
High Performance TCP/IP Networking

Concepts, Issues, and Solutions

Mahbub Hassan

The University of New South Wales

Raj Jain

The Ohio State University



Upper Saddle River, New Jersey 07458

Vice President and Editorial Director, ECS: *Marcia J. Horton*
Publisher: *Alan R. Apt*
Associate Editor: *Toni Dianne Holm*
Editorial Assistant: *Patrick Lindner*
Vice President and Director of Production and Manufacturing, ESM: *David W. Riccardi*
Executive Managing Editor: *Vince O'Brien*
Managing Editor: *Camille Trentacoste*
Production Editor: *Joan Wolk*
Director of Creative Services: *Paul Belfanti*
Creative Director: *Carole Anson*
Art Director: *Art Director's Name*
Cover Manager: *Jayne Conte*
Cover Designer: *Suzanne Behnke*
Cover Image: *Credit if Appropriate*
Managing Editor, AV Management and Production: *Patricia Burns*
Art Editor: *Gregory Dulles*
Manufacturing Manager: *Trudy Piscioti*
Manufacturing Buyer: *Lynda Castillo*
Marketing Manager: *Pamela Shaffer*



© 2004 by Pearson Education, Inc.
Pearson Prentice Hall
Pearson Education, Inc.
Upper Saddle River, NJ 07458

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

Pearson Prentice Hall® is a trademark of Pearson Education, Inc.

The author and publisher of this book have used their best efforts in preparing this book. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, expressed or implied with regard to these programs or the documentation contained in this book. The author and publisher shall not be liable in any event for incidental or consequential damages in connection with or arising out of the furnishing, performance or use of these programs.

Printed in the United States of America
10 9 8 7 6 5 4 3 2 1

ISBN 0-13-064634-2





Pearson Education Ltd., *London*
Pearson Education Australia Pty. Ltd., *Sydney*
Pearson Education Singapore, Pte. Ltd.
Pearson Education North Asia Ltd., *Hong Kong*
Pearson Education Canada, Inc., *Toronto*
Pearson Educación de Mexico, S.A. de C.V.
Pearson Education Japan, *Tokyo*
Pearson Education Malaysia, Pte. Ltd.
Pearson Education, Inc., *Upper Saddle River, New Jersey*

*To my parents, my wife, my son Aaron, and all
readers of this book*

—Mahbub Hassan

To my wife, Neelu, and my sons, Sameer and Amit

—Raj Jain



Contents

| | |
|--|-----------|
| Preface | xv |
| 1 Introduction | 1 |
| 1.1 History of TCP/IP | 1 |
| 1.2 TCP Applications and Services | 2 |
| 1.3 Motivation for Performance Study of TCP/IP | 3 |
| 1.4 What Do We Mean by TCP Performance? | 4 |
| 1.5 Overview of the Remainder of This Book | 6 |
| 1.6 Further Reading | 8 |
| 1.7 Summary | 9 |
| 1.8 Review Questions | 10 |
| 1.9 Case Study: Introduction to Wireless Corporation | 11 |
| 2 TCP/IP Fundamentals | 12 |
| 2.1 TCP | 12 |
| 2.1.1 TCP Services | 12 |
| 2.1.2 Header Format | 13 |
| 2.1.3 Encapsulation in IP | 16 |
| 2.1.4 Acknowledgment Mechanism | 16 |
| 2.1.5 Retransmission Mechanism | 17 |
| 2.1.6 Connection Establishment and Termination | 18 |
| 2.1.7 Flow Control and Sliding Window | 20 |
| 2.1.8 Congestion Control | 21 |
| 2.2 UDP | 23 |
| 2.2.1 UDP Services | 23 |
| 2.2.2 Header Format | 24 |
| 2.2.3 Encapsulation in IP | 25 |
| 2.3 IP | 25 |
| 2.3.1 IP Services | 25 |
| 2.3.2 Fragmentation and Reassembly | 26 |
| 2.3.3 Header Format | 27 |
| 2.3.4 IP Version 6 | 29 |
| 2.4 Further Reading | 29 |
| 2.5 Summary | 29 |
| 2.6 Review Questions | 29 |
| 2.7 Case Study: WCORP Adopts TCP/IP | 30 |
| 3 Performance Measurement of TCP/IP Networks | 32 |
| 3.1 Reasons for Network Measurement | 32 |
| 3.2 Measurement Tasks | 33 |
| 3.3 Classification of Measurement Tools | 33 |
| 3.4 Popular Measurement Tools and Their Applications | 34 |

viii Contents

| | | |
|----------|---|-----------|
| 3.4.1 | Tcpdump | 35 |
| 3.4.2 | Tcpstat | 43 |
| 3.4.3 | Ttcp | 49 |
| 3.4.4 | Netperf | 53 |
| 3.4.5 | NetPIPE | 58 |
| 3.4.6 | Distributed Benchmark System | 62 |
| 3.5 | Selecting the Right Tool | 71 |
| 3.6 | Further Reading | 72 |
| 3.7 | Summary | 73 |
| 3.8 | Review Questions | 73 |
| 3.9 | Hands-on Projects | 74 |
| 3.10 | Case Study: WCORP Monitors Network Traffic | 74 |
| 4 | TCP/IP Network Simulation | 76 |
| 4.1 | The Role of Simulation | 76 |
| 4.2 | Steps of a Systematic Simulation Study | 77 |
| 4.3 | Types of Simulations | 79 |
| 4.3.1 | Continuous versus Discrete Event | 79 |
| 4.3.2 | Terminating versus Steady State | 80 |
| 4.3.3 | Synthetic versus Trace-Driven Simulation | 81 |
| 4.4 | Simulation Validation and Verification | 82 |
| 4.5 | Confidence Level of Simulation Results | 82 |
| 4.5.1 | Confidence Level Formula | 83 |
| 4.5.2 | Terminating Simulation | 84 |
| 4.5.3 | Steady-State Simulation | 84 |
| 4.5.4 | Common Simulation Mistakes | 87 |
| 4.6 | Simulation with Self-Similar Traffic | 88 |
| 4.7 | Classification of Simulation Tools | 89 |
| 4.8 | The “ns” Network Simulator | 90 |
| 4.8.1 | Model Construction and Parameter Setting | 90 |
| 4.8.2 | Data Collection | 95 |
| 4.8.3 | Simulation Execution | 96 |
| 4.8.4 | Presentation of Results | 97 |
| 4.8.5 | Examples of TCP/IP Simulation Using <i>ns</i> | 97 |
| 4.9 | OPNET | 103 |
| 4.9.1 | Model Construction | 104 |
| 4.9.2 | Parameter Setting | 109 |
| 4.9.3 | Data Collection | 110 |
| 4.9.4 | Simulation Execution | 111 |
| 4.9.5 | Presentation of Results | 113 |
| 4.9.6 | Examples of TCP/IP Simulation Using OPNET | 113 |
| 4.10 | Selecting the Right Tool | 120 |
| 4.11 | Further Reading | 121 |
| 4.12 | Summary | 121 |
| 4.13 | Review Questions | 122 |
| 4.14 | Hands-On Projects | 122 |

| | | |
|----------|--|------------|
| 4.15 | Case Study: WCORP Uses Measurement, Analysis, and Simulation to Dimension Sydney–Melbourne Link Capacity | 122 |
| 5 | TCP Modeling | 125 |
| 5.1 | Motivation for Mathematical Modeling of TCP | 125 |
| 5.2 | Essentials of TCP Modeling | 127 |
| 5.2.1 | Window Dynamics | 127 |
| 5.2.2 | Packet-Loss Process | 128 |
| 5.3 | Gallery of TCP Models | 128 |
| 5.3.1 | Periodic Model | 129 |
| 5.3.2 | Detailed Packet Loss Model | 131 |
| 5.3.3 | Stochastic Model with General Loss Process | 137 |
| 5.3.4 | Control System Model | 141 |
| 5.3.5 | Network System Model | 143 |
| 5.4 | Further Reading | 148 |
| 5.5 | Summary | 149 |
| 5.6 | Review Questions | 150 |
| 5.7 | Hands-on Projects | 151 |
| 5.8 | Case Study: Understanding Factors Influencing TCP Throughput | 151 |
| 6 | TCP/IP Performance over Wireless Networks | 153 |
| 6.1 | Wireless Networks | 153 |
| 6.1.1 | Generic Characteristics | 153 |
| 6.1.2 | Wireless Local Area Networks | 154 |
| 6.1.3 | Cellular Communications Networks | 156 |
| 6.2 | TCP Performance Issues over Wireless Links | 157 |
| 6.2.1 | Inappropriate Reduction of Congestion Window | 157 |
| 6.2.2 | Throughput Loss in WLANs | 158 |
| 6.2.3 | Throughput Loss in Cellular Communication Systems | 159 |
| 6.3 | Improving TCP Performance over Wireless Links | 161 |
| 6.3.1 | Splitting TCP Connections | 161 |
| 6.3.2 | Snooping TCP at Base Stations | 161 |
| 6.3.3 | Notifying the Causes of Packet Loss | 162 |
| 6.3.4 | Adding Selective Acknowledgments to TCP | 163 |
| 6.3.5 | Summary and Comparison of Enhancement Schemes | 163 |
| 6.4 | Wireless System Evolution and TCP/IP | 164 |
| 6.4.1 | Trends in Cellular Communication Systems | 164 |
| 6.4.2 | Trends in Wireless LAN Systems | 165 |
| 6.4.3 | TCP/IP over Heterogeneous Wireless Systems | 165 |
| 6.5 | Further Reading | 166 |
| 6.6 | Summary | 167 |
| 6.7 | Review Questions | 167 |
| 6.8 | Hands-On Projects | 168 |
| 6.9 | Case Study: WCORP Installs Wireless LANs | 168 |
| 7 | TCP/IP Performance over Mobile Networks | 170 |
| 7.1 | Cellular and Ad Hoc Networks | 170 |
| 7.2 | TCP Performance in Cellular Networks | 171 |

x Contents

- 7.2.1 Mobile IP 171
- 7.2.2 Impact of Mobility on TCP Performance 173
- 7.2.3 Approaches to Improve TCP Performance 174
- 7.3 TCP Performance in Ad Hoc Networks 177
 - 7.3.1 Dynamic Source Routing 178
 - 7.3.2 Impact of Mobility on TCP Performance 178
 - 7.3.3 Approaches to Improve TCP Performance 179
- 7.4 Further Reading 181
- 7.5 Summary 182
- 7.6 Review Questions 182
- 7.7 Hands-On Projects 183
- 7.8 Case Study: WCORP Tunes TCP to Combat Blackouts in Mobile Networking 183

8 TCP/IP Performance over Optical Networks 185

- 8.1 Evolution of Optical Networks 185
- 8.2 IP over DWDM 186
- 8.3 Multiprotocol Label Switching 188
- 8.4 Multiprotocol Lambda Switching 189
- 8.5 Optical Burst Switching 190
- 8.6 Optical Packet Switching 192
 - 8.6.1 Optical Packet Format 193
 - 8.6.2 Congestion Resolution in Optical Packet Switches 195
- 8.7 Performance of TCP/IP over Optical Networks 197
 - 8.7.1 Optical Packet Network End-to-End Performance 198
 - 8.7.2 Mapping of TCP in Optical Packets 199
 - 8.7.3 Optical Packet Design in the TCP/IP Environment 201
- 8.8 Further Reading 204
- 8.9 Summary 204
- 8.10 Review Questions 204
- 8.11 Hands-On Projects 205

9 TCP/IP Performance over Satellite Networks 206

- 9.1 A Brief History of Data Satellites 206
- 9.2 Motivations for Using Satellites 207
- 9.3 Types of Satellites 208
- 9.4 Satellite Internet Architectures 209
- 9.5 Satellite Characteristics Affecting TCP 211
 - 9.5.1 Long Feedback Loop 211
 - 9.5.2 Link Impairment 214
 - 9.5.3 Bandwidth-Delay Product 215
 - 9.5.4 Bandwidth Asymmetry 216
 - 9.5.5 Variable Delays 217
 - 9.5.6 LEO Handoff 217
 - 9.5.7 Spectral Congestion 217
 - 9.5.8 Security 217

| | | |
|-----------|--|------------|
| 9.6 | Goals for TCP Performance Enhancement Solutions | 217 |
| 9.7 | TCP Enhancements for Satellite Networks | 219 |
| 9.7.1 | Path MTU Discovery | 219 |
| 9.7.2 | TCP for Transactions | 219 |
| 9.7.3 | Window Scaling | 220 |
| 9.7.4 | Large Initial Window | 221 |
| 9.7.5 | Byte Counting | 223 |
| 9.7.6 | Delayed ACKs after Slow Start | 224 |
| 9.7.7 | Explicit Congestion Notification | 225 |
| 9.7.8 | Multiple Connections | 225 |
| 9.7.9 | Pacing TCP Segments | 226 |
| 9.7.10 | TCP/IP Header Compression | 226 |
| 9.7.11 | Security Issues | 227 |
| 9.7.12 | Conclusions for TCP Enhancements | 228 |
| 9.8 | Advanced Enhancements and New Versions of TCP | 229 |
| 9.8.1 | Quick-Start TCP | 229 |
| 9.8.2 | HighSpeed TCP | 230 |
| 9.8.3 | TCP Peach | 230 |
| 9.8.4 | Explicit Transport Error Notification | 231 |
| 9.8.5 | TCP Westwood | 231 |
| 9.8.6 | XCP | 231 |
| 9.9 | New Transport Protocols for Satellite Links | 232 |
| 9.9.1 | Satellite Transport Protocol | 232 |
| 9.9.2 | Space Communications Protocol Specifications-Transport Protocol | 233 |
| 9.10 | Performance Enhancing Proxy | 233 |
| 9.10.1 | Motivations for the Use of PEP on Satellite Networks | 233 |
| 9.10.2 | Types of Performance Enhancing Proxies | 234 |
| 9.10.3 | Mechanisms Used in Performance Enhancing Proxies | 236 |
| 9.10.4 | Implications of Using Performance Enhancing Proxies | 237 |
| 9.10.5 | Security with Performance Enhancing Proxies | 238 |
| 9.10.6 | Commercial PEP (SkyX) | 239 |
| 9.11 | Further Reading | 240 |
| 9.12 | Summary | 241 |
| 9.13 | Review Questions | 242 |
| 9.14 | Hands-On Projects | 243 |
| 9.15 | Case Study: Improving TCP Performance over Satellite Using SkyX | 243 |
| 10 | TCP/IP Performance over Asymmetric Networks | 247 |
| 10.1 | Types of Network Asymmetry | 248 |
| 10.1.1 | Bandwidth Asymmetry | 248 |
| 10.1.2 | Media-Access Asymmetry | 249 |
| 10.1.3 | Loss Rate Asymmetry | 249 |
| 10.2 | Impact of Asymmetry on TCP Performance | 249 |
| 10.2.1 | Bandwidth Asymmetry | 249 |
| 10.2.2 | Media-Access Asymmetry | 251 |
| 10.3 | Improving TCP Performance over Asymmetric Networks | 255 |

xii Contents

| | | |
|-----------|---|------------|
| 10.3.1 | Uplink Bandwidth Management | 255 |
| 10.3.2 | Handling Infrequent ACKs | 257 |
| 10.4 | Experimental Evaluation of Performance Improvement Techniques | 259 |
| 10.4.1 | Experiments with Bandwidth Asymmetry | 259 |
| 10.4.2 | Experiments with Media-Access Asymmetry | 260 |
| 10.5 | Further Reading | 260 |
| 10.6 | Summary | 262 |
| 10.7 | Review Questions | 262 |
| 10.8 | Hands-On Projects | 262 |
| 10.9 | Case Study: Improving TCP Performance over ADSL | 262 |
| 11 | New TCP Standards and Flavors | 264 |
| 11.1 | Duplicate Acknowledgments and Fast Retransmit | 264 |
| 11.2 | Fast Recovery and TCP Reno | 265 |
| 11.3 | TCP NewReno | 266 |
| 11.4 | TCP with Selective Acknowledgments | 267 |
| 11.5 | Forward Acknowledgments | 268 |
| 11.6 | TCP Vegas | 268 |
| 11.7 | Overview of Other Features and Options | 269 |
| 11.8 | Performance Comparison of TCP Flavors | 270 |
| 11.9 | Further Reading | 278 |
| 11.10 | Summary | 278 |
| 11.11 | Review Questions | 278 |
| 11.12 | Hands-On Projects | 278 |
| 11.13 | Case Study: High Performance TCP for Computational Grid | 279 |
| 12 | Active Queue Management in TCP/IP Networks | 281 |
| 12.1 | Passive Queue Management | 281 |
| 12.1.1 | Tail-Drop | 282 |
| 12.1.2 | Drop-From-Front | 282 |
| 12.1.3 | Push-Out | 283 |
| 12.1.4 | Problems with Passive Queue Management | 283 |
| 12.2 | Active Queue Management | 284 |
| 12.2.1 | Random Early Detection | 284 |
| 12.2.2 | Classifying the RED Variants | 287 |
| 12.2.3 | RED Variants with Aggregate Control | 288 |
| 12.2.4 | RED Variants with Per-Flow Accounting | 294 |
| 12.3 | Performance Evaluation and Comparison of AQM Schemes | 299 |
| 12.3.1 | Throughput and Fairness | 299 |
| 12.3.2 | Delay and Jitter | 301 |
| 12.3.3 | Time Response | 301 |
| 12.3.4 | Traffic Oscillation | 302 |
| 12.3.5 | Performance Summary of AQM Schemes | 302 |
| 12.4 | AQM and Differentiated Service | 302 |
| 12.5 | Further Reading | 304 |
| 12.6 | Summary | 305 |
| 12.7 | Review Questions | 306 |

- 12.8 Hands-On Project 306
- 12.9 Case Study: WCORP Deploys Active Queue Management 307
- 13 TCP Implementation 308**
- 13.1 TCP Implementation Overview 309
 - 13.1.1 Buffering and Data Movement 311
 - 13.1.2 Accessing User Memory 311
 - 13.1.3 TCP Data Exchange 313
 - 13.1.4 Retransmissions 316
 - 13.1.5 Congestion 316
- 13.2 High Performance TCP 317
 - 13.2.1 High-Bandwidth-Delay Products 317
 - 13.2.2 Round-Trip Estimation 318
 - 13.2.3 Path MTU Discovery 318
- 13.3 Reducing End-System Overhead 319
 - 13.3.1 Overhead, CPU Utilization, and Bandwidth 320
 - 13.3.2 The Role of Application Processing 321
 - 13.3.3 Sources of Overhead for TCP/IP 322
 - 13.3.4 Per-Packet Overhead 323
 - 13.3.5 Interrupts 324
 - 13.3.6 Checksums 325
 - 13.3.7 Connection Management 326
- 13.4 Copy Avoidance 326
 - 13.4.1 Page Remapping 327
 - 13.4.2 Scatter/Gather I/O 329
 - 13.4.3 Remote Direct Memory Access 330
- 13.5 TCP Offload 331
- 13.6 Further Reading 333
- 13.7 Summary 333
- 13.8 Review Questions 333
- 13.9 Hands-On Projects 334
- A M/M/1 Queues 335**
- B FreeBSD 338**
- B.1 Installation 338
- B.2 Configuration 338
 - B.2.1 Network Card Configuration 338
 - B.2.2 Starting Network Services 339
- B.3 Kernel Modification 339
 - B.3.1 Modify Kernel Configuration Files 339
 - B.3.2 Modify Kernel Source Files 340
 - B.3.3 Build and Install the New Kernel 341
 - B.3.4 Kernel Error Recovery Technique 341
 - B.3.5 Kernel Modification Example 342
- C TCP Auto-Tuning 345**
- C.1 Motivation for TCP Auto-Tuning 345

xiv Contents

| | | |
|-----|---|------------|
| C.2 | TCP Auto-Tuning Techniques and Products | 345 |
| C.3 | Selecting Auto-Tuned TCP | 346 |
| C.4 | Further Reading | 347 |
| | Bibliography | 348 |
| | Index | 367 |

Preface

The world is undergoing a revolution in information and communication technology. Not only the lives of citizens but also the networking technology are profoundly affected by this revolution. Traditional wired networks are being replaced or complemented by networks based on wireless, optical, satellite, and other media. TCP/IP has emerged as the global Internet-working solution allowing communication over a wide variety of media and networks. These new networking media and the new ways of communication over these networks have given rise to a host of new performance issues and concepts. To adapt and contribute effectively to such changes, engineers and computer scientists must acquire a solid foundation and understanding of the fundamental concepts that affect performance in TCP/IP networks.

Existing texts on TCP/IP focus on the presentation of the protocol details with little coverage of the performance issues and concepts. These texts are good for a first course on TCP/IP networking but do not provide sufficient material for those advanced readers interested in acquiring in-depth knowledge of the performance aspects of TCP/IP, especially in the emerging networking environment. To address this need, we have written *High Performance TCP/IP Networking: Concepts, Issues, and Solutions*, with a clear focus on the performance fundamentals of TCP/IP.

High Performance TCP/IP Networking: Concepts, Issues, and Solutions is a comprehensive guide to the study of its topic. Our book provides an in-depth coverage of (1) tools and techniques for the performance evaluation of TCP/IP networks, (2) performance concepts and issues for running TCP/IP over wireless, mobile, optical, and satellite networks, (3) congestion-control algorithms in hosts and routers, and (4) high performance implementation of TCP/IP protocol stack. This text has been created with an emphasis on fundamental concepts, such as network measurement and simulation techniques, mathematical modeling of TCP dynamics, and management of implementation overhead, which will continue to guide new developments in TCP/IP. Although many specific networks, tools, and protocols are discussed in the text, a continuous effort has been made to emphasize the underlying performance issues and concepts.

CONTRIBUTING AUTHORS

This book contains contributions from many leading experts actively working on specific performance issues in TCP/IP networks. In addition to the two editors (Hassan and Jain), who themselves wrote parts of the book, there are a total of 24 authors who wrote specific chapters of the book. One of the most challenging tasks was to integrate these individual submissions into a coherent book. As part of the integration effort, the editors have introduced a range of additional materials, including learning objectives, review questions, hands-on projects, and case studies. The editors maintained close liaison with the chapter authors throughout the manuscript preparation process. The manuscript was reviewed and revised twice to address the concerns of the reviewers. Substantial material was added in each revision to further integrate the chapters and improve the quality of the book. While most individual chapter authors were contacted for the revisions, the editors

xvi Preface

themselves revised some of the chapters. The list of authors who contributed to this book follows (chapters that were substantially revised by the editors are marked with an asterisk).

Chapter 1 Introduction

Mahbub Hassan, University of New South Wales, Australia
Raj Jain, Ohio State University, USA

Chapter 2 TCP/IP Fundamentals*

Sanjay Jha, University of New South Wales, Australia

Chapter 3 Performance Measurement of TCP/IP Networks*

Yukio Murayama, Kurashilki University of Science and Arts, Japan
Suguru Yamaguchi, Nara Institute of Science and Technology, Japan

Chapter 4 TCP/IP Network Simulation

Mahbub Hassan, University of New South Wales, Australia
Sonia Fahmy, Purdue University, USA
Jim Wu, University of New South Wales, Australia
Abdul Aziz, University of New South Wales, Australia

Chapter 5 TCP Modeling

Sven Östring, University of Cambridge, United Kingdom
Harsha Sirisena, University of Canterbury, New Zealand

Chapter 6 TCP/IP Performance over Wireless Networks

George Xylomenos, Athens University of Economics and Business, Greece
George Polyzos, Athens University of Economics and Business, Greece
Petri Mähönen, Aachen University, Germany
Mika Saarinen, Nokia Mobile Phones, Finland

Chapter 7 TCP/IP Performance over Mobile Networks

Raghupathy Sivakumar, Georgia Institute of Technology, USA

Chapter 8 TCP/IP Performance over Optical Networks

Franco Callegati, Universita' di Bologna, Italy
Maurizio Casoni, Universita' di Modena and Reggio Emilia, Italy
Carla Raffaelli, Universita' di Bologna, Italy

Chapter 9 TCP/IP Performance over Satellite Networks*

Arjan Durrezi, Louisiana State University, USA
Sastri Kota, Loral Skynet, USA

Chapter 10 TCP/IP Performance over Asymmetric Networks

Venkat Padmanabhan, Microsoft Research, USA
Hari Balakrishnan, Massachusetts Institute of Technology, USA

Chapter 11 New TCP Standards and Flavors

Sonia Fahmy, Purdue University, USA

Chapter 12 Active Queue Management in TCP/IP Networks

Mohammed Atiquzzaman, University of Oklahoma, USA
Bing Zheng, New Focus, Inc., USA

Chapter 13 Software Implementation of TCP

Jeff Chase, Duke University, USA

Appendix A M/M/1 Queue

Mahbub Hassan, University of New South Wales, Australia
Raj Jain, Ohio State University, USA

Appendix B FreeBSD

Rui Zhao, University of New South Wales, Australia

Appendix C TCP Auto-Tuning

Mahbub Hassan, University of New South Wales, Australia

ORGANIZATION AND OUTLINE

The book is organized into five parts.

- **Part I: Background.** Part I provides an introduction to the book. It contains two chapters. Chapter 1 provides a rationale for the book. Chapter 2 reviews some of the key features of TCP/IP protocols used in later chapters in the book to explain many performance issues. Chapter 2 reviews only the key features of TCP/IP. A comprehensive treatment of TCP/IP protocol stack is beyond the scope of the book.
- **Part II: Performance Evaluation.** Part II consists of Chapters 3, 4, and 5 and provides detailed coverage of the tools and techniques for performance evaluation of TCP/IP networks. Chapter 3 discusses the performance measurement tools available for monitoring, analyzing, and benchmarking the performance of TCP/IP networks. Chapter 4 introduces simulation techniques and discusses two popular simulation tools. Chapter 5 is devoted to the mathematical modeling of TCP congestion control algorithms.
- **Part III: Performance in Emerging Networks.** Chapters 6 through 10 examine the performance concepts and issues for running TCP/IP in the emerging networking environment. Although many of us think modems and Ethernet when we think Internet and TCP/IP, this is no longer the reality. Yes, it is true that nearly every home has a modem for Internet connection, and nearly every organization has some version of the wired Ethernet connectivity to the desktop. Many homes, however, are subscribing to Digital Subscriber Loop (DSL) technologies for high-speed Internet connection, and many organizations are deploying wireless LANs for flexibility. In the wide area, too, we are witnessing new networking technologies such as mobile cellular data networks (e.g., GPRS), high-speed optical backbones, and increasing use of satellite links for long-distance and global coverage. Each of these new technologies has given rise to some new concepts and issues for TCP/IP performance. We have therefore dedicated a separate chapter to deal with each of these technologies: Chapter 6 for wireless, Chapter 7 for mobility, Chapter 8 for optical, Chapter 9 for satellite, and Chapter 10 for asymmetric networks (e.g., ADSL).
- **Part IV: Congestion Control.** With the increase in networking complexities and traffic dynamics, congestion-control algorithms employed at the end hosts and in the network routers continue to evolve. The new congestion-control algorithms in the TCP protocol resulted in many different TCP flavors (e.g.,

xviii Preface

Tahoe, Reno, Vegas, and so on). Part IV consists of Chapters 11 and 12 and presents an in-depth coverage of the congestion-control algorithms proposed so far. Chapter 11 discusses various TCP flavors, and Chapter 12 examines the new queue management schemes proposed for the network routers to combat congestion in highly dynamic environment.

- **Part V: Implementation.** For emerging high-speed networks (e.g., 10 Gbps Ethernet), the end-system implementation of TCP can become a performance bottleneck. Part V (Chapter 13) summarizes critical performance issues for TCP implementation in end systems and surveys solutions for improving bulk transfer performance.

HOW TO USE THIS BOOK

The book is designed for use in a second course on networking with a prerequisite course on introductory networking or data communications. Some of the possible courses for which this book can be used include Advanced Computer Networks, Advanced TCP/IP Networks, High Performance Networks, and Internet-working. There is enough material in the book for a one-semester or one-quarter course with 12 or 13 weeks of lecture. Depending on the background of the students, two possible course compositions are given here.

Computer science students with limited background in mathematics and hardware design can exclude Chapter 5 (mathematical modeling) and study Chapter 4 (Simulation) in more detail. Engineering students graduating in computer or electrical engineering can spend fewer weeks on Chapter 4 and one extra week on Chapter 5.

Professionals working as network engineers, R & D managers, research scientists, and network administrators will also find this book valuable as a reference to the most recent advances in TCP performance research.

LEARNING AIDS

There are many learning aids in this book:

- **Learning Objectives.** Each chapter starts with a list of learning objectives. The learning objectives highlight the fundamental concepts (skills) students should understand (master) as a result of reading the chapter and help them organize their study goals. They assist instructors in pointing out lecture objectives.
- **Further Reading Lists.** Annotated reading lists at the end of the chapters provide students with valuable resources for independent exploration on specific topics of interest. These lists are particularly useful for professionals.
- **Chapter Summaries.** Summaries offer students a chance to review their understanding of key concepts in the chapter before moving on.
- **Review Questions.** End-of-chapter review questions evaluate the degree to which the student achieved the learning objectives and force the students to think about the key concepts in the chapter. Answers to most of the review questions can be found directly from the chapter; therefore, students are

forced to reread parts of the chapter to locate the answers. Such rereading is often required to gain a clear understanding of many difficult concepts. The instructor can use some of these questions for classroom discussions or class tests.

- **Hands-On Projects.** For each chapter, a list of performance evaluation experiments are provided for advanced students seeking to gain a deeper understanding of some of the key concepts and solutions described in the chapter. These experiments can be carried out on open platforms using freely available software. The hands-on projects in this book cover a range of difficulty. Some experiments can be completed in a few weeks using *ns-2* simulation software, without requiring any kernel-level programming. Other experiments require modification of existing TCP/IP stacks in FreeBSD operating system kernel. These experiments are quite challenging and can be given to students as whole semester projects. Students attempting these experiments are expected to have a good background in programming and operating systems. (Appendix B provides a brief tutorial on FreeBSD for students with no prior background in kernel programming.)
- **Case Studies.** A case study is introduced in Chapter 1 based on a fictitious, but realistic organization with TCP/IP networking infrastructure. The same case study is then used in subsequent chapters with some modifications to introduce new performance problems. The running case study holds together different chapters in the text, provides students a realistic context in which to apply the concepts and techniques learned in the relevant chapters, and yields a classroom discussion topic for the instructor.
- **Figures and Illustrations.** Many concepts throughout the book are explained using illustrations. These illustrations help students understand complex performance issues and concepts.
- **Examples.** Examples have been used where applicable to explain the use of techniques learned from the text.

ACKNOWLEDGMENTS

The book would never exist without the contributions from the individual chapter authors. We take this opportunity to thank all chapter authors for their expertise and time and for putting up with our many requests throughout the preparation of the manuscript. We are indebted to the anonymous reviewers for reading the whole manuscript or part of the earlier versions of the manuscript and making useful comments. Their constructive suggestions significantly influenced the revisions of the manuscript. We thank Professor Krzysztof Pawlikowski of the University of Canterbury, New Zealand, for providing early feedback on Chapter 4. The author of Chapter 11 thanks Tapan Karwa, Venkatesh Prabhakar, Farnaz Erfan, and Minseok Kwon for their help with the simulation experiments in that chapter. Jim Wu, a coauthor of Chapter 4, has been instrumental in fixing some of the problems we faced in preparing the manuscript in LaTeX. We gratefully acknowledge the support of the entire production team at Prentice Hall. Finally, the first editor (Hassan)

xx Preface

would like to thank Professor Arun Sharma (previous head of school) and Professor Paul Compton (current head of school) at the University of New South Wales for providing a pleasant and stimulating environment in which to work.

Mahbub Hassan
Raj Jain