



Gerald R. Ash


Traffic Engineering and QoS Optimization of Integrated Voice & Data Networks

TN92
A919

Traffic Engineering and QoS Optimization of Integrated Voice & Data Networks

GERALD R. ASH

Fellow and Senior Technical Consultant, AT&T Labs



E200604207



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

Morgan Kaufmann Publishers is an imprint of Elsevier



MORGAN KAUFMANN PUBLISHERS

Acquisitions Editor	Rick Adams
Publishing Services Manager	George Morrison
Production Editor	Dawnmarie Simpson
Associate Acquisitions Editor	Rachel Roumeliotis
Production Assistant	Melinda Ritchie
Cover Design	Eric DeCicco
Composition	Integra Software Services, Pvt Ltd.
Technical Illustration	Integra Software Services, Pvt Ltd.
Copyeditor	Melissa Revelle
Proofreader	Phyllis Coyne et al. Proofreading
Indexer	Broccoli Information Management
Interior Printer	The Maple-Vail Book Manufacturing Group
Cover Printer	Phoenix Color Corporation

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. 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

Ash, Gerald R.

Traffic engineering and Qos optimization of integrated voice & data networks/Gerald R. Ash.
p. cm.

Includes bibliographical references and index.

ISBN-13: 978-0-12-370625-6 (hardcover : alk. paper)

ISBN-10: 0-12-370625-4 (hardcover : alk. paper)

1. Telecommunication—Traffic—Management.
 2. Computer networks—Quality control.
 3. Internet telephony—Quality control. I. Title.
- TK5105.8865.A84 2006
621.382'1—dc22

2006020078

ISBN 13: 978-0-12-370625-6

ISBN 10: 0-12-370625-4

For information on all Morgan Kaufmann publications, visit our
Web site at www.mkp.com or www.books.elsevier.com

Printed in the United States of America

06 07 08 09 10 5 4 3 2 1

Working together to grow
libraries in developing countries

www.elsevier.com | www.bookaid.org | www.sabre.org

ELSEVIER

BOOK AID
International

Sabre Foundation

Traffic Engineering and QoS Optimization of Integrated Voice & Data Networks

The Morgan Kaufmann Series in Networking

Series Editor, David Clark, M.I.T.

Traffic Engineering and QoS Optimization of Integrated Voice & Data Networks
Gerald R. Ash

IPv6 Core Protocols Implementation
Qing Li, Tatuya Jinmei, and Keiichi Shima

Smart Phone and Next-Generation Mobile Computing
Pei Zheng and Lionel Ni

GMPLS: Architecture and Applications
Adrian Farrel and Igor Bryskin

Network Security: A Practical Approach
Jan L. Harrington

Content Networking: Architecture, Protocols, and Practice
Markus Hofmann and Leland R. Beaumont

Network Algorithmics: An Interdisciplinary Approach to Designing Fast Networked Devices
George Varghese

Network Recovery: Protection and Restoration of Optical, SONET-SDH, IP, and MPLS
Jean Philippe Vasseur, Mario Pickavet, and Piet Demeester

Routing, Flow, and Capacity Design in Communication and Computer Networks
Michal Pióro and Deepankar Medhi

Wireless Sensor Networks: An Information Processing Approach
Feng Zhao and Leonidas Guibas

Communication Networking: An Analytical Approach
Anurag Kumar, D. Manjunath, and Joy Kuri

The Internet and Its Protocols: A Comparative Approach
Adrian Farrel

Modern Cable Television Technology: Video, Voice, and Data Communications, 2e
Walter Ciciora, James Farmer, David Large, and Michael Adams

Bluetooth Application Programming with the Java APIs
C Bala Kumar, Paul J. Kline, and Timothy J. Thompson

Policy-Based Network Management: Solutions for the Next Generation
John Strassner

Computer Networks: A Systems Approach, 3e
Larry L. Peterson and Bruce S. Davie

Network Architecture, Analysis, and Design, 2e
James D. McCabe

MPLS Network Management: MIBs, Tools, and Techniques
Thomas D. Nadeau

Developing IP-Based Services: Solutions for Service Providers and Vendors
Monique Morrow and Kateel Vijayananda

Telecommunications Law in the Internet Age
Sharon K. Black

Optical Networks: A Practical Perspective, 2e
Rajiv Ramaswami and Kumar N. Sivarajan

Internet QoS: Architectures and Mechanisms
Zheng Wang

TCP/IP Sockets in Java: Practical Guide for Programmers
Michael J. Donahoo and Kenneth L. Calvert

TCP/IP Sockets in C: Practical Guide for Programmers
Kenneth L. Calvert and Michael J. Donahoo

Multicast Communication: Protocols, Programming, and Applications
Ralph Wittmann and Martina Zitterbart

MPLS: Technology and Applications
Bruce Davie and Yakov Rekhter

High-Performance Communication Networks, 2e
Jean Walrand and Pravin Varaiya

Internetworking Multimedia
Jon Crowcroft, Mark Handley, and Ian Wakeman

Understanding Networked Applications: A First Course
David G. Messerschmitt

Integrated Management of Networked Systems: Concepts, Architectures, and their Operational Application
Heinz-Gerd Hegering, Sebastian Abeck, and Bernhard Neumair

Virtual Private Networks: Making the Right Connection
Dennis Fowler

Networked Applications: A Guide to the New Computing Infrastructure
David G. Messerschmitt

Wide Area Network Design: Concepts and Tools for Optimization
Robert S. Cahn

For further information on these books and for a list of forthcoming titles, please visit our Web site at <http://www.mkp.com>.

*In memory of my parents,
who gave me life and shaped who I am.
And dedicated to my wife, children, and grandchildren,
who gave my life meaning, fulfillment, and joy.*

Foreword

There is no question that there has been a radical shift in the type of services carried over the Internet. Changes in technology and bandwidth availability have been leveraged by new Internet protocols to realize new services, all of which are implemented over the same core network. This is truly a widening of the functions available to a customer, as the simple data delivery mechanisms have been extended down the food-chain to offer virtual private networks, virtual LANs, and pseudowires to carry transport circuits over the connectionless infrastructure of the Internet. At the same time, the convergence of voice, video, and data has become a reality, with many millions of individuals using voice over IP connections, telecoms companies migrating their telephony to IP, video being streamed point-to-point and point-to-multipoint across the Internet, and the more established data services like Web access continuing to grow. All of these different service types place very different demands for Quality of Service (QoS) and result in the network operator needing to make widely different contractual Service Level Agreements (SLAs) with the customers.

The term *Traffic Engineering* has come to mean very different things to different people. Some view it as a relatively static, off-line, planning activity that is used to dimension capacity planning within a transport network. Others see it as a dynamic mechanism for placing traffic within a network. In either case the objective is the same: to optimize the use of the network resources so that maximum revenue is derived from minimum expenditure. Clearly this objective is met by avoiding congestion, placing existing traffic so that capacity is available to meet further services, providing adequate QoS so that differentiated services can be sold to the customers, and planning capacity so that resources will be available in good time to meet the demands.

Historically, some of the division in the interpretation of traffic engineering stems from the separation between a transport-centric view of networks where long-term circuits are provisioned only after a rigorous planning exercise, and a data-centric view of networking where data are dynamically directed to the available bandwidth, and around network hot-spots. But the networking world is converging, and those who are not willing to embrace both concepts of traffic engineering will be left behind.

Network engineers must be willing to accept the necessity both of careful capacity planning and on-line, dynamic traffic engineering. Modern networks must be able to adapt flexibly to rapid variations in traffic demand, and must be able to offer significant discrimination between service types by meeting substantially different SLAs for each service's traffic.

This flexibility critically includes the ability to perform network planning and traffic engineering across multiple network layers. When we consider that connections within one network layer provide capacity within a higher network layer (for example, a

TDM circuit may realize a TE link in an MPLS network), it should be clear that capacity planning in the higher layer requires network engineering in the lower layer. Thus, as capacity planning becomes more dynamic in response to flexible service demands, the network engineering of lower layers develops into dynamic traffic engineering. Fully coupling these processes across layer boundaries enables multi-layer traffic engineering that can make optimal use of network resources at all layers, and can ensure that multiple client layers can be integrated over a single, lower, server layer.

All of this demands that considerably more attention be paid to the techniques of traffic engineering. We need a formal understanding of the issues within the network and the mechanisms available to provide adequate QoS. We need a thorough analysis of the various possible approaches to traffic management and capacity planning, and we need to tie these methods to our many years of experience with network optimization and the latest operational techniques for predicting traffic behavior, including forecasting, performance monitoring, and fault analysis. With access to this information and the full toolset, we will be ready to take our networks forward to meet the demands of convergence between network layers, as well as the demands of convergence of services onto a common network infrastructure.

Adrian Farrel
Llangollen
Old Dog Consulting
Co-chair of the IETF's CCAMP, PCE, and L1VPN working groups
June 2006

Preface

Why This Book?

About a decade ago AT&T completed its worldwide evolution/revolution to dynamic routing in its global circuit-switched voice/ISDN network. It had taken nearly 2 decades to make that all happen. The first implementation began on July 14, 1984, Bastille Day, which celebrates the French Revolution and itself represented a major revolution in network technology. Both revolutions aimed at introducing more freedom and fairness, but fortunately, the revolution that occurred in 1984 did not result in any chopped-off heads, as did the first revolution: the scientists and engineers responsible for the routing revolution only were subjected to hats-off treatment! Dynamic routing was big news on cutover day, and the national news in the United States covered the event. Over the next 2 decades, the transition to a fully deployed global dynamic routing network was enormous, and the payoff was dramatic. There was much celebrating, backslapping, and crowing along the way. Eventually I published a very large book on the technology, but at about the same time, and as it continuously does, the world was changing fundamentally: the Internet was rising on the horizon.

Astounding breakthroughs by the NetHeads (who we suppose come from Geekia) gave us the Internet: intelligent end-user devices that communicate with packet switching and can define new services, an automated, distributed, and self-organizing network, protocols that are end to end and open, but where quality of service (QoS) is not assured. This Internet revolution trumped 100 years of crown-jewel technological innovation by the BellHeads (who we suppose come from Telephonia), which yielded the greatest machine ever devised: a global, intelligent telephone network, densely connected by circuit switching, protocols that are often proprietary, and where QoS is assured by careful engineering and management.

This Internet technology revolution fully eclipsed the voice/ISDN dynamic routing revolution, and also my book, and over time these were gradually forgotten. But while one door closed, another door opened: constant and unchangeable through the changing world were the traffic engineering and QoS optimization (TQO) principles used to design dynamic routing protocols. TQO controls a network's response to traffic demands and other stimuli, such as network failures, and encompasses traffic management through optimization and control of routing functions, and capacity management through optimization and control of network design. TQO principles have been used to design the revolutionary dynamic routing implementations worldwide

and can be used to design integrated voice/data dynamic routing networks in any technology, particularly Internet technology.

This book explains, illustrates, and applies these design principles, which include class-of-service routing, connection admission control, source-based dynamic path selection, dynamic resource allocation/protection, dynamic transport routing, queuing priority mechanisms, integrated services performance realization, and others. I have been deeply involved in applying these principles to AT&T's network evolution, where I lead the internal routing/addressing/traffic-engineering strategy ("RATS") team that spear-headed the evolution studies. RATS conducted detailed modeling/analysis and case studies for a wide range of alternative architectures under consideration, which are presented throughout this book for both intranetwork and internetwork TQO/dynamic routing design. Network evolution stemming from the RATS effort revolutionized the reliability, performance, traffic handling efficiency, and revenue generation capability of the integrated voice/ISDN network and inspired a worldwide migration to TQO/dynamic routing by many other carriers. We provide detailed case studies of the optimization of multiprotocol label switching (MPLS)/generalized MPLS (GMPLS)-based integrated voice/data dynamic routing networks. These studies illustrate the application of the TQO design principles and provide a basis for generic TQO (GTQO) protocol requirements for MPLS/GMPLS-enabled technologies.

Approach

TQO is important because networks are subject to overloads and failures, no matter how big or how fat ("overprovisioned") we build them and/or how sophisticated we design their management and control technology. We've all experienced communication network overloads and failures, they happen all the time. Web sites go down and congest, terrorist attacks and hurricanes knock out data centers and server farms, and cellular networks are unusable during events such as 9/11 and the devastating 2005 hurricanes. Overload/failure events are typical and unavoidable in all types of networks, no matter what the technology, size, etc. These negative effects of overloads and failures can be greatly mitigated by TQO methods, which have been used in data networks since Morse's invention of the telegraph in 1835 and in voice networks since Bell's invention of the telephone in 1876. There is of course proof for the benefits of TQO, and we show that. TQO/dynamic networks achieve essentially zero traffic loss performance under normal traffic/network conditions, capital savings from efficient network design/optimization, new services revenue, and operational cost efficiencies through automation and real-time network control. New services can be designed and introduced based on the class-of-service routing concept, enabling services to be defined through provisioning of tables and parameters in the traffic router nodes rather than through new software/hardware development.

We describe analysis, design, and simulation models developed by the author and his colleagues over the past 2 decades and use these to analyze the various components of TQO, to conduct large-scale case studies of converged networks, to

design the GTQO protocol, and to quantify the benefits. We use a layered model of TQO [traffic/application layer, MPLS label switched path (LSP)/connection layer, GMPLS LSP/logical link layer, physical network layer, and operations/management layer] and formulate the TQO design problem and discuss its solution at each layer. A comprehensive coherent vision is analyzed at each layer of a converged network architecture, which considers the deep technical issues as well as the impact of the divergent BellHead/NetHead views. As to the latter, a cultural dynamic that became apparent in RATS was the large gap between the BellHead culture, representing the voice/ISDN circuit-switching technologies, and NetHead culture, representing the IP-, ATM-, and frame-relay-based technologies. NetHeads live and build computer data networks and have their geek culture and beliefs, while BellHeads live and build telephone voice networks and have their rather square culture and beliefs. These two worlds have been at war for nigh-on to 40 years, and this “war of two worlds” is the battlefield setting for this book.

Regarding TQO, NetHeads believe that networks should be very fat and designed to carry any traffic the network might encounter—careful engineering is neither desired nor required. BellHeads believe that networks should be carefully engineered to carry the expected traffic—they have sophisticated theories to do careful engineering to assure QoS. NetHeads are scornful of BellHeads’ careful engineering and QoS assurance, “capacity is free,” just put in an infinite amount and stop worrying! BellHeads counter that those who proclaim “capacity is free” aren’t the ones paying for it, and that 100 years of careful traffic engineering have proven successful and wise! This book reflects the lessons and wisdom of both worlds, and in RATS these cultures did work well together, despite the gap, and reached consensus on innovative architecture directions. NetHeads will see a lot of BellHead philosophy in this book, and perhaps declare this a “BellHead book,” but BellHeads will see much NetHead thinking as well, and perhaps declare this a “NetHead book.”

In the end, we speak of truce and focus on the convergence of these two disparate worlds, bridging the gap between BellHeads and NetHeads in TQO space. Both Geekian and Telephonian views are taken into account since both are right. If indeed GTQO methods are developed and implemented, NetHeads should be happy to see that their ingenious protocols, particularly MPLS, GMPLS, and others, are central to the GTQO requirements. BellHeads will recognize that their time-honored and successful networking principles—bandwidth reservation, dynamic alternate routing, traffic management, capacity management, and others—are included as well. So everyone will be happy at least some of the time and no one will be unhappy all of the time with this approach.

A main avenue to realize requirements is through the standards process, and needed standards extensions are discussed, including end-to-end QoS signaling with NSIS (next steps in signaling), PCE (path computation element), DSTE (DiffServ-aware MPLS traffic engineering), MPLS crankback, and others. Alternatives to the GTQO approach, including distributed virtual network approaches, flow-aware networking, centralized TQO approaches, and game theoretic approaches, are presented for thoughtful comparisons and discussion.

Audience

This book is targeted at practitioners, network designers, and software engineers who want an in-depth understanding of TQO of converged IP/MPLS/GMPLS networks and provides an excellent supplementary text for academic courses at the graduate or undergraduate level in computer networking, network design, network routing, optimization, and emerging standards. Practitioners will find descriptions of current TQO trends and solutions and will find answers to many of their everyday questions and problems. This book also provides a comprehensive resource for researchers in traffic engineering, optimization, network design, network routing, voice/data network technology, and network convergence standards.

This book assumes a working knowledge of networks, protocols (particularly Internet protocols), and optimization techniques. The reader should be fully familiar with the concepts of Internet protocol (IP), IP routing, and MPLS/GMPLS signaling basics. These topics are reviewed briefly, and references are provided for more information; however, there is no intent to provide an exhaustive treatment of these topics.

Content

Chapter 1 begins with a general model for TQO functions at each network layer, as well as traffic management and capacity management operational functions. Network layers include the traffic/application layer, connection/MPLS LSP layer, logical link/GMPLS LSP layer, physical network layer, and operations/management layer. We formulate the TQO design problem addressed at each layer and outline the solution approach, where the latter includes an analysis of TQO design and operational experience, as well as design/analysis studies. The analysis of TQO design and operational experience traces the evolution and benefits of TQO methods using ARPANET to illustrate TQO evolution in data networks and the AT&T network to illustrate TQO evolution in voice networks. TQO design principles are identified, and TQO benefits are quantified based on the operational experience. We present the key results and conclusions of the modeling and analysis studies, case studies, and GTQO protocol design.

In Chapter 2, we present models for call/session routing, which entails number/name translation to a routing address associated with service requests, and also compare various connection (bearer-path) routing methods. We introduce a full-scale, 135-node national network model and a multiservice traffic demand model, which are used throughout the book to study various TQO scenarios and trade-offs in TQO optimization, including (a) fixed routing, time-dependent routing, state-dependent routing (SDR), and event-dependent routing (EDR) path selection, (b) two-link and multilink path selection, (c) resource management and connection admission control methods, (d) service priority differentiation of key services, normal services, and best-effort services, and (e) single-area flat topologies versus multiarea hierarchical topologies. The TQO modeling shows that (a) multilink routing in sparse topology networks

provides better overall performance under overload than meshed topology networks, but performance under failure may favor the meshed topology options with more alternate routing choices, and (b) EDR path selection methods exhibit comparable or better network performance compared to SDR methods.

In Chapter 3, we examine QoS resource management methods and illustrate per-flow versus per-virtual-network (VNET) resource management and multiservice integration with priority routing services. QoS resource management includes class-of-service routing, connection admission control, priority routing, bandwidth allocation/protection/reservation, priority queuing, and other related functions. Class-of-service routing provides a means to define network services through table-driven concepts rather than software development and new network deployment. The conclusions reached include (a) bandwidth reservation is critical to stable and efficient network performance and for multiservice bandwidth allocation, protection, and priority treatment and (b) per-VNET bandwidth allocation is essentially equivalent to per-flow bandwidth allocation in network performance and efficiency.

In Chapter 4, we discuss routing table management approaches and provide information exchange requirements needed for interworking across network types. Routing table management entails the automatic generation of routing tables based on information such as topology update, status update, and routing recommendations. This information is used in applying routing table design rules to determine path choices in the routing table. Link-state routing protocols such as open shortest path first (OSPF) use topology-state update mechanisms to build the topology database at each node, typically conveying the topology status through flooding of control messages containing link, node, and reachable-address information. Congestion in link-state protocols can result in widespread loss of topology database information and overload in flooding of topology database information. These and other routing table management information exchange issues are examined in this chapter. Results show that per-VNET QoS resource management, sparse, single-area flat topology, multilink routing, and EDR path selection methods lead to dramatically lower routing table management overhead.

In Chapter 5 we describe methods for dynamic transport routing, which can be realized by the capabilities of GMPLS and optical cross-connect devices to dynamically rearrange transport network capacity. GMPLS technology enables a revolutionary new approach to integrated control of the layer 3 dynamic connection routing and layer 2 dynamic transport routing to shift transport bandwidth among node pairs and services. This allows simplicity of design and robustness to load variations and network failures and provides automatic link provisioning, diverse link routing, and rapid link restoration for improved transport capacity utilization and performance under stress. We conclude that GMPLS-based dynamic transport routing provides greater network throughput and, consequently, enhanced revenue, achieves efficient network design and capital savings, and greatly enhances network performance under failure and overload.

In Chapter 6 we discuss optimization methods and principles for routing design optimization, including shortest path models and discrete event simulation models.

We also discuss optimization methods and principles for capacity design optimization, including (a) discrete event flow optimization (DEFO), (b) traffic load flow optimization, (c) virtual trunk flow optimization, and (d) dynamic transport routing capacity design. We quantify the impacts of traffic variations on network capacity design, including minute-to-minute, hour-to-hour, day-to-day, and forecast uncertainty/reserve capacity design impacts. We illustrate the use of the DEFO model for various comparative analyses, including (a) per-flow versus per-VNET design, (b) multilink versus two-link routing design, (c) single-area flat topologies versus two-level hierarchical topology design, (d) EDR versus SDR design, and (e) dynamic transport routing versus fixed transport routing network design. The conclusions show that (a) sparse topologies with multilink dynamic routing lead to capital cost advantages compared with two-link routing in meshed topologies, (b) EDR methods exhibit comparable design efficiencies to SDR, and (c) dynamic transport routing achieves capital savings by concentrating capacity on fewer, high-capacity physical fiber links. DEFO design models are shown to be extremely flexible and successful in the design of complex routing algorithms and as a basis for network capacity design methods.

In Chapter 7 we present TQO operational requirements for traffic management and capacity management functions in both data and voice networks, including (a) performance management, which collects and analyzes real-time network status and performance data and detects and corrects abnormal network conditions, (b) fault management, which deals with problems and emergencies, such as router failures and power losses, and (c) capacity management, which gathers statistics on equipment and facility use and analyzes trends to project required network upgrades and capacity augments. Traffic management controls are described, including code blocks, connection request gapping, and reroute controls, and we illustrate the conditions that warrant activation of these controls. Capacity management processes are described, including capacity forecasting, daily and weekly performance monitoring, and short-term network adjustment. We illustrate these functions with examples, and in particular we illustrate an MPLS network management implementation by taking an example from AT&T's MPLS operations architecture.

In Chapters 8 and 9 we present several case studies of TQO protocol design in operational networks: (a) circuit-switched integrated voice/ISDN dynamic routing network design for intranetwork applications, (b) circuit-switched integrated voice/ISDN dynamic routing network design for access and internetwork applications, (c) two case studies where TQO designs went astray, but which provide valuable lessons for future designs, and (d) MPLS/GMPLS-based integrated voice/data dynamic routing network design. In the final case study, we develop a 71-node national network model and multiservice traffic demand model for the TQO protocol design and again use the DEFO model for the design and optimization. We illustrate the optimization of an EDR-based path selection protocol from among a large set of candidates, show that a separate emergency-services queue is needed to assure emergency-services performance for scenarios where the normal priority queue congests, and quantify the significant benefits attainable in loss/delay performance for both voice and data traffic. These case studies provide an important basis for the GTQO protocol requirements.

Chapter 10 summarizes the results of studies presented in this book and, based on the results of these studies and operational experience, a GTQO protocol is described for application to MPLS/GMPLS-enabled technologies. Some of the important conclusions derived from the analysis models are (a) EDR path selection is preferred to SDR path selection and (b) aggregated per-VNET bandwidth allocation is preferred to per-flow bandwidth allocation. These design choices reduce control overhead, thereby increasing scalability, whereas the GMPLS-based dynamic transport routing capabilities provide greater network reliability, throughput, revenue, and capital savings. The GTQO requirements apply to access, core, intranetwork, and internetwork architectures and include end-to-end QoS signaling, class-of-service routing, per-VNET QoS resource management, dynamic bandwidth reservation, DSTE bandwidth allocation, EDR path selection, differentiated services (DiffServ) queuing priority, separate high-priority queue for emergency services, and GMPLS-based dynamic transport routing. In addition to the GTQO protocol, we present several other TQO approaches that may well be deployed and identify various standards extensions needed to accommodate the GTQO requirements and capabilities, including end-to-end QoS signaling, path computation element, DiffServ-enabled traffic engineering, MPLS crankback, and others.

Appendix A reviews some of the key TQO technologies: MPLS, GMPLS, QoS mechanisms, integrated services (IntServ), resource reservation protocol (RSVP), DiffServ, and MPLS-based QoS mechanisms. This is intended as a quick refresher and/or a brief introduction for those unfamiliar with these technologies. Ample references are provided for more detailed coverage of these important topics.

Acknowledgments

There are a large number of people I wish to acknowledge and profusely thank for assistance, collaboration, help, encouragement, great ideas, and fun over the past 15 years or so.

Very helpful comments and suggestions were made on the book proposal by Adrian Farrel, Luyuan Fang, Bur Goode, Michal Pioro, Ragu Raghuram, and JP Vasseur. Their review and suggestions definitely improved the presentation and I'm indebted to them for their help and encouragement in the course of developing this work. Thanks to Bur Goode for providing some wake-up-call-type comments on preliminary drafts of the manuscript. Bur's comments helped inspire the book's NetHead/BellHead themes and quest to seek "this is a NetHead book"-type reactions in the final manuscript.

Adrian Farrel, Lorne Mason, and Michal Pioro reviewed the draft manuscript and made extensive comments and suggestions. I feel their candid evaluations and inputs have definitely helped improve this book, and I sincerely appreciate their time and effort. Special thanks to Adrian for all his help and encouragement, from beginning to end, and for writing the Foreword. Rick Adams, Rachel Roumeliotis, and Dawnmarie Simpson provided support, encouragement, helpful suggestions, and patience throughout the course of producing the book, and for that I'm sincerely thankful.

Within AT&T I'd like to especially thank the following people for their excellent contributions, collaborations, support, and on-going discussions: Bruce Blake, Deborah Brungard, Chris Chase, Jin-Shi Chen, Angela Chiu, Suching Chou, Gagan Choudhury, Li Chung, Martin Dolly, Kevin D'Souza, Chuck Dvorak, Jerry Ezrol, Luyuan Fang, Saul Fishman, Tom Frost, Liza Fung, Bur Goode, Mohammed Hamami, Jim Hand, Chuck Kalmanek, Christopher Kwan, Wai Sum Lai, Chin Lee, Yoni Levy, Clayton Lockhart, Joyce Migdall, John Mulligan, Al Morton, Quynh Nguyen, John Oetting, Ragu Raghuram, Scott Sayers, Mostafa Hashem Sherif, Percy Tarapore, Diana Woo, Yung Yu, and David Zerling.

In the IETF I'd like to extend warmest thanks to the following people for all their collaboration, helpful comments, assistance, and encouragement: Arthi Ayyangar, Dan Awduche, Attila Bader, Lou Berger, Nabil Bitar, David Black, Jim Boyle, Igor Bryskin, Anna Charny, Dean Cheng, Yacine El Mghazli, Adrian Farrel, Matthias Friedrich, Xiaoming Fu, Durga Gangiseti, Robert Hancock, Atsushi Iwata, Lars-Erik Jonsson, Cornelia Kappler, Georgios Karagiannis, Kenji Kumaki, Chris Lang, Francois Le Faucheur, Jean-Louis Le Roux, Andy Malis, Andrew McDonald, Thomas Morin, Eiji Oki, Dave Oran, Dimitri Papadimitriou, Tom Phelan, Hal Sandick, Tom Scott,

Hannes Tschofenig, Sven van den Bosch, JP Vasseur, Lars Westberg, and Raymond Zhang.

In the ITU I'd like to particularly thank the following people for their many contributions, insights, discussions, and comments in the course of this work: Mikael Ahman, Anne Elvidge, Davide Grillo, Tommy Petersen, Bruce Pettitt, Jim Roberts, Michael Tuxen, and Manuel Villèn-Altamirano.

Within the ATM Forum I'm indebted to the following people for their significant input and valuable help throughout the course of this work: Carl Rajsic, Peter Roberts, John Ruttmiller, and Mickey Spiegel.

I owe a debt of gratitude to the many other esteemed colleagues for their pioneering, creative work over the past 15 years and who were kind enough to provide information contained in this book: Dave Allan, Joyce Bradley, Terry Brown, Kenneth Chan, Fu Chang, Prosper Chemouil, Jiayu Chen, Joshua Dayanim, Joachim Dressler, Jack Dudash, Victoria Fineberg, Alan Frey, Lindsay Hiebert, BaoSheng Huang, Frank Kelly, Peter Key, K. R. Krishnan, John Labourdette, Kenichi Mase, Lorne Mason, Jerry McCurdy, David McGuigan, Deep Medhi, Chris Metz, Arne Oestlie, Jennifer Rexford, Steven Schwartz, Harmen Van der Linde, and Ben Vos.

I'm deeply indebted to all of these people for their collaborations over the past 15 or more years. It's been my privilege to work with these talented individuals, and to all these people, my heartfelt thanks.