EDITION 4

COMPUTER NETWORKS

A SYSTEMS APPROACH

LARRY L. PETERSON & BRUCE S. DAVIE

MORGAN KAUFMANN PUBLISHERS
An imprint of Elsevier

COMPUTER NETWORKS, 4/e
Peterson and Davie

Morgan Kaufmann Publishers
An imprint of Elsevier
500 Sansome Street, Suite 400,
San Francisco, CA 94111

Original ISBN: 978-0-12-374013-7

First Printed in India 2007
Reprinted 2008 (Thrice)
Reprinted 2009

Indian Reprint ISBN: 978-81-312-1045-1

This edition is for sale in Bangladesh, Bhutan, India, Maldives, Nepal, Pakistan, Sri Lanka and designated countries in South-East Asia through Elsevier (Singapore) Pte. Ltd. Sale and purchase of this book outside of these countries is unauthorised by the publisher.

Published by Elsevier, a division of Reed Elsevier India Private Limited,
Logix Park, Ist Floor, A-4 & 5, Sector-16, Noida-201301, UP (India)

Printed and bound in India by Sheel Print-N-Pack
D-132, Hosiery Complex, Phase II, NOIDA 201 305.

To Lee Peterson and Robert Davie

FOREWORD

David D. Clark
Massachusetts Institute of Technology

It is now ten years since this classic book first appeared. Looking back, it is amazing what has happened in that time. We have seen the transformation of the Web from a small experiment to a World Wide phenomenon. We have seen the emergence of voice over IP and peer-to-peer content sharing. We have seen technology speed up a hundred-fold, the emergence of broadband to the home, and the rise of botnets and other horrid security problems. Many things have changed, technology has come and gone, and (perhaps equally amazing) much of the basics of the Internet are still there.

This book, too, has changed much in ten years, with four editions to keep up. But the basic value of the book remains the same as the first edition. This book gives you the facts you need, and puts those facts into the larger context so that the knowledge you gain will be of value even as the details change. Reading this book informs you about today and prepares you for tomorrow. One new feature is a set of sidebars that illustrate the context of ideas being presented in the text—the *why* of the ideas. Why did an idea fail? Why did it succeed?

What has changed in the book? Some technologies have faded from sight, and get less attention in this edition. We bid a fond farewell to FDDI and ATM LANs. Some technologies have mutated and emerged in new forms. Remote Procedure Call is no longer a LAN-based low-level invocation mechanism, but the foundation of Internet-wide Web Services. We welcome gigabit Ethernet, an updated and expanded section on wireless, and more on router implementation. The material on TCP is up to date, with discussion of new acknowledgment schemes and extensions for high speed.

With the increasing concern with security, there is a completely revised chapter with a new emphasis on a systems approach to security, and a discussion of threats and how to counter them. And at the end, there is a chapter that helps you "put it all together," using case studies at the application layer (VoIP, multimedia, and peer to peer) to show how all the concepts from the previous chapters combine to provide the system that supports these applications.

The evolution of networks is not going to slow down. Soon we will be talking about the impact of television over IP, the collision of the Internet and sensor networks, and lots of other very new and exciting ideas. But relax—if you read this book today you will have the insights you need for tomorrow.

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 ex-

perience 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, 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 ".com" stocks have gone through an entire boom and bust cycle. A host of new technologies ranging from optical switches to wireless networks are now becoming mainstream. 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 three decades is precisely the right place to start. Like the previous editions, the third 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 Fourth 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 now has some familiarity 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.

As we did in the second and third 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:

- Comprehensively revised and updated coverage of security, with a focus on building secure *systems*, not just on cryptographic algorithms;
- Expanded and updated coverage of XML (extensible markup language);
- An updated section on overlay networks, including "peer-to-peer" networking and "content distribution networks";
- A new section on web services, including the SOAP and REST (Representational State Transfer) architectures;

- Updated material on wireless technology, including the 802.11 (WiFi) and 802.16 (WiMAX) standards as well as cellular wireless technologies including the 3G (third generation) standards;
- Expanded coverage of interdomain routing;
- Expanded coverage on protocols and quality of service for multimedia applications such as voice over IP (VoIP) and video streaming;
- Updated coverage of congestion control mechanisms, particularly for high bandwidth-delay product networks.

In addition, we have added a new feature to this edition: "Where are they now?" sidebars. These short discussions 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 20-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:

- Rather than accept existing artifacts as gospel, we start first with 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. We also include code segments taken from a working protocol stack to illustrate how you might implement certain protocols and algorithms.
- 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.

Pedagogy and Features

The fourth edition retains several features from prior editions, and adds one more, 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.
- *"Where are they now?" sidebars.* These new elements 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.
- *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. Motivated by widespread 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 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 two prevalent switching technologies—switched Ethernet and 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. Unicast, multicast, and interdomain routing are covered.
- 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 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 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. 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 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, 4.4, and 4.5).

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 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 6.4 in the context of Chapter 5.

Exercises

Significant effort has gone into improving the exercises with each new edition. In the second edition we greatly increased the number of problems and, based on class testing, dramatically improved their quality. In the third edition we made two other important changes, which we retained here:

- For those exercises that we felt are particularly challenging or require special knowledge not provided in the book (e.g., probability expertise), we have added an icon ★ to indicate the extra level of difficulty.
- In each chapter we added some extra representative exercises for which worked solutions are provided at the back of the book. These exercises, marked ✓, are intended to provide some help in tackling the other exercises in the book.

In this edition we have added new exercises to reflect the updated content. 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, socket-based programming assignments, as well as simulation labs, 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 and programming assignments are available through the Morgan Kaufmann website at http://www.mkp.com/pd4e. 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.

And finally, as with the third edition, a set of laboratory experiments supplement the book. These labs, developed by Professor Emad Aboelela from the University of Massachusetts Dartmouth, use simulation to explore the behavior, scalability, and performance of protocols covered in the book. Sections that discuss material covered by the laboratory exercises are marked with the icon shown in the margin. The simulations use the OPNET simulation toolset, which is available for free to any one using *Computer Networks* in their course.

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 netbugsPD4e@elsevier.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 all or parts of the manuscript. In addition to those who reviewed prior editions, we wish to thank David Maltz, Bobby Bhattacharjee, and Sarvesh Kaulkarni for their thorough reviews. Thanks also to Ric Pruss and Mike Takefman for their reviews of various sections. We also wish to thank all those who provided feedback and input to help us decide what to do in this edition: Tim Batten, Julio Pontes, and Kevin Mills.

Several members of the Network Systems Group at Princeton contributed ideas, examples, corrections, data, and code to earlier editions of this book. In particular, we would like to thank Andy Bavier, Tammo Spalink, Mike Wawrzoniak, Zuki Gottlieb, George Tzanetakis, and Chad Mynhier. KyoungSoo Park provided valuable help on the exercise solutions, instructor's manual, and lecture slides. As before, we want to thank the Defense Advanced Research Projects Agency, the National Science Foundation, Intel Corporation, and Cisco Systems, Inc. for supporting our networking research over the past several years. Thanks also to Cisco for providing the time for one of us to work on the book.

This edition could not have been produced without the substantial contributions of Mark Abbott, who crafted a great deal of new material for this book in return for not much more than these few lines of thanks.

Finally, 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 original sponsoring editor, Jennifer Young; our editor for this edition, Rick Adams; our developmental editor, Rachel Roumeliotis; and to Karyn Johnson, assistant editor on prior editions. The whole crew at MKP has been a delight to work with over the lifetime of this book.

CONTENTS