

经 典 原 版 书 库

网络互连

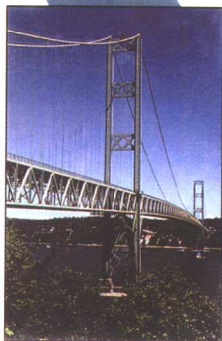
网桥·路由器·交换机和互连协议

(英文版·第2版)

Interconnections Second Edition

Bridges, Routers, Switches,
and Internetworking Protocols

Radia Perlman



ADDISON-WESLEY PROFESSIONAL COMPUTING SERIES

(美) Radia Perlman 著



机械工业出版社
China Machine Press



Addison-Wesley

经 典 原 版 书 库

(英文版·第2版)

网 络 互 连

网桥·路由器·交换机和互连协议

Interconnections

Bridges, Routers, Switches,
and Internetworking Protocols

(Second Edition)

(美) Radia Perlman 著



机械工业出版社
China Machine Press

English reprint edition copyright © 2002 by PEARSON EDUCATION NORTH ASIA LIMITED and CHINA MACHINE PRESS.

Interconnections (Second Edition) : Bridges, Routers, Switches, and Internetworking Protocols by Radia Perlman, Copyright © 2000.

All rights reserved.

Published by arrangement with Pearson Education, Inc.

本书英文影印版由美国Addison Wesley公司授权机械工业出版社在中国大陆境内独家出版发行, 未经出版者许可, 不得以任何方式抄袭、复制或节录本书中的任何部分。

版权所有, 侵权必究。

本书版权登记号: 图字: 01-2001-3756

图书在版编目(CIP)数据

网络互连: 网桥·路由器·交换机和互连协议(英文版·第2版)/(美)珀尔曼(Perlman, R.)著.-北京: 机械工业出版社, 2002.1
(经典原版书库)

ISBN 7-111-09506-5

I. 网… II. 珀… III. 计算机网络-英文 IV. TP393

中国版本图书馆CIP数据核字(2001)第081807号

机械工业出版社(北京市西城区百万庄大街22号 邮政编码 100037)

责任编辑: 华章

北京市密云县印刷厂印刷·新华书店北京发行所发行

2002年1月第1版第1次印刷

850mm×1168mm 1/32·17.625印张

印数: 0 001-2 000册

定价: 36.00元

凡购本书, 如有倒页、脱页、缺页, 由本社发行部调换

出版者的话

文艺复兴以降，源远流长的科学精神和逐步形成的学术规范，使西方国家在自然科学的各个领域中取得了垄断性的优势；也正是这样的传统，使美国在信息技术发展的六十多年间名家辈出、独领风骚。在商业化的进程中，美国的产业界与教育界越来越紧密地结合，计算机学科中的许多泰山北斗同时身处科研和教学的最前线，由此而产生的经典科学著作，不仅擘划了研究的范畴，还揭橥了学术的源变，既遵循学术规范，又自有学者个性，其价值并不会因年月的流逝而减退。

近年，在全球信息化大潮的推动下，我国的计算机产业发展迅猛，对专业人才的需求日益迫切。这对计算机教育界和出版界都既是机遇，也是挑战；而专业教材的建设在教育战略上显得举足轻重。在我国信息技术发展时间较短、从业人员较少的现状下，美国等发达国家在其计算机科学发展的几十年间积淀的经典教材仍有许多值得借鉴之处。因此，引进一批国外优秀计算机教材将对我国计算机教育事业的发展起积极的推动作用，也是与世界接轨、建设真正的世界一流大学的必由之路。

机械工业出版社华章图文信息有限公司较早意识到“出版要为教育服务”。自1998年始，华章公司就将工作重点放在了遴选、移译国外优秀教材上。经过几年的不懈努力，我们与Prentice Hall, Addison-Wesley, McGraw-Hill, Morgan Kaufmann等世界著名出版公司建立了良好的合作关系，从它们现有的数百种教材中甄选出Tanenbaum, Stroustrup, Kernighan, Jim Gray等大师名家的一批经典作品，以“计算机科学丛书”为总称出版，供读者学习、研究及度藏。大理石纹理的封面，也正体现了这套丛书的品位和格调。

“计算机科学丛书”的出版工作得到了国内外学者的鼎力襄助，国内的专家不仅提供了中肯的选题指导，还不辞劳苦地担任了翻译和审校的工作；而原书的作者也相当关注其作品在中国的传播，有的还专诚为其书的中译本作序。迄今，“计算机科学丛书”已经出版了近百个品种，这些书籍在读者中树立了良好的口碑，并被许多高校采用为正式教材和参考书籍，为进一步推广与发展打下了坚实的基础。

随着学科建设的初步完善和教材改革的逐渐深化，教育界对国外计算机教材的需求和应用都步入一个新的阶段。为此，华章公司将加大引进教材的力度，在“华章教育”的总规划之下出版三个系列的计算机教材：针对本科生的核心课程，剔抉外版菁华而成“国外经典教材”系列；对影印版的教材，则单独开辟出“经典原版书库”；定位在高级教程和专业参考的“计算机科学丛书”还将保持原来的风格，继续出版新的品种。为了保证这三套丛书的权威性，同时也为了更好地为学校和老师们的服务，华章公司聘请了中国科学院、北京大学、清华大学、国防科技大学、复旦大学、上海交通大学、南京大学、浙江大学、中国科技大学、哈尔滨工业大学、西安交通大学、中国人民大学、北京航空航天大学、北京邮电大学、中山大学、解放军理工大学、郑州大学、湖北工学院、中国国家信息安全测评认证中心等国内重点大学和科研机构在计算机的各个领域的著名学者组成“专家指导委员会”，为我们提供选题意见和出版监督。

“经典原版书库”是响应教育部提出的使用原版国外教材的号召，为国内高校的计算机教学度身订造的。在广泛地征求并听取丛书的“专家指导委员会”的意见后，我们最终选定了这30多种篇幅内容适度、讲解鞭辟入里的教材，其中的大部分已经被M.I.T.、Stanford、U.C. Berkley、C.M.U.等世界名牌大学采用。丛书不仅涵盖了程序设计、数据结构、操作系统、计算机体系结构、数据库、编译原理、软件工程、图形学、通信与网络、离散数学等国内大学计算机专业普遍开设的核心课程，而且各具特色——有的出自语言设计者之手、有的历三十年而不衰、有的已被全世界的几百所高校采用。在这些圆熟通博的名师大作的指引之下，读者必将在计算机科学的宫殿中由登堂而入室。

权威的作者、经典的教材、一流的译者、严格的审校、精细的编辑，这些因素使我们的图书有了质量的保证，但我们的目标是尽善尽美，而反馈的意见正是我们达到这一终极目标的重要帮助。教材的出版只是我们的后续服务的起点。华章公司欢迎老师和读者对我们的工作提出建议或给予指正，我们的联系方式如下：

电子邮件：hzedu@hzbook.com

联系电话：(010) 68995265

联系地址：北京市西城区百万庄南街1号

邮政编码：100037

专家指导委员会

(按姓氏笔画顺序)

尤晋元
石教英
张立昂
邵维忠
周克定
郑国梁
高传善
裘宗燕

王珊
吕建
李伟琴
陆丽娜
周傲英
施伯乐
梅宏
戴葵

冯博琴
孙玉芳
李师贤
陆鑫达
孟小峰
钟玉琢
程旭

史忠植
吴世忠
李建中
陈向群
岳丽华
唐世渭
程时端

史美林
吴时霖
杨冬青
周伯生
范明
袁崇义
谢希仁

Praise for *Interconnections, Second Edition*:

"Many networking books are little more than condensations of product literature and standard documents. Radia's book goes into the concepts and principles behind the technology, and in some cases, the history. This information will be very useful to students learning the technology, networking managers trying to improve their networks and product developers hoping to avoid past mistakes. If network architects and standards developers absorb this book perhaps the next generation of networking will be less complex and more effective. The book includes historical anecdotes and controversial recommendations. These make it entertaining as well as informative."

—Tony Lauck

"While there are other books on routing, I can think of no other that covers bridging at any depth. And the fact that it deals with both in one book is helpful to show the relationship, strengths, and weaknesses of each technology."

—Paul Koning
Zedia Corporation

"*Interconnections, Second Edition* does an excellent job of presenting a wide variety of solutions to a common problem . . . allowing the reader to appreciate the trade-offs and key issues."

—Peter Memishian

Praise for *Interconnections*:

"No other book ever written is better suited to build a strong understanding of networking concepts. Radia Perlman has been intimately involved with networking during its evolution, and her experience carries over into this entertaining book. Radia's book is perfect for beginners who wish to build a strong base of networking knowledge. It's also a great book for experienced network professionals who have a difficult time finding books that have new information."

—Tony Northrup
GTE Internetworking

"This is one of the best sources of information on the design of network routing protocols that exists. The book is indispensable if you are involved in a software development project that benefits from routing, at any level."

—Thomas H. Ptacek
Sonicity, Inc.

"A wonderful and richly detailed introduction to bridging and routing from one of the creators of the field."

—Craig Partridge, author of *Gigabit Networking*,
Addison-Wesley, 1993

"This book is the one-stop introduction to the field. It explains in detail the transparent and source routing bridges, the spanning tree protocol, and the various IP/OSI routing protocols: RIP, OSPF, BGP, ISIS. This is a must read for anyone interested in gaining a good grounding in internetworking protocols."

—Great Mahershi

Preface

Interconnections, Second Edition is about what goes on inside the boxes that move data around the Internet. These boxes are variously called bridges, routers, switches, and hubs. The book also describes the devices that connect to the network.

There is considerable confusion in this area. Most of the terminology is ill defined and is used in conflicting ways. The terminology and the specifications tend to be daunting. Some knowledge is spread among many different documents; much is unwritten folk wisdom. Adding to the confusion is dogma. Beliefs are accepted as truth, and questioning any of the dogma is often greeted with hostility. But good engineering demands that we understand what we're doing and why, keep an open mind, and learn from experience.

In *Interconnections, Second Edition*, instead of diving right into the details of one protocol, I first focus on the problems to be solved. I examine various solutions to each of these problems and discuss the engineering trade-offs involved. Then I look at a variety of solutions that have been deployed and compare the approaches. I give technical arguments for any opinions, and if you think I have missed any arguments I welcome email discussion. My email address is at the back of the book, which I hope you will find after having read the book cover to cover.

In the first edition, my intention was to help people understand the problems and the general types of solutions, assuming that they would read the specifications to get the details of specific protocols. But people used the book as a reference in addition to using it to understand the issues. So in this edition I have documented many more of the protocols in detail.

I believe that to understand something deeply you need to compare it to something else. The first edition was "minimalist" in that I always used only two examples: two types of bridges, bridges versus routers, connection-oriented versus connectionless network layer protocols, and two examples of connectionless protocols (CLNP and IP). In this edition I add a lot more examples, including ATM, IPv6, IPX, AppleTalk, and DECnet. I did this in part because these protocols exist, and it is hard to get information about them. But mostly I did it because the protocols embody interesting ideas that

should not be lost. When we design new protocols, we should learn from previous ideas, both good and bad. Also, it takes very little additional effort, after the problem is described generically, to describe several examples.

The Tao of network protocols: If all you see is IP, you see nothing.

—Greg Minshall

Roadmap to the Book

The first four chapters are not significantly different from their counterparts in the first edition, but the rest of the book has been largely rewritten. Chapters 1 through 4 cover general networking concepts, data link issues such as addressing and multiplexing, transparent bridges and the spanning tree algorithm, and source routing bridges. Chapter 5 is completely new and explains how the notion of a switch evolved into a rediscovery of the bridge. It also covers VLANs and fast Ethernet.

The remainder of the book concentrates on layer 3 (the network layer). Chapter 6 gives an overview of the network layer. Chapter 7 covers connection-oriented networks, including ATM and X.25. Chapter 8 discusses the issues in a generic connectionless network layer. Chapter 9 covers layer 3 addressing generically and gives a detailed comparison of IP, IPv6, CLNP, DECnet, AppleTalk, and IPX. Chapter 10 covers the information that should appear in a network layer header and contrasts the headers of several protocols.

Chapter 11 covers autoconfiguration and neighbor discovery, including protocols such as ARP and DHCP. Chapter 12 covers routing algorithms generically.

Chapter 13 discusses the problem of doing longest-prefix matching, which is required in order to forward IP packets quickly. Chapter 14 discusses the specifics of various routing protocols including RIP, IS-IS, OSPF, PNNI, NLSP, and BGP. Chapter 15 covers network layer multicast. Chapter 16 explains how to design a network that is invulnerable to sabotage, an idea whose time may come.

The final two chapters summarize the book, and I hope they will be mostly light and entertaining reading. Chapter 17 probes the mystery of what, if anything, distinguishes a router from a bridge. Chapter 18 attempts to capture folk wisdom about how to design a protocol.

Finally, there is an extensive glossary. I try to define terms when I first use them, but if I ever fail to do that, you will probably find them in the glossary.

Acknowledgments

Writing this section is scary because I am afraid I will leave people out. I'd like to thank the people who reviewed all or part of the book: Peter Memishian, Paul Koning, Tony Lauck,

Craig Partridge, Dan Pitt, Brian Kernighan, Paul Bottorff, Joel Halpern, Charlie Kaufman, Mike Speciner, Andy Tanenbaum, Phil Rosenzweig, Dan Senie, William Welch, Craig Labovitz, Chase Bailey, George Varghese, and Suchi Raman. Other people who have been helpful by answering questions are Ariel Hendel, Rich Kubota, Stuart Cheshire, Tom Maufer, Steve Deering, and John Moy. The first time I sent an email question in the middle of the night (when I did most of my work on this book) to Craig Partridge, the co-series editor for this book, the beep indicating incoming mail happened so immediately that I assumed it was an automatic mail responder informing me he was on vacation. But it was an answer to my question. I assume he doesn't have an automatic mail responder so clever that it can answer technical questions, so I thank him for being so prompt and available. Brian Kernighan, the other series editor, also had detailed and helpful comments on the entire book.

The people at Addison-Wesley have been amazingly patient with me for the many years in which I've been working on this edition. I'm not sure they had any alternative besides patience, but it was nice that they believed I'd finish even when I wasn't so sure. So thank you to Mary Hart, Karen Gettman, Jacquelyn Doucette, and Jason Jones. And I'd also like to thank my copy editor, Betsy Hardinger. She of all people will have read every word of the book, while maintaining the concentration to note inconsistencies and ways of removing excess words here and there. I know it's her job, but I'm still impressed.

Mike Speciner helped me figure out the mysteries of Framemaker. Ray Perlner made sure that I maintained some humor in the book and watched over my shoulder while I typed the last chapter to see that I had enough funny bad real-life protocols. Dawn Perlner has been terrifically supportive, convincing her friends and even strangers in bookstores to buy my books. She used to be my child. Now she's a wonderful friend.

Contents

Preface

Roadmap to the Book
Acknowledgments

Chapter 1 Essential Networking Concepts 1

- 1.1 Layers 1
- 1.2 Service Models 7
- 1.3 Important Properties of a Network 9
- 1.4 Reliable Data Transfer Protocols 12

Chapter 2 Data Link Layer Issues 19

- 2.1 Generic LANs 19
- 2.2 IEEE 802 LANs 22
- 2.3 Names, Addresses, Routes 23
- 2.4 LAN Addresses 25
- 2.5 Multicast versus Unicast Addresses 27
- 2.6 The Broadcast Address 28
- 2.7 Multiplexing Field 29
- 2.8 Bit Order 32
- 2.9 Logical Link Control 34
- 2.10 Issues in 802.3 35
- 2.11 Issues in 802.5 38
- 2.12 Packet Bursts 40
- 2.13 Reasons for Bridges 40
- 2.14 Point-to-Point Links 41

Chapter 3 Transparent Bridges 45

- 3.1 The No-Frills Bridge 45
- 3.2 The Learning Bridge 47
- 3.3 Spanning Tree Algorithm 58
- 3.4 Spanning Tree Algorithm Refinements 63

3.5	Bridge Message Formats	77
3.6	Other Bridge Issues	79
3.7	Remote Bridges	87
Chapter 4 Source Routing Bridges 95		
4.1	Pure Source Routing	96
4.2	SR-TB Bridges	104
4.3	SRT Bridges	108
4.4	End-system Algorithms	109
4.5	Source Routing versus Transparent Bridging	116
4.6	Ideas for Improving Source Route Bridging	119
Chapter 5 Hubs, Switches, Virtual LANs, and Fast Ethernet 127		
5.1	Hubs	127
5.2	Faster LANs	132
5.3	Virtual LANs (VLANs)	134
Chapter 6 Network Interface: Service Models 145		
6.1	What Is the Network Layer?	145
6.2	Network Service Types	146
Chapter 7 Connection-oriented Nets: X.25 and ATM 153		
7.1	Generic Connection-oriented Network	153
7.2	X.25: Reliable Connection-oriented Service	155
7.3	Implementing X.25 Inside the Net	167
7.4	Asynchronous Transfer Mode	168
Chapter 8 Generic Connectionless Service 179		
8.1	Data Transfer	179
8.2	Addresses	180
8.3	Hop Count	180
8.4	Service Class Information	181
8.5	Network Feedback	183
8.6	Fragmentation and Reassembly	183
8.7	Maximum Packet Size Discovery	185
Chapter 9 Network Layer Addresses 189		
9.1	Hierarchical Addresses with Fixed Boundaries	190
9.2	Hierarchical Addresses with Flexible Boundaries	192
9.3	Owning versus Renting Addresses	193
9.4	Types of Addresses	194
9.5	IP	194
9.6	IPX	199

- 9.7 IPX+ 203
- 9.8 IPv6 204
- 9.9 CLNP Network Layer Addresses 211
- 9.10 AppleTalk Network Layer Addresses 215
- 9.11 DECnet Phases III and IV 216
- 9.12 NAT/NAPT 218

Chapter 10 Connectionless Data Packet Formats 221

- 10.1 Pieces of a Connectionless Network Layer 221
- 10.2 Data Packets 222
- 10.3 Summary of Packet Formats for Easy Reference 222
- 10.4 Technical Issues and Comparisons in Data Packet Formats 230
- 10.5 Source Routing 250
- 10.6 The Great IPX Frame Format Mystery 253
- 10.7 Error Reports and Other Network Feedback to the Endnode 256

Chapter 11 Neighbor Greeting and Autoconfiguration 265

- 11.1 Endnodes Attached via Point-to-Point Links 266
- 11.2 Endnodes Attached via LANs 267
- 11.3 Endnodes Attached via Nonbroadcast Multiaccess Media 284
- 11.4 Finding Things 293

Chapter 12 Routing Algorithm Concepts 299

- 12.1 Distance Vector Routing 299
- 12.2 Link State Routing 307
- 12.3 Comparison of Link State and Distance Vector Routing 320
- 12.4 Load Splitting 325
- 12.5 Link Costs 326
- 12.6 Migrating Routing Algorithms 328
- 12.7 LANs 331
- 12.8 Types of Service 334
- 12.9 Partition Repair: Level 1 Subnetwork Partition 341

Chapter 13 Fast Packet Forwarding 347

- 13.1 Using an Additional Header 347
- 13.2 Address Prefix Matching 348
- 13.3 Longest Prefix Match with Trie 350
- 13.4 Binary Search 361

Chapter 14 Specific Routing Protocols 367

- 14.1 A Brief History of Intradomain Routing Protocols 367
- 14.2 RIP 369
- 14.3 RTMP, IPX-RIP, and DECnet 372

- 14.4 IS-IS, OSPF, NLSP, and PNNI 373
- 14.5 Interdomain Routing Protocols 425

Chapter 15 WAN Multicast 449

- 15.1 Introduction 449
- 15.2 Multicast in IP 453

Chapter 16 Sabotage-proof Routing 479

- 16.1 The Problem 480
- 16.2 All You Need to Know about Cryptography 480
- 16.3 Overview of the Approach 482
- 16.4 Detailed Description of the Approach 484
- 16.5 Summary 491
- 16.6 For Further Reading 491

Chapter 17 To Route, Bridge, or Switch: Is That the Question? 493

- 17.1 Switches 493
- 17.2 Bridges versus Routers 493
- 17.3 Extensions to Bridges 496
- 17.4 Extensions to Routers 498

Chapter 18 Protocol Design Folklore 503

- 18.1 Simplicity versus Flexibility versus Optimality 503
- 18.2 Knowing the Problem You're Trying to Solve 505
- 18.3 Overhead and Scaling 506
- 18.4 Operation Above Capacity 506
- 18.5 Compact IDs versus Object Identifiers 507
- 18.6 Optimizing for the Most Common or Important Case 507
- 18.7 Forward Compatibility 508
- 18.8 Migration: Routing Algorithms and Addressing 513
- 18.9 Parameters 514
- 18.10 Making Multiprotocol Operation Possible 517
- 18.11 Running over Layer 3 versus Layer 2 518
- 18.12 Robustness 519
- 18.13 Determinism versus Stability 520
- 18.14 Performance for Correctness 521
- 18.15 In Closing 521

Glossary 523

Index 533

Chapter 1

Essential Networking Concepts

This chapter introduces concepts that are essential to understanding the specific subfield of computer networking that includes bridges and routers. It covers the OSI reference model, including layering and service models, because this model is a useful basis for some vocabulary. It also discusses various dimensions along which network designs can differ, such as scope, scalability, robustness, and autoconfigurability. Chapter 1 also describes the typical techniques involved in providing reliable two-party communication because some of the techniques used by routers can interact with techniques used by other layers.

1.1 Layers

Understanding, designing, and building a computer network would be too difficult a task unless the problem were partitioned into smaller subtasks, traditionally by dividing the problem into layers. The idea behind layering is that each layer is responsible for providing a service to the layer above by using the services of the layer below.

Each layer communicates with its *peer* layer in another node through the use of a *protocol*. This communication is accomplished through direct communication with the layer below. The communication between layer n and layer $n-1$ is known as an *interface*.

The OSI (Open Systems Interconnection) Reference Model defines seven layers, as shown in Figure 1.1. There is nothing magic about the number seven or the functionality in the layers. The reference model was designed before the protocols themselves, and then committees were set up to design each of the layers. Many of the layers were subsequently subdivided into further layers. The distinction between the layers is not always clear. Bridges and routers are a good example of a case in which people should rightfully be confused about which layers are which. But semantic arguments about layers are not

very productive. Instead, the layering should be viewed as a useful framework for discussion and not as a bible.

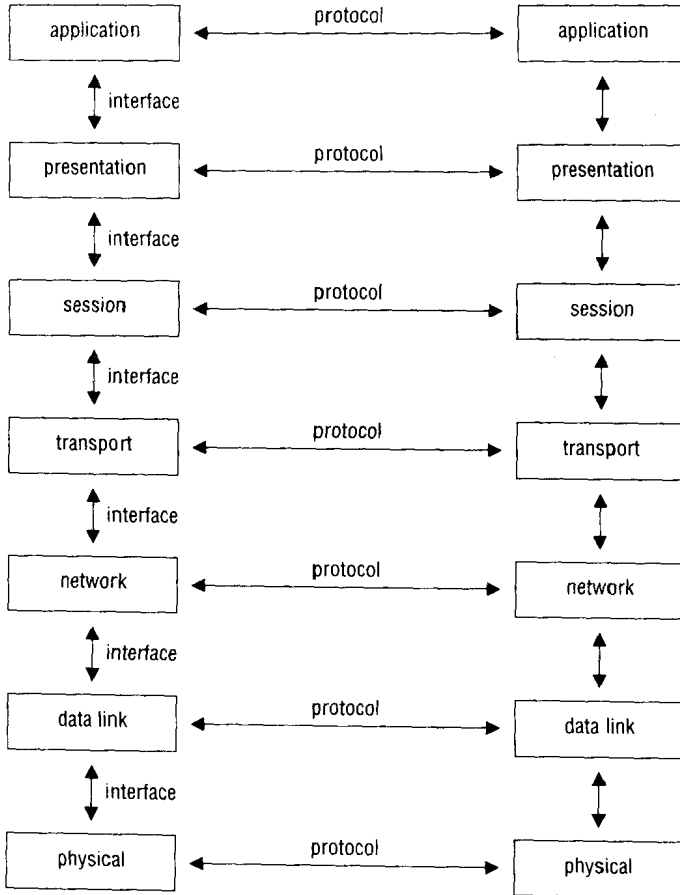


Figure 1.1 The OSI Reference Model

Layers defined by ISO:

1. *Physical layer:* The physical layer transmits bits of information across a link. It deals with such problems as size and shape of connectors, assignment of functions to pins, conversion of bits to electrical signals, and bit-level synchronization. It is usual for several different types of physical layers to exist within a network and even for multiple different types of physical layers to exist within a node, because each technology requires its own physical layer.