address Conventions</i>	66
4.12	<i>Internet Addressing Authority</i>	66
4.13	<i>An Example</i>	67
4.14	<i>Network Byte Order</i>	69
4.15	<i>Summary</i>	70

Chapter 5 Mapping Internet Addresses To Physical Addresses (ARP) 73

5.1	<i>Introduction</i>	73
5.2	<i>The Address Resolution Problem</i>	73
5.3	<i>Two Types Of Physical Addresses</i>	74
5.4	<i>Resolution Through Direct Mapping</i>	74

5.5	<i>Resolution Through Dynamic Binding</i>	75
5.6	<i>The Address Resolution Cache</i>	76
5.7	<i>ARP Refinements</i>	77
5.8	<i>Relationship Of ARP To Other Protocols</i>	77
5.9	<i>ARP Implementation</i>	77
5.10	<i>ARP Encapsulation And Identification</i>	79
5.11	<i>ARP Protocol Format</i>	79
5.12	<i>Summary</i>	81

Chapter 6 Determining An Internet Address At Startup (RARP) 83

6.1	<i>Introduction</i>	83
6.2	<i>Reverse Address Resolution Protocol (RARP)</i>	84
6.3	<i>Timing RARP Transactions</i>	86
6.4	<i>Primary And Backup RARP Servers</i>	86
6.5	<i>Summary</i>	87

Chapter 7 Internet Protocol: Connectionless Datagram Delivery 89

7.1	<i>Introduction</i>	89
7.2	<i>A Virtual Network</i>	89
7.3	<i>Internet Architecture And Philosophy</i>	90
7.4	<i>The Concept Of Unreliable Delivery</i>	90
7.5	<i>Connectionless Delivery System</i>	91
7.6	<i>Purpose Of The Internet Protocol</i>	91
7.7	<i>The Internet Datagram</i>	91
7.8	<i>Internet Datagram Options</i>	100
7.9	<i>Summary</i>	106

Chapter 8 Internet Protocol: Routing IP Datagrams 109

8.1	<i>Introduction</i>	109
8.2	<i>Routing In An Internet</i>	109
8.3	<i>Direct And Indirect Delivery</i>	111
8.4	<i>Table-Driven IP Routing</i>	113
8.5	<i>Next-Hop Routing</i>	113
8.6	<i>Default Routes</i>	115
8.7	<i>Host-Specific Routes</i>	115
8.8	<i>The IP Routing Algorithm</i>	116
8.9	<i>Routing With IP Addresses</i>	116
8.10	<i>Handling Incoming Datagrams</i>	118

8.11	<i>Establishing Routing Tables</i>	119
8.12	<i>Summary</i>	119

Chapter 9 Internet Protocol: Error And Control Messages (ICMP) 123

9.1	<i>Introduction</i>	123
9.2	<i>The Internet Control Message Protocol</i>	123
9.3	<i>Error Reporting vs. Error Correction</i>	124
9.4	<i>ICMP Message Delivery</i>	125
9.5	<i>ICMP Message Format</i>	126
9.6	<i>Testing Destination Reachability And Status (Ping)</i>	127
9.7	<i>Echo Request And Reply Message Format</i>	128
9.8	<i>Reports Of Unreachable Destinations</i>	128
9.9	<i>Congestion And Datagram Flow Control</i>	130
9.10	<i>Source Quench Format</i>	130
9.11	<i>Route Change Requests From Routers</i>	131
9.12	<i>Detecting Circular Or Excessively Long Routes</i>	133
9.13	<i>Reporting Other Problems</i>	134
9.14	<i>Clock Synchronization And Transit Time Estimation</i>	134
9.15	<i>Information Request And Reply Messages</i>	136
9.16	<i>Obtaining A Subnet Mask</i>	136
9.17	<i>Summary</i>	137

Chapter 10 Subnet And Supernet Address Extensions 139

10.1	<i>Introduction</i>	139
10.2	<i>Review Of Relevant Facts</i>	139
10.3	<i>Minimizing Network Numbers</i>	140
10.4	<i>Transparent Routers</i>	141
10.5	<i>Proxy ARP</i>	142
10.6	<i>Subnet Addressing</i>	143
10.7	<i>Flexibility In Subnet Address Assignment</i>	146
10.8	<i>Implementation Of Subnets With Masks</i>	147
10.9	<i>Subnet Mask Representation</i>	148
10.10	<i>Routing In The Presence Of Subnets</i>	149
10.11	<i>The Subnet Routing Algorithm</i>	150
10.12	<i>A Unified Routing Algorithm</i>	151
10.13	<i>Maintenance Of Subnet Masks</i>	152
10.14	<i>Broadcasting To Subnets</i>	152
10.15	<i>Supernet Addressing</i>	153
10.16	<i>The Effect Of Supernetting On Routing</i>	154
10.17	<i>Summary</i>	155

Chapter 11 Protocol Layering	159
11.1 Introduction	159
11.2 The Need For Multiple Protocols	159
11.3 The Conceptual Layers Of Protocol Software	160
11.4 Functionality Of The Layers	163
11.5 X.25 And Its Relation To The ISO Model	164
11.6 Differences Between X.25 And Internet Layering	167
11.7 The Protocol Layering Principle	169
11.8 Layering In The Presence Of Network Substructure	171
11.9 Two Important Boundaries In The TCP/IP Model	173
11.10 The Disadvantage Of Layering	174
11.11 The Basic Idea Behind Multiplexing And Demultiplexing	174
11.12 Summary	176
 Chapter 12 User Datagram Protocol (UDP)	 179
12.1 Introduction	179
12.2 Identifying The Ultimate Destination	179
12.3 The User Datagram Protocol	180
12.4 Format Of UDP Messages	181
12.5 UDP Pseudo-Header	182
12.6 UDP Encapsulation And Protocol Layering	183
12.7 Layering And The UDP Checksum Computation	185
12.8 UDP Multiplexing, Demultiplexing, And Ports	185
12.9 Reserved And Available UDP Port Numbers	186
12.10 Summary	188
 Chapter 13 Reliable Stream Transport Service (TCP)	 191
13.1 Introduction	191
13.2 The Need For Stream Delivery	191
13.3 Properties Of The Reliable Delivery Service	192
13.4 Providing Reliability	193
13.5 The Idea Behind Sliding Windows	195
13.6 The Transmission Control Protocol	198
13.7 Ports, Connections, And Endpoints	199
13.8 Passive And Active Opens	201
13.9 Segments, Streams, And Sequence Numbers	201
13.10 Variable Window Size And Flow Control	202
13.11 TCP Segment Format	203

13.12	<i>Out Of Band Data</i>	205
13.13	<i>Maximum Segment Size Option</i>	206
13.14	<i>TCP Checksum Computation</i>	207
13.15	<i>Acknowledgements And Retransmission</i>	208
13.16	<i>Timeout And Retransmission</i>	209
13.17	<i>Accurate Measurement Of Round Trip Samples</i>	211
13.18	<i>Karn's Algorithm And Timer Backoff</i>	212
13.19	<i>Responding To High Variance In Delay</i>	213
13.20	<i>Response To Congestion</i>	214
13.21	<i>Establishing A TCP Connection</i>	216
13.22	<i>Initial Sequence Numbers</i>	217
13.23	<i>Closing a TCP Connection</i>	217
13.24	<i>TCP Connection Reset</i>	219
13.25	<i>TCP State Machine</i>	219
13.26	<i>Forcing Data Delivery</i>	221
13.27	<i>Reserved TCP Port Numbers</i>	221
13.28	<i>TCP Performance</i>	221
13.29	<i>Silly Window Syndrome And Small Packets</i>	223
13.30	<i>Avoiding Silly Window Syndrome</i>	224
13.31	<i>Summary</i>	227

Chapter 14 Routing: Cores, Peers, And Algorithms (GGP)

231

14.1	<i>Introduction</i>	231
14.2	<i>The Origin Of Routing Tables</i>	232
14.3	<i>Routing With Partial Information</i>	233
14.4	<i>Original Internet Architecture And Cores</i>	234
14.5	<i>Core Routers</i>	235
14.6	<i>Beyond The Core Architecture To Peer Backbones</i>	238
14.7	<i>Automatic Route Propagation</i>	240
14.8	<i>Vector Distance (Bellman-Ford) Routing</i>	240
14.9	<i>Gateway-To-Gateway Protocol (GGP)</i>	242
14.10	<i>GGP Message Formats</i>	243
14.11	<i>Link-State (SPF) Routing</i>	245
14.12	<i>SPF Protocols</i>	246
14.13	<i>Summary</i>	246

Chapter 15 Routing: Autonomous Systems (EGP)

249

15.1	<i>Introduction</i>	249
15.2	<i>Adding Complexity To The Architectural Model</i>	249
15.3	<i>A Fundamental Idea: Extra Hops</i>	250

15.4	<i>Autonomous System Concept</i>	252
15.5	<i>Exterior Gateway Protocol (EGP)</i>	254
15.6	<i>EGP Message Header</i>	255
15.7	<i>EGP Neighbor Acquisition Messages</i>	256
15.8	<i>EGP Neighbor Reachability Messages</i>	257
15.9	<i>EGP Poll Request Messages</i>	258
15.10	<i>EGP Routing Update Messages</i>	259
15.11	<i>Measuring From The Receiver's Perspective</i>	261
15.12	<i>The Key Restriction Of EGP</i>	262
15.13	<i>Technical Problems</i>	264
15.14	<i>Decentralization Of Internet Architecture</i>	264
15.15	<i>Beyond Autonomous Systems</i>	264
15.16	<i>Summary</i>	265

Chapter 16 Routing: In An Autonomous System (RIP, OSPF, HELLO) 267

16.1	<i>Introduction</i>	267
16.2	<i>Static Vs. Dynamic Interior Routes</i>	267
16.3	<i>Routing Information Protocol (RIP)</i>	270
16.4	<i>The Hello Protocol</i>	276
16.5	<i>Combining RIP, Hello, And EGP</i>	278
16.6	<i>The Open SPF Protocol (OSPF)</i>	279
16.7	<i>Routing With Partial Information</i>	286
16.8	<i>Summary</i>	286

Chapter 17 Internet Multicasting (IGMP) 289

17.1	<i>Introduction</i>	289
17.2	<i>Hardware Broadcast</i>	289
17.3	<i>Hardware Multicast</i>	290
17.4	<i>IP Multicast</i>	291
17.5	<i>IP Multicast Addresses</i>	291
17.6	<i>Mapping IP Multicast To Ethernet Multicast</i>	292
17.7	<i>Extending IP To Handle Multicasting</i>	293
17.8	<i>Internet Group Management Protocol</i>	294
17.9	<i>IGMP Implementation</i>	294
17.10	<i>Group Membership State Transitions</i>	295
17.11	<i>IGMP Message Format</i>	296
17.12	<i>Multicast Address Assignment</i>	297
17.13	<i>Propagating Routing Information</i>	297
17.14	<i>The MROUTED Program</i>	298
17.15	<i>Summary</i>	300