

教育部高等教育司推荐
国外优秀信息科学与技术系列教学用书

计算机网络

—— 自顶向下方法与 Internet 特色

(影印版)

COMPUTER NETWORKING
A Top-Down Approach Featuring the Internet

■ James F. Kurose
Keith W. Ross



高等教育出版社
Higher Education Press
Pearson Education
出版集团

教育部高等教育司推荐
国外优秀信息科学与技术系列教学用书

计算机网络

——自顶向下方法与 Internet 特色
(影印版)

COMPUTER NETWORKING
A Top-Down Approach Featuring the Internet

James F. Kurose
Keith W. Ross



高等教育出版社



Pearson Education 出版集团

图字：01-2001-2173 号

English Reprint Copyright © 2001 by PEARSON EDUCATION NORTH ASIA LIMIED and HIGHER EDUCATION PRESS

Computer Networking: A Top-down Approach Featuring the Internet from Addison Wesley Longman's edition of the work

Computer Networking: A Top-down Approach Featuring the Internet, 1st edition by James F. Kurose, Keith W. Ross, Copyright © 2001

All Rights Reserved

Published by arrangement with ADDISON WESLEY LONGMAN, a Pearson Education company

This edition is authorized for sale only in the People's Republic of China (excluding the Special Administrative Regions of Hong Kong and Macau)

图书在版编目(CIP)数据

计算机网络：自顶向下方法与 Internet 特色 / (美)
库罗斯 (Kurose, J.F.), (美) 罗斯 (Ross, K.W.)
影印本. —北京：高等教育出版社, 2001
ISBN 7-04-010100-9

I. 计... II. ①库... ②罗... III. 计算机网络 - 高等
学校 - 教材 - 英文 IV. TP393

中国版本图书馆 CIP 数据核字 (2001) 第 045472 号

计算机网络 —— 自顶向下方法与 Internet 特色 (影印版)

James F.Kurose 等

出版发行 高等教育出版社

社 址 北京市东城区沙滩后街 55 号 邮政编码 100009
电 话 010-64054588 传 真 010-64014048

网 址 <http://www.hep.edu.cn>
<http://www.hep.com.cn>

经 销 新华书店北京发行所

印 刷 北京民族印刷厂

开 本 787 × 1092 1/16

版 次 2001 年 8 月影印版

印 张 46.25

印 次 2001 年 8 月第 1 次印刷

字 数 1 104 000

定 价 39.00 元

本书如有缺页、倒页、脱页等质量问题, 请到所购图书销售部门联系调换。

版权所有 侵权必究

前 言

20 世纪末，以计算机和通信技术为代表的信息科学和技术，对世界的经济、军事、科技、教育、文化、卫生等方面的发展产生了深刻的影响，由此而兴起的信息产业已经成为世界经济发展的支柱。进入 21 世纪，各国为了加快本国的信息产业，加大了资金投入和政策扶持。

为了加快我国信息产业的进程，在我国《国民经济和社会发展第十个五年计划纲要》中，明确提出“以信息化带动工业化，发挥后发优势，实现社会生产力的跨越式发展。”信息产业的国际竞争将日趋激烈。在我国加入 WTO 后，我国信息产业将面临国外竞争对手的严峻挑战。竞争成败最终将取决于信息科学和技术人才的多少与优劣。

在 20 世纪末，我国信息产业虽然得到迅猛发展，但与国际先进国家相比，差距还很大。为了赶上并超过国际先进水平，我国必须加快信息技术人才的培养，特别要培养一大批具有国际竞争能力的高水平的信息技术人才，促进我国信息产业和国家信息化水平的全面提高。为此，教育部高等教育司根据教育部吕福源副部长的意见，在长期重视推动高等学校信息科学和技术教学的基础上，将实施超前发展战略，采取一些重要举措，加快推动高等学校的信息科学和技术等相关专业的教学工作。在大力宣传、推荐我国专家编著的面向 21 世纪和“九五”重点的信息科学和技术课程教材的基础上，在有条件的高等学校的某些信息科学和技术课程中推动使用国外优秀教材的影印版进行英语或双语教学，以缩短我国在计算机教学上与国际先进水平的差距，同时也有助于强化我国大学生的英语水平。

为了达到上述目的，在分析一些出版社已影印相关教材，一些学校已试用影印教材进行教学的基础上，教育部高等教育司组织并委托高等教育出版社开展国外优秀信息科学和技术优秀教材及其教学辅助材料的引进研究与影印出版的试点工作。为推动用影印版教材进行教学创造条件。

本次引进的系列教材的影印出版工作，是在对我国高校信息科学和技术专业的课程与美国高校的对比分析的基础上展开的；所影印出版的教材均由我国主要高校

1995/10

的信息科学和技术专家组成的专家组，从国外近两年出版的大量最新教材中精心筛选评审通过的内容新、有影响的优秀教材；影印教材的定价原则上应与我国大学教材价格相当。

教育部高等教育司将此影印系列教材推荐给高等学校，希望有关教师选用，使用后有什么意见和建议请及时反馈。也希望有条件的出版社，根据影印教材的要求，积极参加此项工作，以便引进更多、更新、更好的外国教材和教学辅助材料。

同时，感谢国外有关出版公司对此项引进工作的配合，欢迎更多的国外公司关心并参与此项工作。

教育部高等教育司

二〇〇一年四月


About the Authors

Jim Kurose

Jim Kurose is Professor and Chair of the Department of Computer Science at the University of Massachusetts, Amherst.

He is the eight-time recipient of the Outstanding Teacher Award from the National Technological University, the recipient of the Outstanding Teacher Award from the College of Natural Science and Mathematics at the University of Massachusetts, and the recipient of the 1996 Outstanding Teaching Award of the Northeast Associations of Graduate Schools. He has been the recipient of a GE Fellowship, an IBM Faculty Development Award, and a Lilly Teaching Fellowship.

Dr. Kurose is a former Editor-In-Chief of the IEEE Transactions on Communications and of the IEEE/ACM Transactions on Networking. He is active in the program committees for IEEE Infocom, ACM SIGCOMM, and ACM SIGMETRICS. He holds a Ph.D. in Computer Science from Columbia University.



Keith Ross

Keith Ross is Chair of the Multimedia Communications Department at Institut Eurecom. From 1985 through 1997 he was a Professor at the University of Pennsylvania where he held appointments in both the Department of Systems Engineering and the Wharton School of Business. In 1999, he co-founded the Internet startup Wimba.com.

Dr. Ross has published over 50 papers and written two books. He has served on editorial boards of five major journals, and has served on the program committees of major networking conferences, including IEEE Infocom and ACM SIGCOMM. He has supervised more than ten Ph.D. theses. His research and teaching interests included multimedia networking, asynchronous learning, Web caching, streaming audio and video, and traffic modeling. He received his Ph.D. from the University of Michigan.



To Julie and our three precious
ones—Chris, Charlie, and Nina
JFK

To my wife, Véronique, and our
trois petits pois, Cécile, Claire, and Katie
KWR

Preface

Welcome to the first edition of *Computer Networking: A Top-Down Approach Featuring the Internet*. We believe that this textbook offers a fresh approach to computer networking instruction. Why is a fresh approach needed, you ask? In recent years we have witnessed two revolutionary changes in the field of networking, changes that are not reflected in the networking texts published in the 1980s and 1990s. First, the Internet has taken over the universe of computer networking. Any serious discussion about computer networking today has to be done with the Internet in mind. Second, over the past ten years the biggest growth area has been in networking services and applications, which can be seen with the emergence of the Web, ubiquitous use of e-mail services, audio and video streaming, Internet phone, ICQ, and online commerce.

We have been teaching computer networking for 15 years, and we will probably continue to teach it for many years to come. Because there is no textbook that covers the underlying principles of networking while at the same time emphasizing Internet protocols and network applications, we were motivated to write this textbook. We believe you will find this new and contemporary approach to computer networking instruction has numerous benefits.

This textbook is for a first course on computer networking. It can be used in both computer science and electrical engineering departments. In terms of programming languages, the book assumes only that the student has experience with C, C++, or Java. A student who has programmed only in C or C++ and not Java should not have any difficulty following the application programming material, even though this material is presented in a Java context. Although this book is more precise and analytical than many other introductory computer networking texts, it rarely uses any mathematical concepts that are not taught in high school. We have made a deliberate effort to avoid using any advanced calculus, probability, or stochastic process concepts. The book is therefore appropriate for undergraduate courses and for first-year graduate courses. It should also be useful to practitioners in the telecommunications industry.

What Is Unique about This Textbook?

The subject of computer networking is enormously vast and complex, involving many concepts, protocols, and technologies that are woven together in an intricate manner. To cope with this scope and complexity, many computer networking texts are often organized around the “layers” of a network architecture. With a layered

organization, students can see through the complexity of computer networking—they learn about the distinct concepts and protocols in one part of the architecture while seeing the big picture of how it all fits together. For example, many texts are organized around the seven-layer OSI architecture. Our personal experience has been that such a layered approach to teaching is indeed highly desirable from a pedagogic perspective. Nevertheless, we have found the traditional approach of teaching bottom-up—that is, from the physical layer towards the application layer—is not the best approach for a modern course on computer networking.

A Top-Down Approach

Unlike the other computer networking texts, this text is organized in a top-down manner—that is, it begins at the application layer and works its way down towards the physical layer. The top-down approach has several important benefits. First, it places emphasis on the application layer, which has been the high “growth area” of computer networking. Indeed, many of the recent revolutions in computer networking—including the Web, audio and video streaming, and content distribution—have taken place at the application layer. We believe that the application layer will remain the highest growth area in the field, both in terms of research and actual deployment. An early emphasis on application-layer issues differs from the approaches taken in most other texts, which have only a small (or nonexistent) amount of material on network applications, their requirements, application-layer paradigms (e.g., client/server), and the application programming interfaces.

Second, our experience as instructors has been that teaching networking applications near the beginning of the course is a powerful motivational tool. Students are thrilled to learn about how networking applications work—applications such as e-mail and the Web, which most students use on a daily basis. Once a student understands the applications, the student can then understand the network services needed to support these applications. The student can then, in turn, examine the various ways in which such services might be provided and implemented in the lower layers. Covering applications early thus provides motivation for the remainder of the text.

Third, the top-down approach enables instructors to introduce network application development at an early stage. Students not only see how popular applications and protocols work, but also learn how easy it is to create their own network applications and application-level protocols. Other introductory computer networking texts do not cover application development and socket programming. (Even though there are books devoted to network programming, they are not introductory networking textbooks.) By providing socket programming examples in Java, we highlight the central ideas without confusing students with complex code. Undergraduates in electrical engineering as well as in computer science should not have difficulty following the Java code. Thus, with the top-down approach, students get early exposure to the notions of application programming interfaces (APIs), service models and protocols—important concepts that resurface in all of the subsequent layers.

An Internet Focus

As indicated by the title, this textbook features the Internet. Most existing textbooks give significant weight to a variety of telecommunications networks and protocol suites, and they treat the Internet as just one of many networking technologies. We instead put the Internet in the spotlight, and use the Internet protocols as a vehicle for studying some of the more fundamental computer networking concepts. But why put the Internet in the spotlight; why not some other networking technology such as ATM? First, computer networking is now synonymous with the Internet. This wasn't the case five-to-ten years ago, when there was a lot of talk about ATM LANs and applications directly interfacing with ATM (without passing through TCP/IP). But now we have reached the point where just about all data traffic is carried over the Internet (or intranets). Today, the only other type of network that competes with the Internet is the circuit-switched telephone networks. And this competitor may also disappear. Although today the majority of voice traffic is carried over the telephone networks, networking equipment manufacturers and telephone company operators are currently preparing for a major migration to Internet technology.

Another benefit for spotlighting the Internet is that most computer science and electrical engineering students are eager to learn about the Internet and its protocols. They use the Internet on a daily basis (at least to send e-mail and surf the Web), and they repeatedly hear stories about how the Internet is a revolutionary and disruptive technology that is profoundly changing our world. Given the enormous relevance of the Internet, students are naturally curious about what is under the hood. Thus, it is easy for an instructor to get students excited about the basic principles when using the Internet as the guiding focus.

Because our book has an Internet focus, it is organized around a five-layer Internet architecture rather than around the more traditional seven-layer OSI architecture. These five layers consist of the application, transport, network, link, and physical layers.

Addressing the Principles

The field of networking is now mature enough that a number of fundamentally important issues can be identified. For example, in the transport layer, the fundamental issues include reliable communication over an unreliable network layer, connection establishment/teardown and handshaking, congestion and flow control, and multiplexing. In the network layer, two fundamentally important issues are how to find "good" paths between two routers and how to deal with the interconnection of a large number of heterogeneous systems. In the data-link layer, a fundamental problem is how to share a multiple access channel. This text identifies fundamental networking issues as well as approaches towards addressing these issues. We believe that the combination of using the Internet to get the student's foot in door and then emphasizing the issues and solution approaches will allow the student to quickly understand just about any networking technology.

The Web Site

One of the most unique and innovative features of this textbook is that it has an extensive companion Web site at <http://www.awl.com/kurose-ross>, which includes:

- ◆ *The entire text online!* If an instructor or student does not have the hard copy handy, it can be accessed at anytime from anyplace by simply going to the site's Web page. Being online also enables us to use more fonts and colors (both within the text and in diagrams), making the text both perky and cheerful. Also, the online format will allow us to publish updates, enabling the text to keep pace with this rapidly changing field.
- ◆ *Over five hundred links to relevant material.* As all of us Internet enthusiasts know, much of the best material describing the Internet is in the Internet itself. The hyperlinks of the online version, embedded in a coherent context, provide the reader direct access to some of the best sites relating to computer networking and Internet protocols. The links point not only to RFCs, journal and conference articles, but also to sites that are more pedagogic in nature, including homebrewed pages on particular aspects of Internet technology and articles appearing in online trade magazines. Professors can assign the material behind the links as supplementary or even required reading. We expect to check the accuracy of the links several times a year.
- ◆ *Interactive learning material.* The site contains interactive Java applets, illustrating key networking concepts. It also provides direct access to the programs such as the Traceroute program (through your browser) that shows the path that packets follow in the Internet. Professors can use these interactive features as mini labs. The Web site also provides direct access to search engines for Internet Drafts, and to a newsgroup in which topics of this book are discussed. A search feature allows readers to search through the online version of the text. Finally, the site also makes available interactive quizzes that permit students to check their basic understanding of the subject matter.

We also expect to be constantly expanding the Web site, adding contributed material from instructors and readers, as well as our own new features such as online lectures. Updates should occur about every three months. If you have any problems with the Web site, send mail to aw.cse@awl.com.

The online version of this book and its companion Web site are ideally suited for asynchronous online courses. Such courses are particularly attractive to students who commute to school or have difficulty scheduling classes due to course time conflicts. Using an earlier draft of the online version of this text, the authors have given asynchronous, fully online courses on computer networking. We have found that one successful asynchronous format is to assign weekly online readings and to have students participate in weekly newsgroup discussions about the readings. Instructors can give students a virtual presence by collecting the URLs of all the stu-

dent homepages and listing the URLs on the class Web page. Students can even collaborate on joint projects, such as research papers and network application development, asynchronously over the Internet. If you are interested in learning more about asynchronous online learning, please visit the Asynchronous Learning Network site at <http://www.aln.org>.

Pedagogical Features

We have been teaching computer networking for over 15 years. We bring to this text a combined 30 years of teaching experience to over 3,000 students. We have also been active researchers in computer networking for over 20 years. (In fact, Jim and Keith first met each other as master's students in a computer networking course taught by Mischa Schwartz in 1979 at Columbia University.) We think this gives us a good perspective on where networking has been and where it is likely to go in the future. Nevertheless, we have resisted temptations to bias the material in this book towards our own pet research projects. We figure you can visit our personal Web sites if you are interested in our research. Thus, this book is about modern computer networking—it is about contemporary protocols and technologies as well as the underlying principles behind these protocols and technologies.

Principles in Practice Sidebars

One important characteristic of this book is that it stresses the underlying principles of computer networking and the role of these principles in practice. In each of the chapters we include a special sidebar that highlights an important principle in computer networking. These sidebars will help students appreciate some of the fundamental concepts being applied in modern networking.

Historical Sidebars

The field of computer networking, beginning in the late 1960s, has a rich and fascinating history. We have made a special effort in the text to tell the history of computer networking. This is done with a special historical section in Chapter 1 and with about a dozen historical sidebars sprinkled throughout the chapters. In these historical pieces, we cover the invention of packet switching, the evolution of the Internet, the birth of major networking giants such as Cisco and 3Com, and many other important events. Students will be stimulated by these historical pieces. Also, as historians tell us, history helps us predict the future. And correctly predicting the future in this rapidly changing field is critical for the success of any networking protocol or technology.

Interviews

We have included yet another original feature that should inspire and motivate students—interviews with renowned innovators in the field of networking. We provide interviews with Leonard Kleinrock, Tim Berners-Lee, Sally Floyd, J.J. Garcia-Luna-Aceves, Bob Metcalfe, Henning Schulzrinne, Phillip Zimmermann, and Jeff Case.

Supplements for Instructors

We realize that changing your approach to teaching a course can be demanding on your time. To aid you in the transition, we are providing a complete supplements package, including:

- ◆ *Power Point Slides.* The course Web site provides Power Point slides for all eight chapters. The slides cover each chapter in detail. They use graphics and animations rather than relying only on monotonous text bullets, to makes the slides interesting and visually appealing. We provide the original Power Point slides to instructors so you can customize them to best suit your own teaching needs.
- ◆ *Laboratory assignments.* The Web site also provides several detailed programming assignments, including an assignment on building a multithreaded Web server, on building an e-mail client with a GUI interface, programming the sender and receiver sides of a reliable data transport protocol, and an assignment on routing in the Internet.
- ◆ *Homework Solutions.* The Web site provides a solutions manual for the homework problems in the text. These solutions are for instructors only. They are available by contacting your Addison-Wesley sales representative or by sending an e-mail message to aw.cse@awl.com.

Chapter Dependencies

The first chapter of this text presents a self-contained overview of computer networking. Introducing many key concepts and terminology, this chapter sets the stage for the rest of the book. All of the other chapters directly depend on this first chapter. We recommend that, after completing Chapter 1, instructors cover Chapters 2 through 5 in sequence, thereby teaching according to the top-down philosophy. Each of these five chapters leverages material from the preceding chapters.

After completing the first five chapters, the instructor has quite a bit of flexibility. There are no interdependencies among the last three chapters, so they can be

taught in any order. However, each of last three chapters depends on the material in the first five chapters. Ideally, the instructor will have time to teach selected material from all three final chapters in a full semester course.

We also note that the first chapter of the text, being comprehensive and self-contained, can serve as the foundation of a short course on networking.

One Final Note

We encourage instructors and students to create new Java applets that illustrate the concepts and protocols in this book. If you have an applet that you think would be appropriate for this text, please submit it to the authors. If the applet (including notation and terminology) are appropriate, we will be happy to include it on the text's Web site, with an appropriate reference to the authors of the applet. We also encourage instructors to send us new homework problems (and solutions), that would complement the current homework problems. We will post these on the instructor-only portion of the Web site.

We also encourage students and instructors to e-mail us about any comments they might have about either the hard-copy or online versions. Feel free to send us interesting URLs, to point out typos, to disagree with any of our claims, and to tell us what works and what doesn't work. Tell us what you think should or shouldn't be included in the next edition. Send your e-mail to kurose@cs.umass.edu and ross@eurecom.fr.

Acknowledgments

Since this project began in 1996, many people have given us invaluable help and have been influential in shaping our thoughts on how to best organize and teach a networking course. We want to say A BIG THANKS to everyone that has helped, including the hundreds of students that have tested preliminary versions. Special thanks go out to:

Al Aho (Lucent Bell Laboratories)

Paul Amer (University of Delaware)

Daniel Brushteyn (former student at University of Pennsylvania, who wrote the Ethernet applet)

Jeff Case (SNMP Research International)

John Daigle (University of Mississippi)

Philippe Decuetos (Institute Eurécom)

Michalis Faloutsos (University of California, Riverside)

Wu-chi Feng (Ohio State University)

Sally Floyd (ACIRI)
J.J. Garcia-Luna-Aceves (University of California at Santa Cruz)
Mario Gerla (University of California at Los Angeles)
Phillipp Hoschka (INRIA/W3C)
Albert Huang (former student of University of Pennsylvania, who wrote the fragmentation applet)
Sugih Jamin (University of Michigan)
Jussi Kangasharju (Eurécom Institute, who helped with the homework solutions and some of the online labs)
Hyojin Kim (former student at University of Pennsylvania, who wrote the flow control applet)
Leonard Kleinrock (UCLA)
Tim-Berners Lee (World Wide Web Consortium)
Brian Levine (University of Massachusetts)
William Liang (former Ph.D. student of University of Pennsylvania)
Willis Marti (Texas A&M)
Deep Medhi (University of Missouri, Kansas City)
Bob Metcalfe (International Data Group)
Erich Nahum (formerly of the University of Massachusetts, now at IBM Research)
Craig Partridge (BBN Technologies)
Radia Perlman (Sun Microsystems)
Jitendra Padhye (formerly of the University of Massachusetts, now at ACIRI)
George Polyzos (University of California at San Diego)
Ken Reek (Rochester Institute of Technology)
Martin Reisslein (Arizona State University)
Despina Saporilla (University of Pennsylvania)
Henning Schulzrinne (Columbia University)
Mischa Schwartz (Columbia University)
Prashant Shenoy (University of Massachusetts)
Subin Shrestha (University of Pennsylvania)
Peter Steenkiste (Carnegie Mellon University)
Tatsuya Suda (University of California at Irvine)
Kin Sun Tam (State University of New York at Albany)
Don Towsley (University of Massachusetts)
David Turner (Eurécom Institute, who helped with one of the online labs)
Raj Yavatkar (Intel, formerly at the University of Kentucky)
Yechiam Yemini (Columbia University)
Ellen Zegura (Georgia Institute of Technology)
Hui Zhang (Carnegie Mellon University)
Lixia Zhang (University of California, Los Angeles)
ZhiLi Zhang (University of Minnesota)
Shuchun Zhang (former student at University of Pennsylvania, who wrote the message switching applet)
Phil Zimmermann (Network Associates)

Jim and Keith want to also thank the entire Addison-Wesley team, which has done an absolutely outstanding job (and who have put up with two very finicky and often behind-schedule authors!). Our thanks to: Joyce Cosentino, Susan Hartman, Michael Hirsch, Lisa Kalner, Patty Mahtani, Helen Reebenacker, and Amy Rose. Special thanks go to Susan, our editor at Addison-Wesley. This book could not and would not have been completed without her management, encouragement, patience, and perseverance.

责任编辑 张海波
封面设计 韩 冲
责任印制 陈伟光

