计 算 机 科 学 丛 书

Second Edition

Computer Networks

A Systems Approach

计算机网络

（英文版·第2版）

(美) Larry L. Peterson
Bruce S. Davie 著

机械工业出版社 China Machine Press
MORGAN KAUFMANN

本书版权登记号：图字：01-1999-2869

图书在版编目 (CIP) 数据

计算机网络：英文/（美）彼得森（Peterson, L. L.），（美）戴维（Davie, B. S.)著. -北京：机械工业出版社，2000.1

（计算机科学丛书）

书名原文：Computer Networks: Systems Approach, Second Edition

ISBN 7-111-07526-9

Ⅰ.计... Ⅱ.①彼... ②戴... Ⅲ.计算机网络-英文 Ⅳ.TP393

中国版本图书馆CIP数据核字（1999）第54642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：周 桦

北京牛山世兴印刷厂印刷 · 新华书店北京发行所发行

2000年1月第1版 · 2001年1月第2次印刷

787mm × 1092 mm 1/16 · 48.25印张

印数：2 001-3 000册

定价：65.00元

FOREWORD

David Clark
Massachusetts Institute of Technology

I am pleased to report that this great book has gotten better. The philosophy of the book remains unchanged: to be timely but timeless in the material it presents. The timeless component is still there, better than ever. The book is an excellent introduction to the core concepts and fundamental principles that will remain useful even when the standards change and the technology evolves. If you want to understand how networks work, not just how the packet headers are formatted, this is the book to read.

To keep a book on the Internet timely takes a lot of work because the Internet evolves so rapidly. It is hard to remember that much less than a decade ago, nobody had heard of a URL, and only a few years ago, Internet telephony and Internet radio were idle speculations of the research community. The authors have put a lot of effort into updating the material to make it as current as possible—there are new sections on a number of important topics, ranging from quality of service to security—so the book can be used as a reference volume as well as a textbook for learning basic principles.

Other changes have the goal of making the examples easier to follow. The code fragments are no longer based on a specific operating system, but have been rewritten to assume a more generic execution environment. The objective is to reduce the effort that the student must put in before getting to the meat of the material, so that learning can proceed as rapidly as possible.

Networking is a fast-moving field. Everything, including a book, has to run fast just to keep up. And the book will help you keep up.

FOREWORD TO THE FIRST EDITION

David Clark
Massachusetts Institute of Technology

The term *spaghetti code* is universally understood as an insult. All good computer scientists worship the god of modularity, since modularity brings many benefits, including the all-powerful benefit of not having to understand all parts of a problem at the same time in order to solve it. Modularity thus plays a role in presenting ideas in a book, as well as in writing code. If a book's material is organized effectively—modularly—the reader can start at the beginning and actually make it to the end.

The field of network protocols is perhaps unique in that the "proper" modularity has been handed down to us in the form of an international standard: the seven-layer reference model of network protocols from the ISO. This model, which reflects a layered approach to modularity, is almost universally used as a starting point for discussions of protocol organization, whether the design in question conforms to the model or deviates from it.

It seems obvious to organize a networking book around this layered model. However, there is a peril to doing so, because the OSI model is not really successful at organizing the core concepts of networking. Such basic requirements as reliability, flow control, or security can be addressed at most, if not all, of the OSI layers. This fact has led to great confusion in trying to understand the reference model. At times it even requires a suspension of disbelief. Indeed, a book organized strictly according to a layered model has some of the attributes of spaghetti code.

Which brings us to this book. Peterson and Davie follow the traditional layered model, but they do not pretend that this model actually helps in the understanding of the big issues in networking. Instead, the authors organize discussion of fundamental concepts in a way that is independent of layering. Thus, after reading the book, readers will understand flow control, congestion control, reliability enhancement, data representation, and synchronization, and will separately understand the implications of addressing these issues in one or another of the traditional layers.

This is a timely book. It looks at the important protocols in use today—especially the Internet protocols. Peterson and Davie have a long involvement in and much

experience with the Internet. Thus their book reflects not just the theoretical issues in protocol design, but the real factors that matter in practice. The book looks at some of the protocols that are just emerging now, so the reader can be assured of an up-to-date perspective. But most importantly, the discussion of basic issues is presented in a way that derives from the fundamental nature of the problem, not the constraints of the layered reference model or the details of today's protocols. In this regard, what this book presents is both timely and timeless. The combination of real-world relevance, current examples, and careful explanation of fundamentals makes this book unique.

PREFACE

When the first edition of this book was published in 1996 (nearly a hundred years ago in Internet time), it was a novelty to be able to order merchandise on the Internet, and a company that advertised its domain name was considered cutting edge. Today, Internet commerce is a fact of life, and it's not uncommon to hear of companies that include ".com" in their name—their legal corporate name, not their Internet name. Stockbrokers refer to "the .com stocks," and the valuation of a coffee retailer soars when it announces plans to develop an "Internet strategy." To point out that things change quickly on the Internet is a bit like pointing out that the sky is blue.

Despite these changes the question we asked in the first edition is just as valid today: What are the underlying concepts and technologies that make the Internet work? The answer is that much of the TCP/IP architecture continues to function just as was envisioned by its creators nearly 30 years ago. This isn't to say that the Internet architecture is uninteresting; quite the contrary. Understanding the design principles that underly an architecture that has not only survived but fostered the kind of growth and change that the Internet has seen over the past three decades is precisely the right place to start. Like the first edition, the second edition makes the "why" of the Internet architecture its cornerstone.

Audience

Our intent is that the book should serve as the text for a comprehensive networking class, at either the graduate or upper-division undergraduate level. We also believe that the book's focus on core concepts should be appealing to industry professionals who are retraining for network-related assignments, as well as current network practitioners who want to understand the "whys" behind the protocols they work with every day and to see the big picture of networking.

It is our experience that both students and professionals learning about networks for the first time often have the impression that network protocols are some sort of edict handed down from on high, and that their job is to learn as many TLAs (Three-Letter Acronyms) as possible. In fact, protocols are the building blocks of a complex system developed through the application of engineering design principles.

Moreover, they are constantly being refined, extended, and replaced based on real-world experience. With this in mind, our goal with this book is to do more than survey the protocols in use today. Instead, we explain the underlying principles of sound network design. We feel that this grasp of underlying principles is the best tool for handling the rate of change in the networking field.

Changes in the Second Edition

Our focus on the underlying principles of networking might sound like an excuse for not making changes in the second edition. On the contrary, the second edition tracks many of the most important recent advances in networking with the addition of a significant amount of new material. We have also deleted, reorganized, and changed the focus of existing material to reflect changes that have taken place over the past four years. In addition, we have responded to extensive feedback received from those who used the first edition, and even from some who did not use it. Changes in this edition include

- a new chapter on network security, including coverage of PGP, IPSEC, secure sockets, and firewalls
- a new chapter on network applications, which includes sections on SMTP, HTTP, SNMP, and DNS, as well as the RTP protocol used to build real-time multimedia applications
- comprehensively updated material on quality of service and congestion control, including the latest on Differentiated Services and RSVP
- expanded discussion of routing protocols, including OSPF and BGP
- a new section on wireless technology, including spread spectrum techniques and the emerging 802.11 standard
- expanded discussion of audio and video compression, including MPEG and MP3
- increased coverage of ATM, including LAN emulation (LANE)
- new material on building virtual private networks (VPNs) on top of the public Internet
- coverage of high-speed networking throughout the book rather than confining it to a single chapter

Two other significant changes affect the book as a whole. First, in response to feedback from instructors, we have removed the *x*-kernel as a pedagogical tool

for understanding protocol implementations and replaced it with operating-system-independent C code. The details of this change are discussed below.

We have also significantly increased the number and quality of exercises at the end of each chapter. This work was spearheaded by a dedicated instructor who has taught from the first edition of our book since its inception.

Approach

For an area that's as dynamic and changing as computer networks, the most important thing a textbook can offer is perspective—to distinguish between what's important and what's not, and between what's lasting and what's superficial. Based on our experience over the past 20 years doing research that has led to new networking technology, teaching undergraduate and graduate students about the latest trends in networking, and delivering advanced networking products to market, we have developed a perspective—which we call the *systems approach*—that forms the soul of this book. The systems approach has several implications:

- Rather than accept existing artifacts as gospel, we start with first principles and walk you through the thought process that led to today's networks. This allows us to explain *why* networks look like they do. It is our experience that once you understand the underlying concepts, any new protocol that you are confronted with will be relatively easy to digest.
- Although the material is loosely organized around the traditional network layers, starting at the bottom and moving up the protocol stack, we do not adopt a rigidly layerist approach. Many topics—congestion control and security are good examples—have implications up and down the hierarchy, and so we discuss them outside the traditional layered model. In short, we believe layering makes a good servant but a poor master; it's more often useful to take an end-to-end perspective.
- Rather than explain how protocols work in the abstract, we use the most important protocols in use today—many of them from the TCP/IP Internet—to illustrate how networks work in practice. This allows us to include real-world experiences in the discussion.
- Although at the lowest levels networks are constructed from commodity hardware that can be bought from computer vendors and communication services that can be leased from the phone company, it is the software that allows networks to provide new services and adapt quickly to changing circumstances. It is for this reason that we emphasize how network software is implemented, rather than stopping with a description of the abstract algorithms involved.

- Networks are constructed from many building-block pieces, and while it is necessary to be able to abstract away uninteresting elements when solving a particular problem, it is essential to understand how all the pieces fit together to form a functioning network. We therefore spend considerable time explaining the overall end-to-end behavior of networks, not just the individual components, so that it is possible to understand how a complete network operates, all the way from the application to the hardware.
- The systems approach implies doing experimental performance studies, and then using the data you gather both to quantitatively analyze various design options and to guide you in optimizing the implementation. This emphasis on empirical analysis pervades the book.
- Networks are like other computer systems—for example, operating systems, processor architectures, distributed and parallel systems, and so on. They are all large and complex. To help manage this complexity, system builders often draw on a collection of design principles. We highlight these design principles as they are introduced throughout the book, illustrated, of course, with examples from computer networks.

Software

As noted above, software and its implementation play an important role in a systems approach to understanding computer networks. The first edition used the *x*-kernel—a software framework for implementing network protocols—as a pedalogical tool, but since there was a startup cost in using the *x*-kernel, many people elected to not use it. For this reason, we have removed the *x*-kernel from the second edition. Code segments are still used throughout the book to illustrate how you might implement certain protocols and algorithms—in fact, we have added additional code—but these segments are given as operating-system-independent C code rather than *x*-kernel protocols. For users who want to continue using the *x*-kernel, the material removed from the book is still available online at http://www.cs.princeton.edu/xkernel.

Because we view network software as an essential component of networking, the second edition now includes an example of a simple application socket program in Chapter 1. Programming assignments based on Unix sockets are available online (see below).

Pedagogy and Features

The second edition retains several features that we encourage you to take advantage of:

- *Problem statements.* At the start of each chapter, we describe a problem that identifies the next set of issues that must be addressed in the design of a

network. This statement introduces and motivates the issues to be explored in the chapter.

- *Shaded sidebars.* Throughout the text, shaded sidebars elaborate on the topic being discussed or introduce a related advanced topic. In many cases, these sidebars relate real-world anecdotes about networking.
- *Highlighted paragraphs.* These paragraphs summarize an important nugget of information that we want you to take away from the discussion, such as a widely applicable system design principle.
- *Real protocols.* Even though the book's focus is on core concepts rather than existing protocol specifications, real protocols are used to illustrate most of the important ideas. As a result, the book can be used as a source of reference for many protocols. To help you find the descriptions of the protocols, each applicable section heading parenthetically identifies the protocols described in that section. For example, Section 5.2, which describes the principles of reliable end-to-end protocols, provides a detailed description of TCP, the canonical example of such a protocol.
- *Open issues.* We conclude the main body of each chapter with an important issue that is currently being debated in the research community, the commercial world, or society as a whole. We have found that discussing these issues helps to make the subject of networking more relevant and exciting.
- *Recommended reading.* These highly selective lists appear at the end of each chapter. Each list generally contains the seminal papers on the topics just discussed. We strongly recommend that advanced readers (e.g., graduate students) study the papers in this reading list to supplement the material covered in the chapter.

Road Map and Course Use

The book is organized as follows:

- Chapter 1 introduces the set of core ideas that are used throughout the rest of the text. In particular, it discusses what goes into a network architecture, and it defines the quantitative performance metrics that often drive network design.
- Chapter 2 surveys a wide range of low-level network technologies, ranging from Ethernet to token ring to wireless. It also describes many of the issues that all data link protocols must address, including encoding, framing, and error detection.

- Chapter 3 introduces the basic models of switched networks (datagrams versus virtual circuits) and describes one prevalent switching technology (ATM) in some detail. It also discusses the design of hardware-based switches.
- Chapter 4 introduces internetworking and describes the key elements of the Internet Protocol (IP). A central question addressed in this chapter is how networks that scale to the size of the Internet are able to route packets.
- Chapter 5 moves up to the transport level, describing both the Internet's Transmission Control Protocol (TCP) and Remote Procedure Call (RPC) used to build client-server applications in detail.
- Chapter 6 discusses congestion control and resource allocation. The issues in this chapter cut across both the network level (Chapters 3 and 4) and the transport level (Chapter 5). Of particular note, this chapter describes how congestion control works in TCP, and it introduces the mechanisms used by both the Internet and ATM to provide quality of service.
- Chapter 7 considers the data sent through a network. This includes both the problems of presentation formatting and data compression. The discussion of compression includes explanations of how MPEG video compression and MP3 audio compression work.
- Chapter 8 discusses network security, ranging from an overview of cryptography protocols (DES, RSA, MD5), to protocols for security services (authentication, digital signature, message integrity), to complete security systems (privacy enchanced email, IPSEC). The chapter also discusses pragmatic issues like firewalls.
- Chapter 9 describes a representative sample of network applications, including both traditional applications like email and the Web, and multimedia applications that use the Real-time Transport Protocol (RTP).

For an undergraduate course, extra class time will most likely be needed to help students digest the introductory material in the first chapter, probably at the expense of the more advanced topics covered in Chapters 6 through 8. Chapter 9 then returns to the popular topic of network applications. In contrast, the instructor for a graduate course should be able to cover the first chapter in only a lecture or two—with students studying the material more carefully on their own—thereby freeing up additional class time to cover the last four chapters in depth. Both graduate and undergraduate classes will want to cover the core material contained in the middle four chapters (Chapters 2–5), although an undergraduate class might choose to skim the more advanced sections (e.g., Sections 2.2, 2.9, 3.4, and 4.4).

For those of you using the book in self-study, we believe that the topics we have selected cover the core of computer networking, and so we recommend that the book be read sequentially, from front to back. In addition, we have included a liberal supply of references to help you locate supplementary material that is relevant to your specific areas of interest.

The book takes a unique approach to the topic of congestion control by pulling all topics related to congestion control and resource allocation together in a single place—Chapter 6. We do this because the problem of congestion control cannot be solved at any one level, and we want you to consider the various design options at the same time. (This is consistent with our view that strict layering often obscures important design trade-offs.) A more traditional treatment of congestion control is possible, however, by studying Section 6.2 in the context of Chapter 3 and Section 6.3 in the context of Chapter 5.

Exercises

Significant effort has gone into improving the exercises in the second edition. Specifically, we have enlisted Peter Dordal, who teaches the networking class at Loyola University of Chicago, to revamp the exercises. Peter has greatly increased the number of problems (from 196 to 362) and, based on class testing, dramatically improved their quality. The current set of exercises are of several different styles:

- Analytical exercises that ask the student to do simple algebraic calculations that demonstrate their understanding of fundamental relationships
- Design questions that ask the student to propose and evaluate protocols for various circumstances
- Hands-on questions that ask the student to write a few lines of code to test an idea or to experiment with an existing network utility
- Library research questions that ask the student to learn more about a particular topic

Also, as described in more detail below, *x*-kernel and socket-based programming assignments are available online.

Supplemental Materials and Online Resources

To assist instructors, we have prepared an instructor's manual that contains solutions to selected exercises. The manual is available from the publisher. Additional support materials, including lecture slides, figures from the text, socket-based programming assignments, and sample exams are available through the Morgan Kaufmann Web site at http://www.mkp.com. We suggest that you visit the page for this book every few

weeks, as we will be adding support materials and establishing links to networking-related sites on a regular basis.

Acknowledgments

This book would not have been possible without the help of many people. We would like to thank them for their efforts in improving the end result. Before we do so, however, we should mention that we have done our best to correct the mistakes that the reviewers have pointed out and to accurately describe the protocols and mechanisms that our colleagues have explained to us. We alone are responsible for any remaining errors. If you should find any of these, please send email to our publisher, Morgan Kaufmann, at netbugs@mkp.com, and we will endeavor to correct them in future printings of this book.

First, we would like to thank the many people who reviewed drafts of various chapters. The list is long and includes Ken Calvert of the University of Kentucky, Douglas Jacobson of Iowa State University, Michel Barbeau of the University of Sherbrooke, Ken Klingenstein of the University of Colorado–Boulder, Robert Strader of Stephen F. Austin State University, Lee Hollaar of the University of Utah, James Ten-Eyck of Marist College, Walt Will of Luther College, David Hutchison of Lancaster University, Ivan Marsic of Rutgers University, Lee Leitner of Infocus and Nova Southeastern University, consultant Michael Cochran, Edward Balassanian of BeComm Corporation, Matt Bishop of UC–Davis, and Steve Casner of Cisco.

Second, several members of the Network Systems Group at Princeton and the University of Arizona contributed ideas, examples, corrections, data, and code to this book. In addition to those mentioned in the first edition, George Tzanetakis explained MP3 to us, and Chad Mynhier helped update the bibliography. As before, we want to thank the Defense Advanced Research Projects Agency and the National Science Foundation for supporting our networking research over the past several years.

Third, we would like to thank our series editor, David Clark, as well as all the people at Morgan Kaufmann who helped shepherd us through the book-writing process. A special thanks is due to our sponsoring editor, Jennifer Mann; her assistant, Karyn Johnson; and our production editor, Cheri Palmer. The whole crew at MKP has been a delight to work with.

Finally, we wish to thank our wives, Lynn Peterson and Jody Davie, who continue to support us in our obsession to explain how networks really work.

CONTENTS