

经 典 原 版 书 库

# 计算机网络 系统方法

(美) Larry L. Peterson Bruce S. Davie 著  
普林斯顿大学 Cisco公司

(英文版·第5版)

Larry L. Peterson and Bruce S. Davie

Fifth Edition

# Computer Networks

a systems approach



机械工业出版社  
China Machine Press

MK

经

典

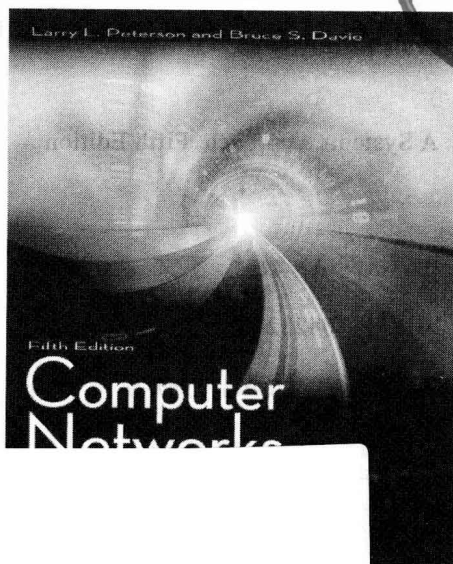
库

# 计算机网络

## 系统方法

(英文版·第5版)

*Computer Networks*  
A Systems Approach (Fifth Edition)



(美) Larry L. Peterson Bruce S. Davie 著  
普林斯顿大学 Cisco公司



机械工业出版社  
China Machine Press

Larry L. Peterson and Bruce S. Davie: Computer Networks: A Systems Approach, Fifth Edition (ISBN 978-0-12-385059-1).

Original English language edition copyright © 2012 by Elsevier Inc. All rights reserved.

Authorized English language reprint edition published by the Proprietor.

Copyright © 2012 by Elsevier (Singapore) Pte Ltd.

Printed in China by China Machine Press under special arrangement with Elsevier (Singapore) Pte Ltd. This edition is authorized for sale in China only, excluding Hong Kong SAR and Taiwan.

Unauthorized export of this edition is a violation of the Copyright Act. Violation of this Law is subject to Civil and Criminal Penalties.

本书英文影印版由Elsevier (Singapore) Pte Ltd. 授权机械工业出版社在中国大陆境内独家发行。本版仅限在中国境内（不包括香港特别行政区及台湾地区）出版及标价销售。未经许可之出口，视为违反著作权法，将受法律之制裁。

封底无防伪标均为盗版

版权所有，侵权必究

本书法律顾问 北京市展达律师事务所

本书版权登记号：图字：01-2012-0223

图书在版编目（CIP）数据

计算机网络：系统方法（英文版·第5版）/（美）彼得森（Peterson, L. L.），（美）戴维（Davie, B. S.）著.  
—北京：机械工业出版社，2012.4

（经典原版书库）

书名原文：Computer Networks: A Systems Approach, Fifth Edition

ISBN 978-7-111-37720-7

I. 计… II. ①彼… ②戴… III. 计算机网络—教材—英文 IV. TP393

中国版本图书馆CIP数据核字（2012）第042321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：迟振春

北京瑞德印刷有限公司印刷

2012年4月第1版第1次印刷

186mm×240mm·57.25印张

标准书号：ISBN 978-7-111-37720-7

定价：139.00元

凡购本书，如有缺页、倒页、脱页，由本社发行部调换

客服热线：(010) 88378991；88361066

购书热线：(010) 68326294；88379649；68995259

投稿热线：(010) 88379604

读者信箱：hzsj@hzbook.com

\* *Experiments were conducted in the Gas Turbine Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139.*

机械工业出版社华章公司较早意识到“出版要为教育服务”。自1998年开始,我们就将工作重点放在了遴选、移译国外优秀教材上。经过多年的不懈努力,我们与Pearson, McGraw-Hill, Elsevier, MIT, John Wiley & Sons, Cengage等世界著名出版公司建立了良好的合作关系,从他们现有的数百种教材中甄选出Andrew S. Tanenbaum, Bjarne Stroustrup, Brain W. Kernighan, Dennis Ritchie, Jim Gray, Alfred V. Aho, John E. Hopcroft, Jeffrey D. Ullman, Abraham Silberschatz, William Stallings, Donald E. Knuth, John L. Hennessy, Larry L. Peterson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、研究及珍藏。大理石纹理的封面,也正体现了这套丛书的品位和格调。

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。随着计算机科学与技术专业学科建设的不断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将步入一个新的阶段，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。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，我们的联系方法如下：

邮政编码: 100037



华章科技图书出版中心

## **In Praise of Computer Networks: A Systems Approach Fifth Edition**

*I have known and used this book for years and I always found it very valuable as a textbook for teaching computer networks as well as a reference book for networking professionals. This Fifth Edition maintains the core value of former editions and brings the clarity of explanation of network protocols in the introduction of the most up-to-date techniques, technologies and requirements of networking. Beyond describing the details of past and current networks, this book successfully motivates the curiosity, and hopefully new research, for the networks of the future.*

**Stefano Basagni**

Northeastern University

*Peterson and Davie have written an outstanding book for the computer networking world. It is a well-organized book that features a very helpful “big picture” systems approach. This book is a must have!*

**Yonshik Choi**

Illinois Institute of Technology

*The Fifth Edition of Computer Networks: A Systems Approach is well-suited for the serious student of computer networks, though it remains accessible to the more casual reader as well. The authors' enthusiasm for their subject is evident throughout; they have a thorough and current grasp of the interesting problems of the field. They explain not only how various protocols work, but also why they work the way they do, and even why certain protocols are the important and interesting ones. The book is also filled with little touches of historical background, from the main text to the “Where Are They Now” sidebars to the papers described in each chapter's “Further Reading” section—these give the reader a perspective on how things came to be the way they are. All in all, this book provides a lucid and literate introduction to networking.*

**Peter Dordal**

Loyola University Chicago

*I have used Computer Networks: A Systems Approach for over five years in an introductory course on communications networks aimed at upper-level undergraduates and first-year Masters students. I have gone through several editions and over the years the book has kept what from the beginning*

*had been its main strength, namely, that it not only describes the 'how,' but also the 'why' and equally important, the 'why not' of things. It is a book that builds engineering intuition, and in this day and age of fast-paced technology changes, this is critical to develop a student's ability to make informed decisions on how to design or select the next generation systems.*

**Roch Guerin**

University of Pennsylvania

*This book is an outstanding introduction to computer networks that is clear, comprehensive, and chock-full of examples. Peterson and Davie have a gift for boiling networking down to simple and manageable concepts without compromising technical rigor. Computer Networks: A Systems Approach strikes an excellent balance between the principles underlying network architecture design and the applications built on top. It should prove invaluable to students and teachers of advanced undergraduate and graduate networking courses.*

**Arvind Krishnamurthy**

University of Washington

*Computer Networks: A Systems Approach has always been one of the best resources available to gain an in-depth understanding of computer networks. The latest edition covers recent developments in the field. Starting with an overview in Chapter 1, the authors systematically explain the basic building blocks of networks. Both hardware and software concepts are presented. The material is capped with a final chapter on applications, which brings all the concepts together. Optional advanced topics are placed in a separate chapter. The textbook also contains a set of exercises of varying difficulty at the end of each chapter which ensure that the students have mastered the material presented.*

**Karkal Prabhu**

Drexel University

*Peterson and Davie provide a detailed yet clear description of the Internet protocols at all layers. Students will find many study aids that will help them gain a full understanding of the technology that is transforming our society. The book gets better with each edition.*

**Jean Walrand**

University of California at Berkeley

# Foreword

Once again, this now-classic textbook has been revised to keep it up-to-date with our evolving field. While the Internet and its protocols now dominate networking everywhere, we see continued evolution in the technology used to support the Internet, with switching at “layer 2” providing rich functionality and powerful tools for network management. The previous edition dealt with switching and routing in two chapters, but a presentation based on layers is not always the best way to convey the essentials of the material, since what we call switching and routing actually play similar and complementary roles. This edition of the book looks at these topics in an integrated way, which brings out their functional similarities and differences. More advanced topics in routing have been moved to a second chapter that can be skipped, depending on the emphasis and level of the class.

I have never been a fan of teaching networking based on a purely layered approach, as my foreword to the first edition indicated (we’ve reprinted it in this edition just for fun.) Some key issues in networking, including security and performance, cannot be solved by assigning them to one layer—there cannot be a “performance” layer. These sorts of topics are both critical and cross-cutting, and the organization of this book continues to treat topics, as well as layers. The organization of this book reflects a great deal of experience using it as a classroom textbook, and as well a preference for an approach that brings out fundamentals as well as current practice.

Some moribund technologies are now missing or minimized, including token ring (one of my old favorites, but clearly it was time to go) and ATM. This edition recognizes that we need to pay more attention to application design, and not just packet forwarding. Wireless and mobility gets more attention as well.

The authors, once again, have worked hard to produce a revision that conveys the essentials of the field in a way that is pedagogically effective. I am pleased to say that I think it is better than ever.

David Clark  
November, 2010

# Foreword to the First Edition

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.

**David D. Clark**

*Massachusetts Institute of Technology*

# Preface

**W**hen the first edition of this book was published in 1996, 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. The primary way for a household to connect to the Internet was via a dial-up modem. Today, Internet commerce is a fact of life, and “.com” stocks have gone through an entire boom and bust cycle. Wireless networks are everywhere and new Internet-capable devices such as smartphones and tablets appear on the market at a dizzying pace. It seems the only predictable thing about the Internet is constant change.

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 more than 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 3 decades is precisely the right place to start. Like the previous editions, the Fifth 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.

We also recognize that there are many different ways that people approach networks. In contrast to when we wrote our first edition, most people will pick up this book having considerable experience as *users* of networks. Some will be looking to become *designers* of networking products or protocols. Others may be interested in *managing* networks, while an increasingly large number will be current or prospective *application developers* for networked devices. Our focus has traditionally been on the designers of future products and protocols, and that continues to be the case, but in this edition we have tried to address the perspectives of network managers and application developers as well.

### Changes in the Fifth Edition

Even though our focus is on the underlying principles of networking, we illustrate these principles using examples from today's working Internet. Therefore, we added a significant amount of new material to track many of the important recent advances in networking. We also deleted, reorganized, and changed the focus of existing material to reflect changes that have taken place over the past decade.

Perhaps the most significant change we have noticed since writing the first edition is that almost every reader is now familiar with networked applications such as the World Wide Web and email. For this reason, we have increased the focus on applications, starting in the first chapter. We use applications as the motivation for the study of networking, and to derive a set of requirements that a useful network must meet if it is to support both current and future applications on a global scale. However, we retain the problem-solving approach of previous editions that starts with the problem of interconnecting hosts and works its way up the layers to conclude with a detailed examination of application layer issues. We believe it is important to make the topics covered in the book relevant by starting with applications and their needs. At the same time,

we feel that higher layer issues, such as application layer and transport layer protocols, are best understood after the basic problems of connecting hosts and switching packets have been explained. That said, we have made it possible to approach the material in a more *top-down* manner, as described below.

As in prior editions, we have added or increased coverage of important new topics, and brought other topics up to date. Major new or substantially updated topics in this edition are:

- Updated material on wireless technology, particularly the various flavors of 802.11 (Wi-Fi) as well as cellular wireless technologies including the third generation (3G) and emerging 4G standards.
- Updated coverage of congestion control mechanisms, particularly for high bandwidth-delay product networks and wireless networks.
- Updated material on Web Services, including the SOAP and REST (Representational State Transfer) architectures.
- Expanded and updated coverage of interdomain routing and the border gateway protocol (BGP).
- Expanded coverage on protocols for multimedia applications such as voice over IP (VOIP) and video streaming.

We also reduced coverage of some topics that are less relevant today. Protocols moving into the “historic” category for this edition include asynchronous transfer mode (ATM) and token rings.

One of the most significant changes in this edition is the separation of material into “introductory” and “advanced” sections. We wanted to make the book more accessible to people new to networking technologies and protocols, without giving up the advanced material required for upper-level classes. The most apparent effect of this change is that Chapter 3 now covers the basics of switching, routing, and Internetworking, while Chapter 4 covers the more advanced routing topics such as BGP, IP version 6, and multicast. Similarly, transport protocol fundamentals are covered in Chapter 5 with the more advanced material such as TCP congestion control algorithms appearing in Chapter 6. We believe this will make it possible for readers new to the field to grasp important foundational concepts without getting overwhelmed by more complex topics.

As in the last edition, we have included a number of “where are they now?” sidebars. These short discussions, updated for this edition, focus on the success and failure of protocols in the real world. Sometimes they

describe a protocol that most people have written off but which is actually enjoying unheralded success; other times they trace the fate of a protocol that failed to thrive over the long run. The goal of these sidebars is to make the material relevant by showing how technologies have fared in the competitive world of networking.

## 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 25-plus 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:

- *First Principles.* 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.
- *Non-layerist.* 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. Similarly, routers and switches have so much in common (and are often combined as single products) that we discuss them in the same chapter. 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.
- *Real-world examples.* Rather than explain how protocols work in the abstract, we use the most important protocols in use today—most 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.

- *Software.* 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. We also include code segments taken from a working protocol stack to illustrate how you might implement certain protocols and algorithms.
- *End-to-end focus.* 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.
- *Performance.* 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.
- *Design Principles.* 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.

## **Pedagogy and Features**

The Fifth Edition retains the key pedagogical features from prior editions, which 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.
- *Where-are-they-now sidebars.* These new elements, a distinctively formatted style of sidebar, trace the success and failure of protocols in real-world deployment.
- *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.
- *What's Next? discussions.* We conclude the main body of each chapter with an important issue that is currently unfolding in the research community, the commercial world, or society as a whole. We have found that discussing these forward-looking 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. Motivated by wide-spread applications, it discusses what goes into a network architecture, provides an

introduction to protocol implementation issues, and defines the quantitative performance metrics that often drive network design.

- Chapter 2 surveys the many ways that a user can get connected to a larger network such as the Internet, thus introducing the concept of *links*. It also describes many of the issues that all link-level protocols must address, including encoding, framing, and error detection. The most important link technologies today—Ethernet and Wireless—are described here.
- Chapter 3 introduces the basic concepts of switching and routing, starting with the virtual circuit and datagram models. Bridging and LAN switching are covered, followed by an introduction to internetworking, including the Internet Protocol (IP) and routing protocols. The chapter concludes by discussing a range of hardware- and software-based approaches to building routers and switches.
- Chapter 4 covers advanced Internetworking topics. These include multi-area routing protocols, interdomain routing and BGP, IP version 6, multiprotocol label switching (MPLS) and multicast.
- 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. The *Real-time Transport Protocol (RTP)*, which supports multimedia applications, is also described.
- Chapter 6 discusses congestion control and resource allocation. The issues in this chapter cut across the link level (Chapter 2), 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 to provide quality of service in IP.
- Chapter 7 considers the data sent through a network. This includes both the problems of presentation formatting and data compression. XML is covered here, and the compression section includes explanations of how MPEG video compression and MP3 audio compression work.
- Chapter 8 discusses network security, beginning with an overview of cryptographic tools, the problems of key distribution, and a

discussion of several authentication techniques using both public and private keys. The main focus of this chapter is the building of secure systems, using examples including Pretty Good Privacy (PGP), Secure Shell (SSH), and the IP Security architecture (IPSEC). Firewalls are also covered here.

- Chapter 9 describes a representative sample of network applications, and the protocols they use, including traditional applications like email and the Web, multimedia applications such as IP telephony and video streaming, and overlay networks like peer-to-peer file sharing and content distribution networks. Infrastructure services—the Domain Name System (DNS) and network management—are described. The Web Services architectures for developing new application protocols are also presented here.

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 4 and 6 through 8. Chapter 9 then returns to the popular topic of network applications. An undergraduate class might reasonably skim the more advanced sections (e.g., Sections 5.3, 9.3.1, 9.3.2 and 9.2.2.)

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 Chapter 4 and the later chapters in depth.

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, and we have included solutions to select exercises.

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