


INCLUDES  
CD-ROM

# NETWORK ROUTING

Algorithms, Protocols, and Architectures



**MK**<sup>®</sup>  
MORGAN KAUFMANN

DEEPANKAR MEDHI

KARTHIKEYAN RAMASAMY

# Network Routing

Algorithms, Protocols, and Architectures

---

Deepankar Medhi

Karthikeyan Ramasamy



AMSTERDAM • BOSTON • HEIDELBERG • LONDON  
NEW YORK • OXFORD • PARIS • SAN DIEGO  
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

Morgan Kaufmann is an imprint of Elsevier



MORGAN KAUFMANN PUBLISHERS

Senior Acquisitions Editor	Rick Adams
Acquisitions Editor	Rachel Roumeliotis
Publishing Services Manager	George Morrison
Senior Production Editor	Dawnmarie Simpson
Cover Design	Eric DeCicco/Yvo Riezebos Design
Cover Image	Getty Images
Composition	VTEX
Copyeditor	SPi
Proofreader	SPi
Indexer	SPi
Interior printer	The Maple-Vail Book Manufacturing Group
Cover printer	Phoenix Color, Inc.

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

This book is printed on acid-free paper.

©2007 by Elsevier Inc. All rights reserved.

Designations used by companies to distinguish their products are often claimed as trademarks or registered trademarks. In all instances in which Morgan Kaufmann Publishers is aware of a claim, the product names appear in initial capital or all capital letters. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopying, scanning, or otherwise—without prior written permission of the publisher.

Permissions may be sought directly from Elsevier's Science & Technology Rights Department in Oxford, UK: phone: (+44) 1865 843830, fax: (+44) 1865 853333, E-mail: [permissions@elsevier.com](mailto:permissions@elsevier.com). You may also complete your request online via the Elsevier homepage (<http://elsevier.com>), by selecting "Support & Contact" then "Copyright and Permission" and then "Obtaining Permissions."

#### Library of Congress Cataloging-in-Publication Data

Medhi, Deepankar.

Network routing : algorithms, protocols, and architectures / Deepankar Medhi, Karthikeyan Ramasamy.

p. cm.

Includes bibliographical references and index.

ISBN-13: 978-0-12-088588-6 (hardcover : alk. paper)

ISBN-10: 0-12-088588-3 (hardcover : alk. paper) 1. Computer networks. 2. Routers (Computer networks) 3. Computer network architectures. I. Ramasamy, Karthikeyan, 1967- II. Title.

TK5105.5.M425 2007

004.6-dc22

2006028700

ISBN 13: 978-0-12-088588-6

ISBN 10: 0-12-088588-3

For information on all Morgan Kaufmann publications,  
visit our Web site at [www.mkp.com](http://www.mkp.com) or [www.books.elsevier.com](http://www.books.elsevier.com)

Printed in the United States of America

07 08 09 10      5 4 3 2 1

Working together to grow  
libraries in developing countries

[www.elsevier.com](http://www.elsevier.com) | [www.bookaid.org](http://www.bookaid.org) | [www.sabre.org](http://www.sabre.org)

ELSEVIER

BOOK AID  
International

Sabre Foundation

# Foreword

My involvement with computer networking started with TheoryNet (1977), an e-mail system for theoretical computer scientists. Later (1981) I helped lead the computer science network (CSNET) project, which eventually connected most academic and many industrial computer research groups. In the early days, our efforts were primarily focused on providing connectivity and being able to use applications such as e-mail, ftp, and telnet. However, even in the simple (by today's standards) environment of the 1970s and early 1980s (Arpanet, CSNET, and other experimental Internet networks), getting routing "right" turned out to be quite challenging.

I was fortunate to be part of the NSFNET regional/backbone model development. This is when I began to fully understand the significance of routing in a large-scale multi-domain network and, in particular, the central role of policy issues in such a decentralized environment. Over the past decade, as the Internet became ubiquitous and global in scale, routing has become ever more important. Packets must be forwarded efficiently from one end of the world to the other with minimal perception of delay. This has required tremendous efforts on many fronts: how to evolve routing protocols for large-scale loosely-coupled networking environments, how to engineer a network for efficient routing from an operational point of view, how to do efficient packet processing at routers, and how to effectively take into account the complexity of policy issues in the determination of routes. And while there have been many exciting advances over the past two decades, much work remains to be done.

In parallel, we have seen tremendous advances in traditional telephony. The underlying telecommunication system has changed from analog to digital and has incorporated the latest advances in optical technologies and, more recently, voice over IP. Throughout these revolutionary changes, routing has continued to play a critical role.

We are now at a crossroad. Various efforts are underway to determine a framework for next generation networks that allow seamless convergence of services and a platform to more easily create new services. Among other things, this requires a fresh look at routing. To be successful, it is important that we understand what has worked to date. To better understand the issues and complexities, we should look at this broadly, considering a variety of *different network architectures*, not just for the Internet. For each such network architecture, we can benefit from understanding its principles, protocols, algorithms, and functions, with a focus on routing. This will help give us perspective as we consider how to design routing for the next-generation network.

In this regard, Deepankar Medhi and Karthikeyan Ramasamy's book, *Network Routing: Algorithms, Protocols, and Architectures*, is very timely. Departing from most other works, it

is unique in providing an in-depth understanding of routing in a wide variety of types of networks. It includes extensive coverage of the evolution of routing over time. Particularly appealing is its in-depth coverage across a spectrum of algorithmic, technical, experiential, and practical issues. In addition, the detailed coverage of routers and switches is particularly valuable, as it helps the reader gain an understanding of why different approaches and components are needed to address packet processing, especially for scalability. In this regard, it is uniquely successful in drawing an important connection between routing and routers.

Medhi and Ramasamy's presentation is clear and approachable, allowing a wide audience to understand and gain an appreciation of network routing. I believe that it will become a core reference book on routing for router developers, network providers, students, and researchers for both today's practitioners and those who are interested in next-generation routing.

LAWRENCE LANDWEBER

Past John P. Morgridge Chair and Past Department Chairman  
Computer Science Department, University of Wisconsin–Madison  
Fellow, Association for Computing Machinery and  
Recipient of IEEE Award on International Communication  
Former President and Chair of the Board of Trustees, Internet Society

# Preface

In the span of a quarter-century, network routing in communication networks has evolved tremendously. Just a quarter-century ago, the public switched telephone network (PSTN) was running hierarchical routing, ARPANET routing was operational, and the telecommunication infrastructure had fixed static transport routes. In the 1980s, we saw the first tremendous growth in routing: Internet routing was deployed under the TCP/IP stack starting, first with the RIP protocol; the telephone network started deploying dynamic call routing schemes; and the telecommunication transport network deployed SONET transport mechanisms, which could reroute in a ring topology in 40 millisecond in the event of a failure. In the past fifteen years, we have seen the need for policy routing because of multiprovider settings, and the need to develop fast lookup algorithms for packet processing that enables efficient routing. We have also seen interdependency between addressing and routing as first addressed through classless interdomain routing (CIDR) and more recently, because of number portability in the PSTN. More importantly, we saw how the way an addressing scheme is deployed can impact routing and lookup algorithms.

Network routing can be broadly divided into three basic fundamental categories: packet routing, circuit-switched routing, and transport routing; certainly, a combination is possible. The evolution over the past quarter-century has brought to the foreground the need to understand and examine where and how different dimensions of routing, from algorithms to protocols to architectures, can differ for different types of networks, and where they intersect. Certainly, the goal is to learn from our past experiences and prepare ourselves for next generation networks and routing.

While numerous papers have been written on the subject of network routing, and several books are now available on routing for specific networks, the field still lacks a comprehensive or systematic guide that encompasses various routing paradigms. Second, even in discussions of a single routing type (for example, either the Internet or PSTN), the focus often appears to be either on protocols or algorithms without tying them together with analysis and implementation; or, the work delves more into router command-line for router configuration; or, being informational without explaining the whys. Furthermore, how the addressing mechanism can affect routing decisions is yet another important topic that is rarely discussed. For efficient routing, how routers are architected—and why—is yet another mystery. Finally, the relation between traffic engineering and efficient routing is also another topic. In the end, one needs to be somewhat of an “expert” in different routing paradigms to get a well-rounded view.



Last, after investigating routing in different networks for a number of years, we have come to the conclusion that network routing is like an economy. Similar to macroeconomics and microeconomics, network routing also has macro- and micro-centric issues. In addition, seemingly different and conflicting systems can and do co-exist. Not all of the issues are purely technical; business relations and regulatory issues are also important to recognize and consider. This book is an attempt to paint a broad picture that encompasses various aspects of network routing in one place.

#### AUDIENCE

Our goal has been to create a book that can be used by a diverse set of audiences, with varied levels of background. Specifically, we set out to create a book that can be used by professionals, as well as students and researchers. In general, this is intended as a self-study. We assume that the reader already has some basic knowledge of networking. Among professionals, the intent has been to cover two broad groups: router developers, including protocol designers and router architects, and network designers and operators, with the overall goal to bring out issues that one group might want to understand that the other group faces. For students, this book is intended to help learn about routing in depth, along with the big picture and lessons from operational and implementation experience. For researchers who want to know what has been done so far and what critical issues to address for next-generation routing, this is intended as a helpful reference. In general, this book has been intended as a one-stop treat for all interested in network routing in different networks.

#### ORGANIZATION AND APPROACH

The book is organized into six parts. Each part starts with a chapter-level summary. We present below a brief overview of each part:

- Part I (four chapters): We cover the basic foundations of routing from algorithms to protocols, along with network flow modeling.
- Part II (five chapters): This part is about IP network routing, from standardized protocols for both intra- and inter-domain routing, to IP traffic engineering and Internet routing architectures.
- Part III (four chapters): This part covers PSTN routing, from hierarchical routing to dynamic routing, and from addressing to traffic engineering, including the role of signaling in routing, along with the impact of number portability in routing.
- Part IV (three chapters): In this part, we cover router architectures for different scale routers for efficient packet processing, along with address lookup algorithms and packet filtering and classification mechanisms.
- Part V (four chapters): As impetuses for next generation routing, we present quality-of-service routing, multiprotocol label switching, generalized multiprotocol label switching, and routing at the intersection of IP-PSTN for voice over IP.
- Part VI (five chapters): This bonus material (available on the CD-ROM) is made up of two sub-parts: the first three chapters continue beyond Part IV by delving more into routers by

presenting efficient switching, packet queueing and scheduling, and traffic conditioning; the remaining two chapters extend Part V by covering transport network routing, optical network routing, and multi-layer routing.

At the beginning of each chapter, a reading guideline is provided. This gives a brief description on the background needed to read the chapter; it also discusses which other chapters this chapter is connected to or has dependency on. In general, it is not necessary to read the chapters in sequential order. Furthermore, the chapters are organized in a way so that the reader who has familiarity with a particular topic can move on and read other chapters of interest. Similarly, there are a few chapters on traffic engineering that require a certain level of mathematical background. They can be read independently if the reader has the background, or can be skipped for later reading, without missing the broad picture. Regardless, each chapter contains a Further Lookup section, which includes a brief discussion on additional reading; followed by a set of exercises that is meant for a wide audience. Notations, conventions, and symbols used in the book are summarized in Appendix A. Miscellaneous refresher topics that are helpful in understanding the material presented in this book are included in Appendix B.

In general, we have given special attention to being concise about describing each topic, while ensuring that the material is approachable for a wider audience. The book is still hefty in size in order to cover routing in different networks. Despite our keen interest, we needed to make the decision to leave out certain important topics instead of cutting corners on the topics presented. The topics *not* covered in the book (except for cursory remarks) are: multicast routing, routing in ATM networks, routing in cellular/wireless networks, routing in sensor networks, and security vulnerabilities in routing. The router command-line-based configuration of protocols is not included in this book, because there are many detailed books available on this aspect for various Internet routing protocols. Finally, there is a direct connection between routing and capacity design and planning. For an in-depth treatment of capacity design and planning, the reader is referred to the companion book [564].

#### BONUS MATERIALS AND ONLINE RESOURCES

The book, in its printed form, has 20 chapters. A CD-ROM is provided with the book that contains an additional five chapters labeled “Advanced Topics.” Of these five chapters, three chapters are related to router architectures: switching packets (Chapter 21), packet queueing and scheduling (Chapter 22), and traffic conditioning (Chapter 23). The remaining two chapters are related to transport and next-generation routing: transport network routing (Chapter 24), and optical network routing and multilayer routing (Chapter 25).

Additional support materials (for example, instructional materials and additional exercises) will be available at <http://www.mkp.com/?isbn=9780120885886> and <http://www.NetworkRouting.net>. The latter site will also serve as a resource site and will provide links to materials available on the web on network routing.



## ACKNOWLEDGMENTS

To quote Jeff Doyle, “An author of a technical book is just a front man for a small army of brilliant, dedicated people.” We could not have said it better.

Our official technical reviewers did a tremendous job of reading carefully and providing detailed comments. We thank Jennifer Rexford (Princeton University), Ibrahim Matta (Boston University), K. R. Krishnan (Telcordia Technologies), and Kannan Varadhan (Juniper Networks) for lending their expertise, time, and effort.

In addition, many afforded their expertise by reading one or more chapters and by providing valuable feedback; we gratefully acknowledge Amit Shukla (Microsoft), Arthi Ayyangar (Nuova Systems), Caterina Scoglio (Kansas State University), Chelian Pandian (Juniper Networks), Dana Blair (Cisco Systems), David Walden (BBN, retired), Debashis Talukdar (Embarq), Dock Williams (Juniper Networks), Driss Benhaddou (University of Houston), Hua Qin (Beijing University of Technology), Hui Zhang (Carnegie Mellon University), Jeff Naughton (University of Wisconsin–Madison), Jignesh M. Patel (University of Michigan), Johannes Gehrke (Cornell University), John Strand (AT&T Labs), Mario Baldi (Politecnico di Torino), Prasad Deshpande (IBM), Prosper Chemouil (France Telecom R&D), Rahul Agrawal (Juniper Networks), Ravi Chandra (Sona Systems), Raymond Reeves (Sprint), Saad Siddiqi (Sprint), Shachi Sharma (Alcatel), Srinivas Seshadri (Kosmix), Steve Dispensa (Positive Networks), Vamsi Valluri (Cisco Systems), Venkatesh Iyengar (Sun Microsystems), and Vijay Talati (Juniper Networks).

The first author’s colleagues in the Networking group at the University of Missouri–Kansas City, Appie van de Liefvoort, Baek-Young Choi, Cory Beard, Jerry Place, Ken Mitchell, and Khosrow Sohraby, served as great resources. They read one or more chapters, were around to have a quick discussion and to provide their theoretical as well as practical expertise when needed. Appie van de Liefvoort and Khosrow Sohraby, in their roles as administrators, provided a much-needed environment for the first author to carry out a project of this magnitude without too many distractions. More than a decade ago, a former colleague, Adrian Tang, was instrumental and believed in the importance of creating a separate course on network routing; with his interest and the nod from Richard Hetherington (the then director), the first author developed and taught a course on network routing encompassing different networks for the first time in fall 1995; he also benefited from the publication of [667] in 1995 that helped jump-start this course. Since then, he has been teaching this course every fall (except when he was on a sabbatical leave). The content has changed significantly in this short span of time to keep up with what has been happening in the field, providing an exciting challenge and opportunity. He gratefully acknowledges having a sabbatical in 2004 to plan for the initial preparation for this book.

The current and recent PhD students of the first author also read many chapters and provided valuable feedback. Many thanks to Amit Sinha, Balaji Krithikaivasan, Dijiang Huang, Gaurav Agrawal, Haiyang Qian, Plarent Tirana, and Shekhar Srivastava.

Several students who took the course, Network Routing, from the first author, in the fall of 2005, read the initial version of the first few chapters. When he taught it again in the fall 2006 semester, the entire manuscript was ready in its draft form; the entire class helped debug it by carefully reading various chapters and providing detailed feedback. For their help, we would like to thank Aditya Walavalkar, Ajay Karanam, Amol Rege, Dong Yoo, Francisco Jose Landeras, Hafeez Razzaq, Haiyang Qian, Hui Chang, Jignesh K. Patel, Jin-Ho Lee,

Jorge Rodriguez, Palani Ramalingam, Phaneesh Gururaj, Ravi Aute, Rehan Ishrat, Ron McManaman, Satoru Yamashita, Sreeram Gudipudi, Swapnil Deshmukh, Shamanth Kengeri, Shashank Manchireddy, Sundeeep Udutha, Tongan Zhao, and Venkat Pagadala. Needless to say, the first author greatly benefited from many questions and discussions from teaching this course over the past decade that altogether attracted more than 300 students. The second author also benefited from his many interactions with colleagues while working at Juniper Networks. As a result, a range of interrelated topics is included in the book to give a broader perspective of network routing.

Over the years, we have both benefited from informative and enlightening discussions on routing in different domains and related topics from many individuals; many also answered queries during the preparation of this book. We like to thank the following: Aekkachai Ratatanadilokochai (Cisco Systems), Åke Arvidsson (Ericsson), Amarnath Mukherjee (Clarifyre), Ananth Nagarajan (Juniper Networks), André Girard (INRS-EMT), Bharani Chadalavada (Juniper Networks), Brion Feinberg (Sereniti), Brunilde Sansò (University of Montréal), David DeWitt (University of Wisconsin–Madison), David Tipper (University of Pittsburgh), David Mills (University of Delaware), David Walden (BBN, retired), Debasis Mitra (Bell Labs), Di Yuan (Linköping Institute of Technology), Fu Chang (Academia Sinica), Gerald Ash (AT&T Labs), Gerald Combs (CACE Technologies, creator of Ethereal/Wireshark), Geoff Huston (APNIC), Götz Gräfe (HP Labs), Hadriel Kaplan (Acme Packet), Indrajanti (Yanti) Sukiman (Cisco Systems), Iraj Saniee (Bell Labs), Jean-François Labourdette (Verizon), Jeff Naughton (University of Wisconsin–Madison), Jim Pearce (Sprint), John Strand (AT&T Labs), Keith Ross (Polytechnic University), Larry Landweber (University of Wisconsin–Madison), Lindsay Hiebert (Cisco Systems), Lorne Mason (McGill University), Michał Pióro (Warsaw University of Technology and Lund University), Mikkel Thorup (AT&T Labs–Research), Mostafa Ammar (Georgia Tech), Mukesh Kacker (NetApp), Nitin Bahadur (Juniper Networks), Oscar González-Soto (ITU), Philip Smith (Cisco Systems), Pramod Srinivasan (Juniper Networks), Prosper Chemouil (France Telecom R&D), Rajat Monga (Contributor), Ravi Chandra (Sonos Systems), Richard Harris (Massey University), Robert Doverpike (AT&T Labs–Research), Ron Skoog (Telcordia Technologies), Saad Siddiqi (Sprint), Samir Shah (Cisco Systems), Saravan Rajendran (Cisco Systems), Sergio Beker (France Telecom R&D), Shankar Satyanarayanan (Cisco Systems), Srinivasa Thirumalasetty (Ciena Corporation), Steve Dispensa (Positive Networks), Steve Robinson (University of Wisconsin–Madison), Toshikane Oda (Nippon Ericsson), Ulka Ranadive (Cisco Systems), Vamsi Valluri (Cisco Systems), Villy Bæk Iversen (Technical University of Denmark), Wayne Grover (TR-Labs & University of Alberta), Wen-Jung Hsin (Park University), Wesam Alanqar (Sprint), Yufei Wang (VPI Systems), and Zhi-Li Zhang (University of Minnesota). Furthermore, the first author benefited from Karen Medhi's insight and expertise in transport network routing and design.

Folks at AS3390 often provided their perspective from the viewpoint of running a stub AS by answering our questions. Our sincere thanks to the following individuals at AS3390: David Johnston, George Koffler, Jim Schonemann, II, and Justin Malyn.

We thank David Clark (M.I.T.), Series Editor for the Morgan Kaufmann series in Networking, for recognizing the importance of having a book that spans network routing in different networks, and for greenlighting our book proposal. We are honored that Larry Landweber

(University of Wisconsin–Madison) gladly accepted our request to write the Foreword for this book.

The first author would like to thank the Defense Advanced Research Project Agency, the National Science Foundation, the University of Missouri Research Board, and Sprint Corporation for supporting his networking research.

Two individuals deserve special note: (1) Jane Zupan took Network Routing from the first author in 2000. She typed his scattered notes, which gave the initial idea for planning this book. Then, at a conference in 2003, the three of us casually joked about doing this book, and the plan finally started. Jane was an initial partner in this project but could not continue to work on it because of time constraints. She, however, continued to provide much help, even at the last stage, by reading and editing, despite her busy schedule. We sincerely thank her for her great help. (2) Balaji Krithikaivasan, who completed a Ph.D. under the first author, passionately read many chapters of the book and provided feedback despite his deteriorating health. Unfortunately, he did not live long enough to see the final publication of the book. Thank you, Bala, wherever you are.

It has been a pleasure to work with Rick Adams and Rachel Roumeliotis of Morgan Kaufmann Publishers/Elsevier. From the initial proposal of the book to final production, they provided guidance in many ways, not to mention the occasional reminder. We appreciate their patience with us during the final stages of the manuscript preparation. Arline Keithe did a nice job of copyediting. In the pre-production phase, we received help from Project Managers Dawnmarie Simpson and Tracy Grace. Folks at VTEX did an excellent job of taking our  $\text{\LaTeX}$  files and transforming them to production quality, and Judy Ahlers did great work on the final proofread. We thank them all.

Our immediate family members suffered the most during our long hours of being glued to our laptops. Throughout the entire duration, they provided all sorts of support, entertainments, and distractions. And often they queried “Are you ever going to get this done?” Deep would like to thank his wife, Karen, and their sons, Neiloy and Robby, for love and patience, and for enduring this route (for the second time). He would also like to thank cc76batch—you know who you are—for their friendship. Karthik would like to thank his wife, Monika, for her love and patience. He would also like to acknowledge his family members, Sudarshan Kumar and Swarn Durgia, Sonu and Rajat Monga, and Tina and Amit Babel for keeping him motivated. In addition, he would like to thank his many friends for their support. Finally, we like to thank our parents and our siblings for their support and encouragement.

DEEPANKAR (DEEP) MEDHI  
Kansas City, Missouri, USA  
dmedhi@umkc.edu

KARTHIKEYAN (KARTHIK) RAMASAMY  
Santa Clara, California, USA  
karthik@cs.wisc.edu

# About the Authors

**Deepankar Medhi** is Professor of Computer Networking in the Computer Science & Electrical Engineering Department at the University of Missouri–Kansas City, USA. Prior to joining UMKC in 1989, he was a member of the technical staff in the traffic network routing and design department at the AT&T Bell Laboratories. He was an invited visiting professor at Technical University of Denmark and a visiting research fellow at the Lund University, Sweden. He is currently a Fulbright Senior Specialist. He serves as a senior technical editor of the *Journal of Network & Systems Management*, and is on the editorial board of *Computer Networks*, *Telecommunication Systems*, and *IEEE Communications Magazine*. He has served on the technical program committees of numerous conferences including IEEE INFOCOM, IEEE NOMS, IEEE IM, ITC, and DRCN. He received B.Sc. (Hons.) in Mathematics from Cotton College, Gauhati University, India, an M.Sc. in Mathematics from the University of Delhi, India, and an M.S. and a Ph.D. in Computer Sciences from the University of Wisconsin–Madison, USA. He has published more than 70 papers, and is co-author of the book *Routing, Flow, and Capacity Design in Communication and Computer Networks*, also published by Morgan Kaufmann (July 2004).

**Karthikeyan Ramasamy** has 15 years of software development experience, including working with companies such as Juniper Networks, Desana Systems, and NCR. His primary areas of technical expertise are networking and database management. As a member of the technical staff at Juniper, he developed and delivered a multitude of features spanning a diverse set of technologies including protocols, platforms, databases, and high availability solutions for the JUNOS routing operating system. As a principal engineer at Desana Systems, he was instrumental in developing and delivering many fundamental components of an L7 switch for managing data centers. While pursuing his doctorate, he worked on a parallel object relational database system, which was spun off as a company and later acquired by NCR. Subsequently, he initiated a project in data warehousing which was adapted by NCR. As a consulting software engineer at NCR, he assisted in the commercialization of these technologies. Currently, he serves as an independent consultant. He received a B.E. in Computer Science and Engineering with distinction from Anna University, India, an M.S. in Computer Science from the University of Missouri–Kansas City, USA, and a Ph.D. in Computer Sciences from the University of Wisconsin–Madison, USA. He has published papers in premier conferences and holds 7 patents.

# Contents

Foreword	xxiii
Preface	xxv
About the Authors	xxxi

<b>Part I: Network Routing: Basics and Foundations</b>	<b>1</b>
<b>1 Networking and Network Routing: An Introduction</b>	<b>2</b>
1.1 Addressing and Internet Service: An Overview	4
1.2 Network Routing: An Overview	5
1.3 IP Addressing	7
1.3.1 Classful Addressing Scheme	8
1.3.2 Subnetting/Netmask	9
1.3.3 Classless Interdomain Routing	10
1.4 On Architectures	11
1.5 Service Architecture	12
1.6 Protocol Stack Architecture	13
1.6.1 OSI Reference Model	13
1.6.2 IP Protocol Stack Architecture	14
1.7 Router Architecture	19
1.8 Network Topology Architecture	20
1.9 Network Management Architecture	21
1.10 Public Switched Telephone Network	21
1.11 Communication Technologies	22
1.12 Standards Committees	24
1.12.1 International Telecommunication Union	24
1.12.2 Internet Engineering Task Force	25
1.12.3 MFA Forum	25
1.13 Last Two Bits	25
1.13.1 Type-Length-Value	25
1.13.2 Network Protocol Analyzer	26

1.14	Summary	26
	Further Lookup	27
	Exercises	27
<b>2</b>	<b>Routing Algorithms: Shortest Path and Widest Path</b>	<b>30</b>
2.1	Background	31
2.2	Bellman–Ford Algorithm and the Distance Vector Approach	33
2.2.1	Centralized View: Bellman–Ford Algorithm	33
2.2.2	Distributed View: A Distance Vector Approach	36
2.3	Dijkstra’s Algorithm	38
2.3.1	Centralized Approach	38
2.3.2	Distributed Approach	40
2.4	Comparison of the Bellman–Ford Algorithm and Dijkstra’s Algorithm	42
2.5	Shortest Path Computation with Candidate Path Caching	43
2.6	Widest Path Computation with Candidate Path Caching	45
2.7	Widest Path Algorithm	47
2.7.1	Dijkstra-Based Approach	47
2.7.2	Bellman–Ford-Based Approach	49
2.8	$k$ -Shortest Paths Algorithm	49
2.9	Summary	51
	Further Lookup	53
	Exercises	53
<b>3</b>	<b>Routing Protocols: Framework and Principles</b>	<b>56</b>
3.1	Routing Protocol, Routing Algorithm, and Routing Table	57
3.2	Routing Information Representation and Protocol Messages	59
3.3	Distance Vector Routing Protocol	60
3.3.1	Conceptual Framework and Illustration	60
3.3.2	Why Timers Matter	66
3.3.3	Solutions	70
3.3.4	Can We Avoid Loops?	74
3.3.5	Distance Vector Protocol Based on Diffusing Computation with Coordinated Update	74
3.4	Link State Routing Protocol	82
3.4.1	Link State Protocol: In-Band Hop-by-Hop Disseminations	83
3.4.2	Link State Protocol: In-Band Based on End-to-End Session	91
3.4.3	Route Computation	92
3.5	Path Vector Routing Protocol	93
3.5.1	Basic Principle	93
3.5.2	Path Vector with Path Caching	97
3.6	Link Cost	102
3.6.1	ARPANET Routing Metrics	102
3.6.2	Other Metrics	103

3.7	Summary	104
	Further Lookup	105
	Exercises	105
<b>4</b>	<b>Network Flow Modeling</b>	<b>108</b>
4.1	Terminologies	109
4.2	Single-Commodity Network Flow	110
4.2.1	A Three-Node Illustration	110
4.2.2	Formal Description and Minimum Cost Routing Objective	111
4.2.3	Variation in Objective: Load Balancing	114
4.2.4	Variation in Objective: Average Delay	116
4.2.5	Summary and Applicability	117
4.3	Multicommodity Network Flow: Three-Node Example	118
4.3.1	Minimum Cost Routing Case	118
4.3.2	Load Balancing	123
4.3.3	Average Delay	125
4.4	Multicommodity Network Flow Problem: General Formulation	128
4.4.1	Background on Notation	129
4.4.2	Link-Path Formulation	130
4.4.3	Node-Link Formulation	135
4.5	Multicommodity Network Flow Problem: Non-Splittable Flow	137
4.6	Summary	138
	Further Lookup	139
	Exercises	139
<b>Part II:</b>	<b>Routing in IP Networks</b>	<b>141</b>
<b>5</b>	<b>IP Routing and Distance Vector Protocol Family</b>	<b>142</b>
5.1	Routers, Networks, and Routing Information: Some Basics	143
5.1.1	Routing Table	143
5.1.2	Communication of Routing Information	146
5.2	Static Routes	146
5.3	Routing Information Protocol, Version 1 (RIPv1)	147
5.3.1	Communication and Message Format	147
5.3.2	General Operation	149
5.3.3	Is RIPv1 Good to Use?	150
5.4	Routing Information Protocol, Version 2 (RIPv2)	150
5.5	Interior Gateway Routing Protocol (IGRP)	153
5.5.1	Packet Format	153
5.5.2	Computing Composite Metric	154
5.6	Enhanced Interior Gateway Routing Protocol (EIGRP)	157
5.6.1	Packet Format	157
5.7	Route Redistribution	160
5.8	Summary	161



Further Lookup	163
Exercises	164
<b>6 OSPF and Integrated IS-IS</b>	<b>166</b>
6.1 From a Protocol Family to an <i>Instance</i> of a Protocol	167
6.2 OSPF: Protocol Features	168
6.2.1 Network Hierarchy	168
6.2.2 Router Classification	168
6.2.3 Network Types	169
6.2.4 Flooding	170
6.2.5 Link State Advertisement Types	171
6.2.6 Subprotocols	171
6.2.7 Routing Computation and Equal-Cost Multipath	172
6.2.8 Additional Features	176
6.3 OSPF Packet Format	177
6.4 Examples of Router LSAs and Network LSAs	183
6.5 Integrated IS-IS	185
6.5.1 Key Features	186
6.6 Similarities and Differences Between IS-IS and OSPF	189
6.7 Summary	191
Further Lookup	191
Exercises	191
<b>7 IP Traffic Engineering</b>	<b>194</b>
7.1 Traffic, Stochasticity, Delay, and Utilization	195
7.1.1 What Is IP Network Traffic?	195
7.1.2 Traffic and Performance Measures	195
7.1.3 Characterizing Traffic	196
7.1.4 Average Delay in a Single Link System	197
7.1.5 Nonstationarity of Traffic	199
7.2 Applications' View	200
7.2.1 TCP Throughput and Possible Bottlenecks	200
7.2.2 Bandwidth-Delay Product	201
7.2.3 Router Buffer Size	202
7.3 Traffic Engineering: An Architectural Framework	203
7.4 Traffic Engineering: A Four-Node Illustration	204
7.4.1 Network Flow Optimization	204
7.4.2 Shortest Path Routing and Network Flow	206
7.5 Link Weight Determination Problem: Preliminary Discussion	211
7.6 Duality of the MCNF Problem	213
7.6.1 Illustration of Duality Through a Three-Node Network	213
7.6.2 General Case: Minimum Cost Routing	215
7.6.3 Minimization of Maximum Link Utilization	219
7.6.4 A Composite Objective Function	221
7.6.5 Minimization of Average Delay	222

7.7	Illustration of Link Weight Determination Through Duality	226
7.7.1	Case Study: I	226
7.7.2	Case Study: II	231
7.8	Link Weight Determination: Large Networks	232
7.9	Summary	234
	Further Lookup	235
	Exercises	235
<b>8</b>	<b>BGP</b>	<b>238</b>
8.1	BGP: A Brief Overview	239
8.2	BGP: Basic Terminology	242
8.3	BGP Operations	243
8.3.1	Message Operations	243
8.3.2	BGP Timers	244
8.4	BGP Configuration Initialization	245
8.5	Two Faces of BGP: External BGP and Internal BGP	247
8.6	Path Attributes	250
8.7	BGP Decision Process	254
8.7.1	BGP Path Selection Process	254
8.7.2	Route Aggregation and Dissemination	256
8.7.3	Recap	257
8.8	Internal BGP Scalability	257
8.8.1	Route Reflection Approach	258
8.8.2	Confederation Approach	261
8.9	Route Flap Dampening	262
8.10	BGP Additional Features	265
8.10.1	Communities	265
8.10.2	Multiprotocol Extension	265
8.11	Finite State Machine of a BGP Connection	266
8.12	Protocol Message Format	270
8.12.1	Common Header	270
8.12.2	Message Type: OPEN	270
8.12.3	Message Type: UPDATE	272
8.12.4	Message Type: NOTIFICATION	274
8.12.5	Message Type: KEEPALIVE	274
8.12.6	Message Type: ROUTE-REFRESH	274
8.12.7	Path Attribute in UPDATE message	276
8.13	Summary	277
	Further Lookup	278
	Exercises	278
<b>9</b>	<b>Internet Routing Architectures</b>	<b>280</b>
9.1	Internet Routing Evolution	281
9.2	Addressing and Routing: Illustrations	283
9.2.1	Routing Packet: Scenario A	285