Computer Networks

Tanenbaum • Feamster • Wetherall

SIXTH EDITION

GLOBAL
EDITION

Computer Networks

Tanenbaum • Feamster • Wetherall

SIXTH EDITION

COMPUTER NETWORKS

SIXTH EDITION

This page is intentionally left blank

COMPUTER NETWORKS

SIXTH EDITION

Global Edition

ANDREW S. TANENBAUM
Vrije Universiteit
Amsterdam, The Netherlands

NICK FEAMSTER
University of Chicago
Chicago, IL

DAVID WETHERALL
Google

Please contact https://support.pearson.com/getsupport/s/contactsupport with any queries on this content.

Pearson Education Limited
KAO Two
KAO Park
Hockham Way
Harlow
CM17 9SR
United Kingdom

and Associated Companies throughout the world

Visit us on the World Wide Web at: www.pearsonglobaleditions.com

© Pearson Education Limited, 2021

The rights of Andrew S. Tanenbaum, Nick Feamster, and David Wetherall to be identified as the authors of this work have been asserted by them in accordance with the Copyright, Designs and Patents Act 1988.

Authorized adaptation from the United States edition, entitled Computer Networks, 6th Edition*, ISBN 978-0-13-676405-2 by Andrew S. Tanenbaum, Nick Feamster, and David Wetherall, published by Pearson Education © 2020.*

All rights reserved. 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, recording or otherwise, without either the prior written permission of the publisher or a license permitting restricted copying in the United Kingdom issued by the Copyright Licensing Agency Ltd, Saffron House, 6–10 Kirby Street, London EC 1N 8TS.

All trademarks used herein are the property of their respective owners. The use of any trademark in this text does not vest in the author or publisher any trademark ownership rights in such trademarks, nor does the use of such trademarks imply any affiliation with or endorsement of this book by such owners. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights and Permissions department, please visit www.pearsoned.com/permissions.

This eBook is a standalone product and may or may not include all assets that were part of the print version. It also does not provide access to other Pearson digital products like MyLab and Mastering. The publisher reserves the right to remove any material in this eBook at any time.

British Library Cataloguing-in-Publication Data
A catalogue record for this book is available from the British Library

ISBN 10: 1-292-37406-3
ISBN 13: 978-1-292-37406-2
eBook ISBN 13: 9781292374017

To Suzanne, Barbara, Daniel, Aron, Nathan, Marvin, Matilde, Olivia, and Mirte (AST)

To Marshini, Mila, and Kira (NF)

To Katrin, Lucy, and Pepper (DJW)

This page is intentionally left blank

CONTENTS

PREFACE **xix**

1 INTRODUCTION **1**

1.1 USES OF COMPUTER NETWORKS 1
- 1.1.1 Access to Information 2
- 1.1.2 Person-to-Person Communication 5
- 1.1.3 Electronic Commerce 6
- 1.1.4 Entertainment 6
- 1.1.5 The Internet of Things 7

1.2 TYPES OF COMPUTER NETWORKS 7
- 1.2.1 Broadband Access Networks 8
- 1.2.2 Mobile and Wireless Access Networks 8
- 1.2.3 Content Provider Networks 11
- 1.2.4 Transit Networks 12
- 1.2.5 Enterprise Networks 13

1.3 NETWORK TECHNOLOGY, FROM LOCAL TO GLOBAL 15
- 1.3.1 Personal Area Networks 15
- 1.3.2 Local Area Networks 16
- 1.3.3 Home Networks 18
- 1.3.4 Metropolitan Area Networks 20
- 1.3.5 Wide Area Networks 21
- 1.3.6 Internetworks 25

1.4 EXAMPLES OF NETWORKS 26
1.4.1 The Internet 26
1.4.2 Mobile Networks 36
1.4.3 Wireless Networks (WiFi) 43

1.5 NETWORK PROTOCOLS 47
1.5.1 Design Goals 47
1.5.2 Protocol Layering 49
1.5.3 Connections and Reliability 53
1.5.4 Service Primitives 56
1.5.5 The Relationship of Services to Protocols 58

1.6 REFERENCE MODELS 59
1.6.1 The OSI Reference Model 59
1.6.2 The TCP/IP Reference Model 61
1.6.3 A Critique of the OSI Model and Protocols 64
1.6.4 A Critique of the TCP/IP Reference Model and Protocols 66
1.6.5 The Model Used in This Book 67

1.7 STANDARDIZATION 68
1.7.1 Standardization and Open Source 68
1.7.2 Who's Who in the Telecommunications World 69
1.7.3 Who's Who in the International Standards World 71
1.7.4 Who's Who in the Internet Standards World 72

1.8 POLICY, LEGAL, AND SOCIAL ISSUES 75
1.8.1 Online Speech 75
1.8.2 Net Neutrality 76
1.8.3 Security 77
1.8.4 Privacy 78
1.8.5 Disinformation 79

1.9 METRIC UNITS 80

1.10 OUTLINE OF THE REST OF THE BOOK 81

1.11 SUMMARY 82

2 THE PHYSICAL LAYER 89

2.1 GUIDED TRANSMISSION MEDIA 90
- 2.1.1 Persistent Storage 90
- 2.1.2 Twisted Pairs 91
- 2.1.3 Coaxial Cable 93
- 2.1.4 Power Lines 94
- 2.1.5 Fiber Optics 95

2.2 WIRELESS TRANSMISSION 100
- 2.2.1 The Electromagnetic Spectrum 101
- 2.2.2 Frequency Hopping Spread Spectrum 103
- 2.2.3 Direct Sequence Spread Spectrum 103
- 2.2.4 Ultra-Wideband Communication 104

2.3 USING THE SPECTRUM FOR TRANSMISSION 104
- 2.3.1 Radio Transmission 104
- 2.3.2 Microwave Transmission 106
- 2.3.3 Infrared Transmission 107
- 2.3.4 Light Transmission 108

2.4 FROM WAVEFORMS TO BITS 109
- 2.4.1 The Theoretical Basis for Data Communication 110
- 2.4.2 The Maximum Data Rate of a Channel 114
- 2.4.3 Digital Modulation 115
- 2.4.4 Multiplexing 123

2.5 THE PUBLIC SWITCHED TELEPHONE NETWORK 131
- 2.5.1 Structure of the Telephone System 131
- 2.5.2 The Local Loop: Telephone Modems, ADSL, and Fiber 134
- 2.5.3 Trunks and Multiplexing 143
- 2.5.4 Switching 149

2.6 CELLULAR NETWORKS 154
- 2.6.1 Common Concepts: Cells, Handoff, Paging 155
- 2.6.2 First-Generation (1G) Technology: Analog Voice 156
- 2.6.3 Second-Generation (2G) Technology: Digital Voice 158
- 2.6.4 GSM: The Global System for Mobile Communications 159
- 2.6.5 Third-Generation (3G) Technology: Digital Voice and Data 162
- 2.6.6 Fourth-Generation (4G) Technology: Packet Switching 166
- 2.6.7 Fifth-Generation (5G) Technology 168

2.7 CABLE NETWORKS 169
2.7.1 A History of Cable Networks: Community Antenna Television 170
2.7.2 Broadband Internet Access Over Cable: HFC Networks 170
2.7.3 DOCSIS 173
2.7.4 Resource Sharing in DOCSIS Networks: Nodes and Minislots 174

2.8 COMMUNICATION SATELLITES 176
2.8.1 Geostationary Satellites 177
2.8.2 Medium-Earth Orbit Satellites 181
2.8.3 Low-Earth Orbit Satellites 181

2.9 COMPARING DIFFERENT ACCESS NETWORKS 184
2.9.1 Terrestrial Access Networks: Cable, Fiber, and ADSL 184
2.9.2 Satellites Versus Terrestrial Networks 186

2.10 POLICY AT THE PHYSICAL LAYER 187
2.10.1 Spectrum Allocation 187
2.10.2 The Cellular Network 190
2.10.3 The Telephone Network 192

2.11 SUMMARY 194

3 THE DATA LINK LAYER 201

3.1 DATA LINK LAYER DESIGN ISSUES 202
3.1.1 Services Provided to the Network Layer 203
3.1.2 Framing 205
3.1.3 Error Control 208
3.1.4 Flow Control 209

3.2 ERROR DETECTION AND CORRECTION 210
3.2.1 Error-Correcting Codes 212
3.2.2 Error-Detecting Codes 217

3.3 ELEMENTARY DATA LINK PROTOCOLS 223
3.3.1 Initial Simplifying Assumptions 223
3.3.2 Basic Transmission and Receipt 224
3.3.3 Simplex Link-Layer Protocols 228

3.4 IMPROVING EFFICIENCY 234
3.4.1 Goal: Bidirectional Transmission, Multiple Frames in Flight 234
3.4.2 Examples of Full-Duplex, Sliding Window Protocols 238

3.5 DATA LINK PROTOCOLS IN PRACTICE 252
3.5.1 Packet over SONET 253
3.5.2 ADSL (Asymmetric Digital Subscriber Loop) 256
3.5.3 Data Over Cable Service Interface Specification (DOCSIS) 259

3.6 SUMMARY 261

4 THE MEDIUM ACCESS CONTROL SUBLAYER 267

4.1 THE CHANNEL ALLOCATION PROBLEM 268
4.1.1 Static Channel Allocation 268
4.1.2 Assumptions for Dynamic Channel Allocation 270

4.2 MULTIPLE ACCESS PROTOCOLS 271
4.2.1 ALOHA 272
4.2.2 Carrier Sense Multiple Access Protocols 276
4.2.3 Collision-Free Protocols 279
4.2.4 Limited-Contention Protocols 283
4.2.5 Wireless LAN Protocols 287

4.3 ETHERNET 290
4.3.1 Classic Ethernet Physical Layer 290
4.3.2 Classic Ethernet MAC Sublayer Protocol 292
4.3.3 Ethernet Performance 296
4.3.4 Switched Ethernet 297
4.3.5 Fast Ethernet 300
4.3.6 Gigabit Ethernet 302
4.3.7 10-Gigabit Ethernet 306
4.3.8 40- and 100-Gigabit Ethernet 307
4.3.9 Retrospective on Ethernet 308

4.4 WIRELESS LANS 309
4.4.1 The 802.11 Architecture and Protocol Stack 310
4.4.2 The 802.11 Physical Layer 311

4.4.3 The 802.11 MAC Sublayer Protocol 314
4.4.4 The 802.11 Frame Structure 321
4.4.5 Services 322

4.5 BLUETOOTH 324
4.5.1 Bluetooth Architecture 325
4.5.2 Bluetooth Applications 326
4.5.3 The Bluetooth Protocol Stack 327
4.5.4 The Bluetooth Radio Layer 328
4.5.5 The Bluetooth Link Layers 329
4.5.6 The Bluetooth Frame Structure 330
4.5.7 Bluetooth 5 331

4.6 DOCSIS 332
4.6.1 Overview 332
4.6.2 Ranging 333
4.6.3 Channel Bandwidth Allocation 333

4.7 DATA LINK LAYER SWITCHING 334
4.7.1 Uses of Bridges 335
4.7.2 Learning Bridges 336
4.7.3 Spanning-Tree Bridges 339
4.7.4 Repeaters, Hubs, Bridges, Switches, Routers, and Gateways 342
4.7.5 Virtual LANs 345

4.8 SUMMARY 351

5 THE NETWORK LAYER 359

5.1 NETWORK LAYER DESIGN ISSUES 360
5.1.1 Store-and-Forward Packet Switching 360
5.1.2 Services Provided to the Transport Layer 361
5.1.3 Implementation of Connectionless Service 362
5.1.4 Implementation of Connection-Oriented Service 363
5.1.5 Comparison of Virtual-Circuit and Datagram Networks 365

5.2 ROUTING ALGORITHMS IN A SINGLE NETWORK 366
5.2.1 The Optimality Principle 368
5.2.2 Shortest Path Algorithm 370

5.2.3 Flooding 372
5.2.4 Distance Vector Routing 374
5.2.5 Link State Routing 377
5.2.6 Hierarchical Routing within a Network 382
5.2.7 Broadcast Routing 384
5.2.8 Multicast Routing 386
5.2.9 Anycast Routing 389

5.3 TRAFFIC MANAGEMENT AT THE NETWORK LAYER 390
5.3.1 The Need for Traffic Management: Congestion 390
5.3.2 Approaches to Traffic Management 393

5.4 QUALITY OF SERVICE AND APPLICATION QOE 406
5.4.1 Application QoS Requirements 406
5.4.2 Overprovisioning 409
5.4.3 Packet Scheduling 410
5.4.4 Integrated Services 417
5.4.5 Differentiated Services 420

5.5 INTERNETWORKING 423
5.5.1 Internetworks: An Overview 423
5.5.2 How Networks Differ 424
5.5.3 Connecting Heterogeneous Networks 425
5.5.4 Connecting Endpoints Across Heterogeneous Networks 428
5.5.5 Internetwork Routing: Routing Across Multiple Networks 430
5.5.6 Supporting Different Packet Sizes: Packet Fragmentation 431

5.6 SOFTWARE-DEFINED NETWORKING 435
5.6.1 Overview 435
5.6.2 The SDN Control Plane: Logically Centralized Software Control 436
5.6.3 The SDN Data Plane: Programmable Hardware 438
5.6.4 Programmable Network Telemetry 440

5.7 THE NETWORK LAYER IN THE INTERNET 441
5.7.1 The IP Version 4 Protocol 444
5.7.2 IP Addresses 448
5.7.3 IP Version 6 461
5.7.4 Internet Control Protocols 470
5.7.5 Label Switching and MPLS 476
5.7.6 OSPF—An Interior Gateway Routing Protocol 479
5.7.7 BGP—The Exterior Gateway Routing Protocol 484
5.7.8 Internet Multicasting 491

5.8 POLICY AT THE NETWORK LAYER 492
5.8.1 Peering Disputes 492
5.8.2 Traffic Prioritization 493

5.9 SUMMARY 494

6 THE TRANSPORT LAYER 501

6.1 THE TRANSPORT SERVICE 501
6.1.1 Services Provided to the Upper Layers 502
6.1.2 Transport Service Primitives 504
6.1.3 Berkeley Sockets 506
6.1.4 An Example of Socket Programming: An Internet File Server 509

6.2 ELEMENTS OF TRANSPORT PROTOCOLS 513
6.2.1 Addressing 514
6.2.2 Connection Establishment 517
6.2.3 Connection Release 523
6.2.4 Error Control and Flow Control 528
6.2.5 Multiplexing 533
6.2.6 Crash Recovery 533

6.3 CONGESTION CONTROL 536
6.3.1 Desirable Bandwidth Allocation 536
6.3.2 Regulating the Sending Rate 540
6.3.3 Wireless Issues 544

6.4 THE INTERNET TRANSPORT PROTOCOLS: UDP 546
6.4.1 Introduction to UDP 547
6.4.2 Remote Procedure Call 549
6.4.3 Real-Time Transport Protocols 552

6.5 THE INTERNET TRANSPORT PROTOCOLS: TCP 557
6.5.1 Introduction to TCP 558
6.5.2 The TCP Service Model 558
6.5.3 The TCP Protocol 561
6.5.4 The TCP Segment Header 562
6.5.5 TCP Connection Establishment 565
6.5.6 TCP Connection Release 567

6.5.7 TCP Connection Management Modeling 567
6.5.8 TCP Sliding Window 570
6.5.9 TCP Timer Management 573
6.5.10 TCP Congestion Control 576
6.5.11 TCP CUBIC 586

6.6 TRANSPORT PROTOCOLS AND CONGESTION CONTROL 587
6.6.1 QUIC: Quick UDP Internet Connections 587
6.6.2 BBR: Congestion Control Based on Bottleneck Bandwidth 588
6.6.3 The Future of TCP 590

6.7 PERFORMANCE ISSUES 590
6.7.1 Performance Problems in Computer Networks 591
6.7.2 Network Performance Measurement 592
6.7.3 Measuring Access Network Throughput 593
6.7.4 Measuring Quality of Experience 594
6.7.5 Host Design for Fast Networks 595
6.7.6 Fast Segment Processing 598
6.7.7 Header Compression 601
6.7.8 Protocols for Long Fat Networks 603

6.8 SUMMARY 607

7 THE APPLICATION LAYER 613

7.1 THE DOMAIN NAME SYSTEM (DNS) 613
7.1.1 History and Overview 614
7.1.2 The DNS Lookup Process 614
7.1.3 The DNS Name Space and Hierarchy 617
7.1.4 DNS Queries and Responses 620
7.1.5 Name Resolution 627
7.1.6 Hands on with DNS 629
7.1.7 DNS Privacy 629
7.1.8 Contention Over Names 631

7.2 ELECTRONIC MAIL 632
7.2.1 Architecture and Services 633
7.2.2 The User Agent 635
7.2.3 Message Formats 637

7.2.4 Message Transfer 642
7.2.5 Final Delivery 647

7.3 THE WORLD WIDE WEB 650
7.3.1 Architectural Overview 651
7.3.2 Static Web Objects 659
7.3.3 Dynamic Web Pages and Web Applications 660
7.3.4 HTTP and HTTPS 664
7.3.5 Web Privacy 676

7.4 STREAMING AUDIO AND VIDEO 680
7.4.1 Digital Audio 682
7.4.2 Digital Video 684
7.4.3 Streaming Stored Media 687
7.4.4 Real-Time Streaming 694

7.5 CONTENT DELIVERY 703
7.5.1 Content and Internet Traffic 705
7.5.2 Server Farms and Web Proxies 707
7.5.3 Content Delivery Networks 711
7.5.4 Peer-to-Peer Networks 715
7.5.5 Evolution of the Internet 721

7.6 SUMMARY 725

8 NETWORK SECURITY 731

8.1 FUNDAMENTALS OF NETWORK SECURITY 733
8.1.1 Fundamental Security Principles 734
8.1.2 Fundamental Attack Principles 736
8.1.3 From Threats to Solutions 738

8.2 THE CORE INGREDIENTS OF AN ATTACK 739
8.2.1 Reconnaissance 739
8.2.2 Sniffing and Snooping (with a Dash of Spoofing) 742
8.2.3 Spoofing (beyond ARP) 744
8.2.4 Disruption 755

8.3 FIREWALLS AND INTRUSION DETECTION SYSTEMS 759
8.3.1 Firewalls 760
8.3.2 Intrusion Detection and Prevention 762

8.4 CRYPTOGRAPHY 766
8.4.1 Introduction to Cryptography 767
8.4.2 Two Fundamental Cryptographic Principles 769
8.4.3 Substitution Ciphers 771
8.4.4 Transposition Ciphers 773
8.4.5 One-Time Pads 774

8.5 SYMMETRIC-KEY ALGORITHMS 779
8.5.1 The Data Encryption Standard 780
8.5.2 The Advanced Encryption Standard 781
8.5.3 Cipher Modes 783

8.6 PUBLIC-KEY ALGORITHMS 787
8.6.1 RSA 788
8.6.2 Other Public-Key Algorithms 790

8.7 DIGITAL SIGNATURES 791
8.7.1 Symmetric-Key Signatures 791
8.7.2 Public-Key Signatures 793
8.7.3 Message Digests 795
8.7.4 The Birthday Attack 797

8.8 MANAGEMENT OF PUBLIC KEYS 799
8.8.1 Certificates 799
8.8.2 X.509 801
8.8.3 Public Key Infrastructures 802

8.9 AUTHENTICATION PROTOCOLS 805
8.9.1 Authentication Based on a Shared Secret Key 806
8.9.2 Establishing a Shared Key: The Diffie-Hellman Key Exchange 811
8.9.3 Authentication Using a Key Distribution Center 813
8.9.4 Authentication Using Kerberos 816
8.9.5 Authentication Using Public-Key Cryptography 819

8.10 COMMUNICATION SECURITY 819
8.10.1 IPsec 820
8.10.2 Virtual Private Networks 824
8.10.3 Wireless Security 825

8.11 EMAIL SECURITY 829
8.11.1 Pretty Good Privacy 829
8.11.2 S/MIME 833

8.12 WEB SECURITY 834
8.12.1 Threats 834
8.12.2 Secure Naming and DNSSEC 835
8.12.3 Transport Layer Security 838
8.12.4 Running Untrusted Code 842

8.13 SOCIAL ISSUES 844
8.13.1 Confidential and Anonymous Communication 844
8.13.2 Freedom of Speech 847
8.13.3 Copyright 851

8.14 SUMMARY 854

9 READING LIST AND BIBLIOGRAPHY 863

9.1 SUGGESTIONS FOR FURTHER READING 863
9.1.1 Introduction and General Works 864
9.1.2 The Physical Layer 865
9.1.3 The Data Link Layer 866
9.1.4 The Medium Access Control Sublayer 867
9.1.5 The Network Layer 868
9.1.6 The Transport Layer 869
9.1.7 The Application Layer 870
9.1.8 Network Security 871

9.2 ALPHABETICAL BIBLIOGRAPHY 872

INDEX 891

PREFACE

This book is now in its sixth edition. Each edition has corresponded to a different phase in the way computer networks were used. When the first edition appeared in 1980, networks were an academic curiosity. When the second edition appeared in 1988, networks were used by universities and large businesses. When the third edition appeared in 1996, computer networks, especially the Internet, had become a daily reality for millions of people. By the fourth edition, in 2003, wireless networks and mobile computers had become commonplace for accessing the Web and the Internet. By the fifth edition, networks were about content distribution (especially videos using CDNs and peer-to-peer networks) and mobile phones. Now in the sixth edition, industry emphasis on is very high performance, with 5G cellular networks, 100-gigabit Ethernet, and 802.11ax WiFi at speeds up to 11 Gbps just around the corner.

New in the Sixth Edition

Among the many changes in this book, the most important one is the addition of Prof. Nick Feamster as a co-author. Prof. Feamster has a Ph.D. from M.I.T. and is now a full professor at the University of Chicago.

Another important change is that Chapter 8 (on security) has been very heavily modified by Prof. Herbert Bos of the Vrije Universiteit in Amsterdam. The focus has moved from cryptography to network security. The issues of hacking, DoS attacks and so much more is front-and-center in the news almost every day, so we are very grateful that Prof. Bos has redone the chapter to deal with these important issues in detail. The chapter discusses vulnerabilities, how to fix them, how hackers respond to the fixes, how the defenders react, and so on ad infinitum. The material on cryptography has been reduced somewhat to make room for the large amount of new material on network security.

Of course, the book also has many other changes to keep up with the ever-changing world of computer networks. A chapter-by-chapter list of the major changes follows.

Chapter 1 serves the same introductory function as in previous editions, but the contents have been revised and brought up to date. Specific updates including adding additional discussions on the Internet of Things and modern cellular architectures, including 4G and 5G networks. Much of the discussion on Internet policy has also been updated, particularly the discussion on net neutrality.

Chapter 2 has been updated to include discussion of more prevalent physical media in access networks, such as DOCSIS and fiber arhictectures. Treatment of modern cellular network architectures and technologies was added, and the section on satellite networks was also substantially updated. Emerging technologies such as virtualization were added, including discussions on mobile virtual network operators and cellular network slicing. The policy discussion was reorganized and updated to include discussion on policy questions in the wireless arena, such as spectrum.

Chapter 3 has been updated to include DOCSIS as a protocol example, as it is a widely used access technology. Much of the error correction codes are, of course, timeless.

Chapter 4 has been brought up to date, with new material on 40- and 100-gigabit Ethernet, 802.11.ac, 802.11ad, and 802.11ax. New material has been added on DOCSIS, explaining the MAC sublayer in cable networks. The material on 802.16 has been removed as it now appears that this technology is going to lose out to the cellular 4G and 5G technologies. The section on RFID has also been removed to make space for new material, but also because it was not directly network related.

Chapter 5 has been updated to clarify and modernize the discussions on congestion management. The sections on traffic management have been updated and clarified, and the discussions on traffic shaping and traffic engineering have been updated. The chapter includes an entirely new section on software-defined networking (SDN), including OpenFlow and programmable hardware (e.g., Tofino). The chapter also includes discussion on emerging applications of SDN, such as in-band network telemetry. Some of the discussion on IPv6 has also been updated.

Chapter 6 has been extensively edited to include new material on modern transport protocols, including TCP CUBIC, QUIC, and BBR. The material on performance measurement has been completely rewritten to focus on the measurement of throughput in computer networks, including an extensive discussion on the challenges of measuring access network throughout as speeds in access ISPs increase. The chapter also includes new material on measuring user quality of experience, an emerging area in performance measurement.

Chapter 7 has been heavily edited. Over 60 pages of material that is no longer relevant to a book on computer networks has been removed. The material on DNS has been almost completely rewritten to reflect modern developments in DNS, including the ongoing trends to encrypt DNS and generally improve its privacy characteristics. Emerging protocols such as DNS-over-HTTPS and other privacy-preserving techniques for DNS are discussed. The discussion of the Web has been extensively updated, to reflect the increasing deployment of encryption on the Web,

as well as extensive privacy issues (e.g., tracking) that are now pervasive on the Web. The chapter includes a completely new section on Web privacy, more extensive discussions of modern content delivery technology (e.g., content delivery networks), and an expanded discussion on peer-to-peer networks. The section on the evolution of the Internet has also been edited to reflect trends towards distributed cloud services.

Chapter 8 has been completely overhauled. In previous editions, the focus of the security chapter was almost exclusively on information security by means of cryptography. However, cryptography is only one aspect of network security and if we look at security incidents in practice, it is generally not the aspect where the problems are. To remedy this, we added new content on security principles, fundamental attack techniques, defenses, and a wide range of systems-related security issues. Moreover, we updated the existing sections by dropping some encryption techniques that are now obsolete and introducing more modern versions of protocols and standards.

Chapter 9 contains a renewed list of suggested readings and a comprehensive bibliography.

In addition, dozens of new exercises and dozens of new references have been added.

List of Acronyms

Computer books are full of acronyms. This one is no exception. By the time you are completely finished reading this one, the following should ring a bell: AES, AMI, ARP, ARQ, ASK, BGP, BSC, CCK, CDM, CDN, CRL, DCF, DES, DIS, DMT, DMZ, DNS, EAP, ECN, EDE, EPC, FDD, FDM, FEC, FSK, GEO, GSM, HFC, HLR, HLS, HSS, IAB, IDS, IGP, IKE, IPS, ISM, ISO, ISP, ITU, IXC, IXP, KDC, LAN, LCP, LEC, LEO, LER, LLD, LSR, LTE, MAN, MEO, MFJ, MGW, MIC, MME, MPD, MSC, MSS, MTU, NAP, NAT, NAV, NCP, NFC, NIC, NID, NRZ, ONF, OSI, PAR, PCF, PCM, PCS, PGP, PHP, PIM, PKI, PON, POP, PPP, PSK, RAS, RCP, RED, RIP, RMT, RNC, RPC, RPR, RTO, RTP, SCO, SDH, SDN, SIP, SLA, SNR, SPE, SSL, TCG, TCM, TCP, TDM, TLS, TPM, UDP, URL, USB, UTP, UWB, VLR, VPN, W3C, WAF, WAN, WDM, WEP, WFQ and WPA. But don't worry. Each will appear in **boldface type** and be carefully defined before it is used. As a fun test, see how many you can identify *before* reading the book, write the number in the margin, then try again *after* reading the book.

Instructors' Resource Materials

The following protected instructors' resource materials are available on the publisher's Web site at *www.pearsonglobaleditions.com*. For a user-name and password, please contact your local Pearson representative.

- Solutions manual
- PowerPoint lecture slides

Students' Resource Materials

Resources for students are available through the open-access Companion Web site link on *www.pearsonglobaleditions.com*, including

- Figures, tables, and programs from the book
- Steganography demo
- Protocol simulators

Acknowledgements

Many people helped us during the course of the sixth edition. We would especially like to thank Phyllis Davis (St. Louis Community College), Farah Kandah (University of Tennessee, Chattanooga), Jason Livingood (Comcast), Louise Moser (University of California, Santa Barbara), Jennifer Rexford (Princeton), Paul Schmitt (Princeton), Doug Sicker (CMU), Wenye Wang (North Carolina State University), and Greg White (Cable Labs).

Some of Prof. Tanenbaum's students have given valuable feedback on the manuscript, including: Ece Doganer, Yael Goede, Bruno Hoevelaken, Elena Ibi, Oskar Klonowski, Johanna Sänger, Theresa Schantz, Karlis Svilans, Mascha van der Marel, Anthony Wilkes, for providing ideas and feedback.

Jesse Donkervliet (Vrije Universiteit) thought of many new end-of-chapter exercises to challenge the reader.

Paul Nagin (Chimborazo Publishing, Inc.) produced the Power Point slides for instructors.

Our editor at Pearson, Tracy Johnson, was her usual helpful self in many ways large and small. Without her advice, drive, and persistence, this edition might never have happened. Thank you Tracy. We really appreciate your help.

Finally, we come to the most important people. Suzanne has been through this 23 times now and still has endless patience and love. Barbara and Marvin now know the difference between good textbooks and bad ones and are always an inspiration to produce good ones. Daniel and Matilde are wonderful additions to our family. Aron, Nathan, Olivia, and Mirte probably aren't going to read this edition, but they inspire me and make me hopeful about the future (AST). Marshini, Mila, and Kira: My favorite network is the one we have built together. Thank you for your support and love (NF). Katrin and Lucy provided endless support and always managed to keep a smile on my face. Thank you (DJW).

ANDREW S. TANENBAUM

NICK FEAMSTER

DAVID J. WETHERALL

1

INTRODUCTION

Each of the past three centuries was dominated by a single new technology. The 18th century was the era of the great mechanical systems accompanying the Industrial Revolution. The 19th century was the age of the steam engine. During the 20th century, the key technology was information gathering, processing, and distribution. Among other developments, we saw the deployment of worldwide telephone networks, the invention of radio and television, the birth and unprecedented growth of the computer industry, the launching of communication satellites, and, of course, the Internet. Who knows what miracles the 21st century will bring?

As a result of this rapid technological progress, these areas are rapidly converging in the 21st century, and the differences between collecting, transporting, storing, and processing information are quickly disappearing. Organizations with hundreds of offices spread over a wide geographical area routinely expect to be able to examine the current status of even their most remote outpost at the push of a button. As our ability to gather, process, and distribute information grows, the demand for more sophisticated information processing grows even faster.

1.1 USES OF COMPUTER NETWORKS

Although the computing industry is still young compared to other technical industries such as automobiles and air transportation, computers have made spectacular progress in a short time. During the first two decades of their existence,

computer systems were highly centralized, usually within a single room. Often, this room had glass windows, through which visitors could gawk at the great electronic wonder inside. A medium-sized company or university might have had one or two computers, while large institutions had at most a few dozen. The idea that within fifty years vastly more powerful computers smaller than postage stamps would be mass produced by the billions was science fiction.

The convergence of computers and communications has had a profound influence on the organization of computer systems. The once-dominant concept of the ''computer center'' as a room with a single large computer to which users bring their work for processing is now obsolete (although data centers holding hundreds of thousands of Internet servers are common). The old model of a single computer serving all of the organization's computational needs has been replaced by one in which a large number of separate but interconnected computers do the job. These systems are called **computer networks**. The design and organization of these networks are the subjects of this book.

Throughout the book, we will use the term ''computer network'' to mean a collection of interconnected, autonomous computing devices. Two computers are said to be interconnected if they can exchange information. Interconnection can take place over a variety of transmission media including copper wire, fiber optic cable, and radio waves (e.g., microwave, infrared, communication satellites). Networks come in many sizes, shapes, and forms, as we will explore throughout the book. They are usually connected to make larger networks, with the **Internet** being the most well-known example of a network of networks.

1.1.1 Access to Information

Access to information comes in many forms. A common method of accessing information via the Internet is using a Web browser, which allows a user to retrieve information from various Web sites, including increasingly popular social media sites. Mobile applications on smartphones now also allow users to access remote information. Topics include the arts, business, cooking, government, health, history, hobbies, recreation, science, sports, travel, and many others. Fun comes in too many ways to mention, plus some ways that are better left unmentioned.

News organizations have largely migrated online, with some even ceasing print operations entirely. Access to information, including the news, is increasingly personalizable. Some online publications even allow you to tell them that you are interested in corrupt politicians, big fires, scandals involving celebrities, and epidemics, but no football, thank you. This trend certainly threatens the employment of 12-year-old paperboys, but online distribution has allowed the distribution of news to reach far larger and broader audiences.

Increasingly, news is also being curated by social media platforms, where users can post and share news content from a variety of sources, and where the news that any given user sees is prioritized and personalized based on both explicit user

preferences and complex machine learning algorithms that predict user preferences based on the user's history. Online publishing and content curation on social media platforms supports a funding model that depends largely on highly targeted behavioral advertising, which necessarily implies gathering data about the behavior of individual users. This information has sometimes been misused.

Online digital libraries and retail sites now host digital versions of content ranging from academic journals to books. Many professional organizations, such as the ACM (*www.acm.org*) and the IEEE Computer Society (*www.computer.org*), already have all their journals and conference proceedings online. Electronic book readers and online libraries may someday make printed books obsolete. Skeptics should take note of the effect the printing press had on the medieval illuminated manuscript.

Much information on the Internet is accessed using a client-server model, where a client explicitly requests information from a server that hosts that information, as illustrated in Fig. 1-1.

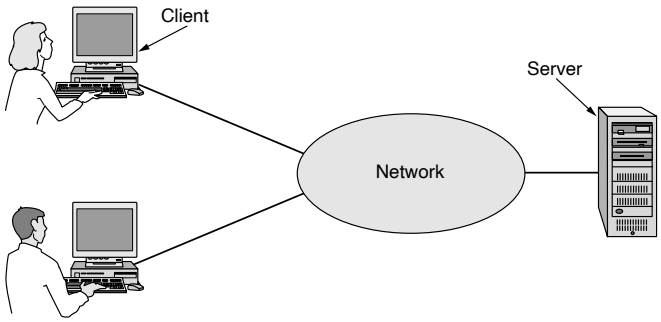

Figure 1-1. A network with two clients and one server.

The **client-server model** is widely used and forms the basis of much network usage. The most popular realization is that of a **Web application**, where a server generates Web pages based on its database in response to client requests that may update the database. The client-server model is applicable not only when the client and server are both in the same building (and belong to the same company), but also when they are far apart. For example, when a person at home accesses a page on the World Wide Web, the same model is employed, with the remote Web server being the server and the user's personal computer being the client. Under most conditions, one server can handle a large number (hundreds or thousands) of clients simultaneously.

If we look at the client-server model, to a first approximation we see that two processes (running programs) are involved, one on the client machine and one on the server machine. Communication takes the form of the client process sending a

message over the network to the server process. The client process then waits for a reply message. When the server process gets the request, it performs the requested work or looks up the requested data and sends back a reply. These messages are shown in Fig. 1-2.

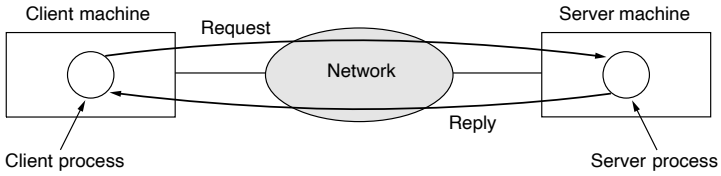

Figure 1-2. The client-server model involves requests and replies.

Another popular model for accessing information is **peer-to-peer** communication (Parameswaran et al., 2001). In this form, individuals who form a loose group can communicate with others in the group, as shown in Fig. 1-3. Every person can, in principle, communicate with one or more other people; there is no fixed division into clients and servers.

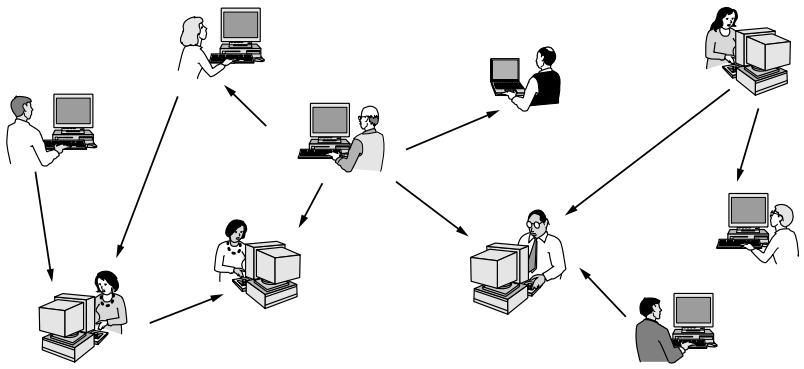

Figure 1-3. In a peer-to-peer system, there are no fixed clients and servers.

Many peer-to-peer systems, such as BitTorrent (Cohen, 2003), do not have a central database of content. Instead, each user maintains a local database of content, as well as a list of other members of the system. A new user can then go to any existing member to see what he has and get the names of other members to inspect for more content and more names. This lookup process can be repeated indefinitely to build up a large local database of what is out there. It is an activity that would get tedious for people, but computers excel at it.

Peer-to-peer communication is often used to share music and videos. It really hit the big time around 2000 with a music sharing service called Napster, which was shut down after a monumental copyright infringement case (Lam and Tan, 2001; and Macedonia, 2000). Legal applications for peer-to-peer communication now exist. These include fans sharing public domain music, families sharing photos and movies, and users downloading public software packages. In fact, one of the most popular Internet applications of all, email, is (conceptually) peer-to-peer. This form of communication is likely to grow considerably in the future.

1.1.2 Person-to-Person Communication

Person-to-person communication is the 21st century's answer to the 19th century's telephone. Email is already used on a daily basis by millions of people all over the world and its use is growing rapidly. It already routinely contains audio and video as well as text and pictures. Smell may take a while.

Many Internet users now rely on some form of **instant messaging** to communicate with other people on the Internet. This facility, derived from the UNIX *talk* program in use since around 1970, allows two people to type messages at each other in real time. There are also multi-person messaging services too, such as the **Twitter** service, which lets people send short messages (possibly including video) called ''tweets'' to their circle of friends or other followers or the whole world.

The Internet can be used by applications to carry audio (e.g., Internet radio stations, streaming music services) and video (e.g., Netflix, YouTube). Besides being an inexpensive way to communicate with your distant friends, these applications can provide rich experiences such as distance learning, meaning attending 8 A.M. classes without the inconvenience of having to get out of bed first. In the long run, the use of networks to enhance human-to-human communication may prove more important than any of the others. It may become hugely important to people who are geographically challenged, giving them the same access to services as people living in the middle of a big city.

Between person-to-person communications and accessing information are **social network** applications. In these applications, the flow of information is driven by the relationships that people declare between each other. One of the most popular social networking sites is **Facebook**. It lets people create and update their personal profiles and shares the updates with other people who they have declared to be their friends. Other social networking applications can make introductions via friends of friends, send news messages to friends, such as Twitter above, and much more.

Even more loosely, groups of people can work together to create content. A **wiki**, for example, is a collaborative Web site that the members of a community edit. The most famous wiki is the **Wikipedia**, an encyclopedia anyone can read or edit, but there are thousands of other wikis.

1.1.3 Electronic Commerce

Online shopping is already popular; users can browse the online catalogs of thousands of companies and have products shipped right to their doorsteps. After the customer buys a product electronically but cannot figure out how to use it, online technical support may be consulted.

Another area in which e-commerce is widely used is access to financial institutions. Many people already pay their bills, manage their bank accounts, and even handle their investments electronically. Financial technology or ''fintech'' applications allow users to conduct a wide variety of financial transactions online, including transferring money between bank accounts, or even between friends.

Online auctions of second-hand goods have become a massive industry. Unlike traditional e-commerce, which follows the client-server model, online auctions are peer-to-peer in the sense that consumers can act as both buyers and sellers, although there is a central server that holds the database of products for sale.

Some of these forms of e-commerce have acquired cute little tags based on the fact that ''to'' and ''2'' are pronounced the same. The most popular ones are listed in Fig. 1-4.

Tag	Full name	Example
B2C	Business-to-consumer	Ordering books online
B2B	Business-to-business	Car manufacturer ordering tires from a supplier
G2C	Government-to-consumer	Government distributing tax forms electronically
C2C	Consumer-to-consumer	Auctioning second-hand products online
P2P	Peer-to-peer	Music or file sharing; Skype

Figure 1-4. Some forms of e-commerce.

1.1.4 Entertainment

Our fourth category is entertainment. This has made huge strides in the home in recent years, with the distribution of music, radio and television programs, and movies over the Internet beginning to rival that of traditional mechanisms. Users can find, buy, and download MP3 songs and high-definition movies and add them to their personal collection. TV shows now reach many homes via **IPTV** (**IP Television**) systems that are based on IP technology instead of cable TV or radio transmissions. Media streaming applications let users tune to Internet radio stations or watch recent episodes of their favorite TV shows or movies. Naturally, all of this content can be moved around your house between different devices, displays, and speakers, usually via a wireless network.

Soon, it may be possible to search for any movie or television program ever made, in any country, and have it be displayed on your screen instantly. New films