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. 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

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.

A Top-Down Pathway

Because most students today come to a networking class familiar with networked applications, a number of classes take the application as their starting point. While we do cover applications at a high level in Chapter 1, it is not until Chapter 9 that application layer issues are discussed in detail. Recognizing that some professors or readers may wish to follow a more top-down ordering, we suggest the following as a possible way to approach the material in this book.

- Chapter 1. This describes applications and their requirements to set the stage for the rest of the material.
- Chapter 9. The sections on traditional applications (Section 9.1) and multimedia applications (Section 9.2) will introduce readers to the concepts of network protocols using the examples of applications with which they are already familiar. Section 9.3.1 (DNS) could also be covered.
- Section 7.2 could be covered next to explain how the data that is generated by multimedia applications is encoded and compressed.
- Chapter 5. Transport protocol basics can now be covered, explaining how the data generated by the application layer protocols can be reliably carried across a network.
- Chapter 3. Switching, Internetworking, and Routing can be understood as providing the infrastructure over which transport protocols run.
- Chapter 2. Finally, the issues of how data is actually encoded and transmitted on physical media such as Ethernets and wireless links can be covered.

Clearly we have skipped quite a few sections in this ordering. For a more advanced course or comprehensive self-study, topics such as resource allocation (Chapter 6), security (Chapter 8), and the advanced topics in Chapter 4 could be added in towards the end. Security could

be covered almost stand-alone, but all these advanced topics will make most sense after IP and TCP have been covered in Chapters 3 and 5 respectively.

Note that the slides made available on our companion site include a set that follows this top-down ordering in addition to the set that follows the order of the book.

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 in 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 Web site at http://mkp.com/computer-networks.

And finally, as with the Fourth 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 an email to our publisher, Morgan Kaufmann, at netbugsPD5e@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 all or parts of the manuscript. In addition to those who reviewed prior editions, we wish to thank Peter Dordal, Stefano Basagni, Yonshik Choi, Wenbing Zhao, Sarvesh Kulkarni, James Menth, and John Doyle (and one anonymous reviewer) for their thorough reviews. Thanks also to Dina Katabi and Hari Balakrishnan 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.

Several members of the Network Systems Group at Princeton contributed ideas, examples, corrections, data, and code to this book. In

particular, we would like to thank Andy Bavier, Tammo Spalink, Mike Wawrzoniak, Stephen Soltesz, and KyoungSoo Park. Thanks also to Shankar M. Banik for developing the two comprehensive sets of slides to accompany the book.

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 original sponsoring editor, Jennifer Young; our editor for this edition, Rick Adams; our developmental editor, Nate McFadden; assistant editor David Bevans; and our production editor, Paul Gottehrer. Thanks also to the publisher at MKP, Laura Colantoni, whose leadership inspired us to embark on this revision.