

ethernet networks

Fourth Edition

- ♦ **Design**
- ♦ **Implementation**
- ♦ **Operation**
- ♦ **Management**

GILBERT HELD

4-Degree Consulting, Macon, Georgia, USA



JOHN WILEY & SONS, LTD

ethernet
networks

Fourth Edition

Books by Gilbert Held, published by Wiley

Quality of Service in a Cisco® Networking Environment
0 470 84425 6 (April 2002)

Bulletproofing TCP/IP-Based Windows NT/2000 Networks
0 471 49507 7 (April 2001)

Understanding Data Communications: From Fundamentals to Networking,
Third Edition
0 471 62745 3 (October 2000)

*High Speed Digital Transmission Networking: Covering T/E-Carrier
Multiplexing, SONET and SDH, Second Edition*
0 471 98358 6 (April 1999)

*Data Communications Networking Devices: Operation, Utilization and LAN
and WAN Internetworking, Fourth Edition*
0 471 97515 X (November 1998)

*Dictionary of Communications Technology: Terms, Definitions and
Abbreviations, Third Edition*
0 471 97517 6 (May 1998)

Internetworking LANs and WANs: Concepts, Techniques and Methods,
Second Edition
0 471 97514 1 (May 1998)

LAN Management with SNMP and RMON
0 471 14736 2 (September 1996)

ethernet networks

Fourth Edition

- ◆ Design
- ◆ Implementation
- ◆ Operation
- ◆ Management

GILBERT HELD

4-Degree Consulting, Macon, Georgia, USA



JOHN WILEY & SONS, LTD

Copyright © 2003

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England

Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk

Visit our Home Page on www.wileyeurope.com or www.wiley.com

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, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770571.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 33 Park Road, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark,
Singapore 129809

John Wiley & Sons Canada Ltd, 22 Worcester Road, Etobicoke, Ontario, Canada M9W 1L1

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN 0-470-84476-0

Typeset in 10.5/13pt Melior by Laserwords Private Limited, Chennai, India
Printed and bound in Great Britain by Biddles Ltd, Guildford and King's Lynn
This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

For the past decade I have been most fortunate in being able to teach graduate courses that were truly enjoyable to teach. In doing so I have been able to tailor my presentation of technical information covering LAN performance and other data communications topics, providing a two-way learning facility and enhancing my presentation skills. Thus, I would be remiss if I did not thank the students at Georgia College and State University as well as Dr Harvey Glover for providing me with the opportunity to teach. In doing so I would like to dedicate this book to those who pursue higher education.

contents

Preface xv

Acknowledgments xix

Chapter 1 Introduction to Networking Concepts 1

- 1.1 WIDE AREA NETWORKS 2
 - COMPUTER-COMMUNICATIONS EVOLUTION 2
 - REMOTE BATCH TRANSMISSION 3
 - IBM 3270 INFORMATION DISPLAY SYSTEM 4
 - NETWORK CONSTRUCTION 5
 - NETWORK CHARACTERISTICS 8
- 1.2 LOCAL AREA NETWORKS 8
 - COMPARISON TO WANs 9
 - TECHNOLOGICAL CHARACTERISTICS 14
 - TRANSMISSION MEDIUM 22
 - ACCESS METHOD 29
- 1.3 WHY ETHERNET 34

Chapter 2 Networking Standards 37

- 2.1 STANDARDS ORGANIZATIONS 37
 - NATIONAL STANDARDS ORGANIZATIONS 38
 - INTERNATIONAL STANDARDS ORGANIZATIONS 39
- 2.2 THE ISO REFERENCE MODEL 41
 - LAYERED ARCHITECTURE 41
 - OSI LAYERS 42
 - DATA FLOW 46

2.3	IEEE 802 STANDARDS	48
	802 COMMITTEES	48
	DATA LINK SUBDIVISION	51
2.4	INTERNET STANDARDS	55
	RFC EVOLUTION	56
	TYPES AND SUBMISSION	56
	OBTAINING RFCs	57
2.5	CABLING STANDARDS	58
	EIA/TIA-568	58
	UTP CATEGORIES	59
	CABLE SPECIFICATIONS	60
	OTHER METRICS	61
	CAT 5E AND CAT 6	63

Chapter 3 Ethernet Networks 65

3.1	ETHERNET	65
	EVOLUTION	66
	NETWORK COMPONENTS	67
	THE 5-4-3 RULE	73
3.2	IEEE 802.3 NETWORKS	74
	NETWORK NAMES	74
	10BASE-5	75
	10BASE-2	79
	10BROAD-36	87
	1BASE-5	89
	10BASE-T	91
3.3	USE OF FIBER-OPTIC TECHNOLOGY	100
	FOIRL	101
	OPTICAL TRANSCEIVER	101
	FIBER HUBS	101
	FIBER ADAPTER	102
	WIRE AND FIBER DISTANCE LIMITS	102
3.4	HIGH-SPEED ETHERNET	108
	ISOCRONOUS ETHERNET	108
	FAST ETHERNET	111
	100VG-ANYLAN	133

- 3.5 GIGABIT ETHERNET 138**
 - COMPONENTS 138**
 - MEDIA SUPPORT 141**
- 3.6 10 GIGABIT ETHERNET 149**
 - RATIONALE 149**
 - ARCHITECTURE 150**
 - OPERATING RATES 153**

Chapter 4 Frame Operations 155

- 4.1 FRAME COMPOSITION 155**
 - PREAMBLE FIELD 157**
 - START-OF-FRAME DELIMITER FIELD 157**
 - DESTINATION ADDRESS FIELD 157**
 - SOURCE ADDRESS FIELD 160**
 - TYPE FIELD 165**
 - LENGTH FIELD 166**
 - DATA FIELD 168**
 - FRAME CHECK SEQUENCE FIELD 168**
 - INTERFRAME GAP 169**
- 4.2 MEDIA ACCESS CONTROL 169**
 - TRANSMIT MEDIA ACCESS MANAGEMENT 172**
 - SERVICE PRIMITIVES 175**
 - PRIMITIVE OPERATIONS 176**
 - HALF- VERSUS FULL-DUPLEX OPERATION 176**
- 4.3 LOGICAL LINK CONTROL 177**
 - TYPES AND CLASSES OF SERVICE 179**
 - SERVICE PRIMITIVES 181**
- 4.4 OTHER ETHERNET FRAME TYPES 181**
 - ETHERNET-802.3 181**
 - ETHERNET-SNAP 182**
 - IEEE 802.1Q FRAME 183**
 - FRAME DETERMINATION 184**
- 4.5 FAST ETHERNET 185**
 - START-OF-STREAM DELIMITER 186**
 - END-OF-STREAM DELIMITER 186**

- 4.6 GIGABIT ETHERNET 186
 - CARRIER EXTENSION 187
 - FRAME BURSTING 189
- 4.7 10 GIGABIT ETHERNET 190

Chapter 5 Networking Hardware and Software 191

- 5.1 WIRED NETWORK HARDWARE COMPONENTS 192
 - REPEATERS 192
 - BRIDGES 195
 - ROUTERS 205
 - BROUTERS 210
 - GATEWAY 213
 - FILE SERVERS 214
 - WIRE HUBS 218
 - INTELLIGENT HUBS 219
 - SWITCHING HUBS 219
- 5.2 WIRELESS NETWORK HARDWARE COMPONENTS 221
 - NETWORK TOPOLOGIES 221
 - ACCESS POINT 222
 - WIRELESS ROUTER 222
 - WIRELESS BRIDGE 223
- 5.3 NETWORKING SOFTWARE 224
 - DOS 224
 - NETWORK SOFTWARE COMPONENTS 225
 - NETWORK OPERATING SYSTEMS 227
 - APPLICATION SOFTWARE 243
- 5.4 THE TCP/IP PROTOCOL SUITE 244
 - OVERVIEW 244
 - PROTOCOL DEVELOPMENT 244
 - THE TCP/IP STRUCTURE 245
 - DATAGRAMS VERSUS VIRTUAL CIRCUITS 247
 - ICMP 249
 - ARP 252
 - TCP 254
 - UDP 259
 - IP 260

DOMAIN NAME SERVICE	269
NAME SERVER	272
TCP/IP CONFIGURATION	272
OPERATING MULTIPLE STACKS	275

Chapter 6 Bridging and Switching Methods and Performance Issues 279

6.1 BRIDGING METHODS	279
ADDRESS ISSUES	280
TRANSPARENT BRIDGING	280
SPANNING TREE PROTOCOL	283
PROTOCOL DEPENDENCY	291
SOURCE ROUTING	292
SOURCE ROUTING TRANSPARENT BRIDGES	297
6.2 BRIDGE NETWORK UTILIZATION	299
SERIAL AND SEQUENTIAL BRIDGING	300
PARALLEL BRIDGING	301
STAR BRIDGING	302
BACKBONE BRIDGING	302
6.3 BRIDGE PERFORMANCE ISSUES	303
TRAFFIC FLOW	303
NETWORK TYPES	304
TYPE OF BRIDGE	304
ESTIMATING NETWORK TRAFFIC	304
PREDICTING THROUGHPUT	311
6.4 LAN SWITCHES	312
RATIONALE	313
BOTTLENECKS	314
CONGESTION-AVOIDANCE OPTIONS	314
LAN SWITCH OPERATIONS	318
6.5 SWITCH BASIC ARCHITECTURE	332
COMPONENTS	332
SWITCH FEATURES	335
SWITCHED-BASED VIRTUAL LANs	348
SWITCH USAGE	360
LAYER 3 AND LAYER 4 SWITCHING	364

Chapter 7 Routers 365

- 7.1 ROUTER OPERATION 365**
 - IP SUPPORT OVERVIEW 365**
 - BASIC OPERATION AND USE OF ROUTING TABLES 368**
 - NETWORKING CAPABILITY 370**
- 7.2 COMMUNICATION, TRANSPORT, AND ROUTING PROTOCOLS 371**
 - COMMUNICATION PROTOCOL 371**
 - ROUTING PROTOCOL 372**
 - HANDLING NONROUTABLE PROTOCOLS 372**
 - TRANSPORT PROTOCOL 373**
- 7.3 ROUTER CLASSIFICATIONS 374**
 - PROTOCOL-DEPENDENT ROUTERS 374**
 - PROTOCOL-INDEPENDENT ROUTERS 377**
- 7.4 ROUTING PROTOCOLS 381**
 - TYPES OF ROUTING PROTOCOLS 381**
 - INTERIOR DOMAIN ROUTING PROTOCOLS 381**
 - EXTERIOR DOMAIN ROUTING PROTOCOLS 382**
 - TYPES OF INTERIOR DOMAIN ROUTING PROTOCOLS 384**
 - ROUTING INFORMATION PROTOCOL 386**
 - CONFIGURATION EXAMPLE 389**
 - ROUTING TABLE MAINTENANCE PROTOCOL 392**
 - INTERIOR GATEWAY ROUTING PROTOCOL 393**
 - LINK STATE PROTOCOLS 394**
- 7.5 FILTERING 399**
 - FILTERING EXPRESSIONS 401**
 - FILTERING EXAMPLES 401**
 - ROUTER ACCESS LISTS 402**
- 7.6 PERFORMANCE CONSIDERATIONS 404**

Chapter 8 Wireless Ethernet 407

- 8.1 OVERVIEW 407**
 - NETWORK TOPOLOGY 409**
 - ROAMING 411**
 - PHYSICAL LAYER OPERATIONS 412**
 - HIGH-SPEED WIRELESS LANs 415**
 - ACCESS METHOD 418**

- 8.2 FRAME FORMATS 421**
 - DATA FRAME 421**
 - CONTROL FIELD 422**
 - CONTROL FRAMES 428**
 - MANAGEMENT FRAMES 430**
 - PHYSICAL PROTOCOL DATA UNITS 432**
- 8.3 DEPLOYMENT 434**
 - WIRELESS PC NETWORK ADAPTER CARDS 434**
 - ACCESS POINT 435**
 - COMBINED ROUTER/ACCESS POINT 437**
 - WIRELESS BRIDGE 439**
 - ROUTER/ACCESS POINT CONFIGURATION 439**
 - CLIENT CONFIGURATION 441**

Chapter 9 Security 447

- 9.1 THE SECURITY ROLE OF THE ROUTER 447**
 - ACCESS CONTROL 448**
 - ACCESS LISTS 457**
 - STANDARD IP ACCESS LISTS 459**
 - EXTENDED IP ACCESS LISTS 462**
 - ANTI-SPOOFING STATEMENTS 471**
 - NAMED ACCESS LISTS 472**
 - DYNAMIC ACCESS LISTS 475**
 - REFLEXIVE ACCESS LISTS 478**
 - TIME-BASED ACCESS LISTS 482**
 - CONTEXT BASED ACCESS CONTROL 483**
- 9.2 THE ROLE OF THE FIREWALL 494**
 - ACCESS-LIST LIMITATIONS 494**
 - PROXY SERVICES 496**
 - FIREWALL LOCATION 498**
 - THE TECHNOLOGIC INTERCEPTOR 504**
 - CHECKPOINT FIREWALL-1 510**
- 9.3 THE ROLE OF THE VIRUS SCANNER AND ENCRYPTION 516**
 - VIRUS OVERVIEW 516**
 - TYPES OF VIRUSES 517**
 - INFECTION PREVENTION 519**

DESKTOP SCANNING	520
EMAIL SCANNING	525
RECOGNIZING INFECTION SYMPTOMS	528

Chapter 10 Managing the Network 531

10.1	SNMP	531
	BASIC COMPONENTS	532
	OPERATION	534
10.2	REMOTE MONITORING	535
	OPERATION	535
	THE RMON MIB	536
	MANAGING REMOTE NETWORKS	539
10.3	OTHER NETWORK MANAGEMENT FUNCTIONS	541
	CONFIGURATION MANAGEMENT	542
	PERFORMANCE MANAGEMENT	543
	FAULT MANAGEMENT	543
	ACCOUNTING MANAGEMENT	544
	SECURITY MANAGEMENT	544
10.4	REPRESENTATIVE NETWORK MANAGEMENT PROGRAMS	544
	TRITICOM ETHERVISION	545
	CINCO NETWORK'S WEBXRAY	554
	WILDPACKETS ETHERPEEK	559

Chapter 11 The Future of Ethernet 567

11.1	ETHERNET TRENDS	567
	NETWORK ADAPTER CARD COST	567
	FUTURE PRICE DIRECTION	568
11.2	NETWORK PERFORMANCE CONSIDERATIONS	571
	SUPPLEMENTING AN EXISTING NETWORK	571
	SUMMARY	579

Index 581

preface

In a prior edition of this book the preface commenced with the paraphrase of an old adage in an era of evolving local area networking technology: Ethernet is dead — long live Ethernet!

Although advances in communications technology continue to occur at a rapid pace, that paraphrase continues to be valid. Within the past decade, the bandwidth of 10 Mbps Ethernet was advanced by a factor of one thousand with the introduction of a series of enhancements to the original Ethernet specification. First, Fast Ethernet resulted in the bandwidth of Ethernet increasing by a factor of 10 to 100 Mbps. The introduction of Gigabit Ethernet resulted in another order of magnitude increase in bandwidth to 1 Gbps. Although many persons felt that a transmission capacity of 1 Gbps would be more than sufficient for the foreseeable future, another adage states that many applications will grow to use all available bandwidth. While most organizations may be hard pressed to use 1 Gbps of bandwidth, other organizations, including Internet Service Providers and corporations and universities with large backbone LANs, were able to literally fill the 1 Gbps pipe, resulting in the development of 10 Gbps Ethernet. Thus, over the past decade Ethernet's 10 Mbps operation has increased by a factor of 1000 to 10 Gbps.

This new edition provides a significant amount of additional material to most of the chapters of this book's previous edition. New information added includes coverage of the transmission of Gigabit over copper conductors, the evolution of cabling standards that facilitate support of higher Ethernet operating rates, and the manner by which LAN switches operate on Ethernet frames transporting information at higher layers in the Open System Interconnection Reference Model.

Recognizing the importance of networking without wires, a new chapter is focused upon wireless Ethernet. This chapter describes and discusses the series of IEEE 802.11 standards and provides practical information concerning the setup and operation of a wireless LAN. Recognizing the importance of security in the modern era of networking resulted in the movement of most security related topics to a new chapter focused on this topic. This chapter considerably expands the prior disparate coverage of security by adding information covering the use of firewalls in both a wired and wireless

environment. In addition, information concerning the use of router access lists is considerably expanded, while new information covering authentication, authorization and accounting has been added to the chapter.

Other topics that have been added or significantly revised in this new edition include the operation of new versions of Windows on Ethernet LANs, the operation and utilization of LAN switches above layer 2 in the ISO Reference Model, new gateway methods you can consider to connect workstation users to mainframes, and the use of both copper and fiber optic to transport high-speed Ethernet. Thus, the scope and depth of material have been significantly revised and updated to continue to provide you with detailed information concerning the design, implementation, operation and management of different types of Ethernet networks.

This book incorporates into one reference source the material you will need to understand how Ethernet networks operate, the constraints and performance issues that affect their design and implementation, and how their growth and use can be managed both locally and as part of an enterprise network. Assuming readers have varied backgrounds in communications terms and technology, the first two chapters were written to provide a common foundation of knowledge. Those chapters cover networking concepts and network standards—two topics on which material in succeeding chapters is based. Succeeding chapters examine Ethernet concepts: frame operations; network construction; the use of bridges, routers, hubs, switches, and gateways; Internet connectivity; network backbone construction; Wireless Ethernet; Security; and the management of Ethernet networks.

In writing this book, my primary goal was to incorporate practical information you can readily use in designing, operating, implementing, and managing an Ethernet network. Although Ethernet had its origins in the 1970s and can be considered a relatively “old” technology, in reality, the technology is anything but old. Only a few years ago, the standardization of what is now known as 10BASE-T (a twisted-wire version of Ethernet) resulted in a considerable expansion in the use of this type of local area network. By 1994 the use of intelligent switches greatly enhanced the operational capability of 10BASE-T networks, providing multiple simultaneous 10 Mbps connectivity. During 1995 high-speed Ethernet technology in the form of Fast Ethernet products provided users with the ability to upgrade their Ethernet networks to satisfy emerging multimedia requirements. Within a few years industry realized that emerging applications, as well as the growth in the use of the Internet, required higher-speed backbone LANs as a mechanism to support Internet access and similar high-speed networking requirements. This realization resulted in the deployment of Gigabit Ethernet hubs and switches during

1997, which was quickly followed by 10 Gbps operations a few years later. Thus, Ethernet technology can be expected to continue to evolve to satisfy the communications requirements of business, government, and academia.

For over 30 years I have worked as a network manager responsible for the design, operation, and management of an enterprise network in which local area networks throughout the United States are interconnected through the use of different wide area network transmission facilities. This challenging position has provided me with the opportunity to obtain practical experience in designing, operating, and interconnecting Ethernet networks to Token-Ring, SNA, the Internet, and other types of networks—experience which I have attempted to share with you. This book will help you consider the practicality of different types of routing protocols, LAN switches, and gateway methods. These and other network design issues are crucial to the efficient and effective expansion of a local Ethernet so that users on that network can access resources on other networks.

As a professional author, I very much value readers' comments. Those comments provide me with feedback necessary to revise future editions so that they better reflect the information requirements of readers. I look forward to receiving your comments, as well as suggestions for information you would like to see in a future edition of this book. You can write to me directly or through my publisher, whose address you can find on page 4 of this book or communicate with me directly via electronic mail at gil_held@yahoo.com.

Gilbert Held
Macon, GA

acknowledgments

This book would not have been possible without the work of two people whose foresight and pioneering efforts were instrumental in the development of the technology upon which Ethernet is based.

One of the key concepts behind Ethernet—that of allocating the use of a shared channel—can be traced to the pioneering efforts of Dr Norman Abramson and his colleagues at the University of Hawaii during the early 1970s. The actual development of Ethernet is due to the foresight of Dr Robert Metcalfe. Working at the Xerox Palo Alto Research Center in Palo Alto, California, Dr Metcalfe headed a development team that interconnected over 100 computers on a 1-km cable using a carrier sense multiple access collision detection (CSMA/CD) protocol. In addition to pioneering the technical development of Ethernet, Dr Metcalfe coined its name, after the luminiferous ether through which electromagnetic radiation was once thought to propagate. I would be remiss if I did not thank Dr Abramson, Dr Metcalfe, and their colleagues for their visionary efforts in developing the technology through which hundreds of millions of people now communicate.

Writing and producing a book about technology requires not only the technology itself, but also the efforts of many individuals. First and foremost, I would like to thank my family for their understanding for the nights and weekends I disappeared to write this book. Once again, I am indebted to Mrs Linda Hayes and Mrs Susan Corbitt for taking my notes and drawings and converting them into a manuscript. Last, but not least, I would like to thank Birgit Gruber and Ann-Marie Halligan as well as the production staff at John Wiley & Sons for backing the new edition of this book, as well as in facilitating the conversion of my manuscript into the book you are reading.

Introduction to Networking Concepts

One of the most logical assumptions an author can make is that readers will have diverse backgrounds of knowledge and experience. Making this book as valuable as possible to persons with different backgrounds requires an introductory chapter that covers basic networking concepts. Unfortunately, basic concepts for one person may not be the same as basic concepts for another person, which presents an interesting challenge for an author.

To meet this challenge, this book takes two courses of action. First, it assumes that some readers will have limited knowledge about the different types of communications systems available for transporting information, the relationship between wide area networks (WANs) and local area networks (LANs), and the relationships among different types of local area networks. Thus, this introductory chapter was written to provide those readers with a basic level of knowledge concerning these important topics. Secondly, readers who are already familiar with these basic concepts may wish to consult individual chapters separately, rather than reading through the entire book. To satisfy those readers, each chapter was written to be as independent as possible from preceding and succeeding chapters. Thus, readers who are familiar with wide and local area networking concepts, as well as the technical characteristics of LANs, may elect to skim or bypass this chapter. For other readers, information contained in this chapter will provide a level of knowledge that will make succeeding chapters more meaningful.

In this introductory chapter, we will first focus our attention on the key concepts behind the construction of wide area networks and local area networks. In doing so, we will examine each type of network to obtain an understanding of its primary design goal. Next, we will compare and contrast their operations and utilizations to obtain an appreciation for the rationale behind the use of different types of local area networks.

Although this book is about Ethernet networks, there are other types of local area networks that provide a viable data transportation highway for millions of users. By reviewing the technological characteristics of different types of LANs, we will obtain an appreciation for the governing characteristics behind the use of different local area networks. In addition, because many local area networks are connected to other LANs and WANs, we will conclude this chapter by focusing on the technological characteristics of local area networks. This will form a foundation for discussing a variety of Ethernet networking issues in succeeding chapters of this book.

1.1 Wide Area Networks

The evolution of wide area networks can be considered to have originated in the mid- to late 1950s, commensurate with the development of the first generation of computers. Based on the use of vacuum tube technology, the first generation of computers were large, power-hungry devices whose placement resulted in a focal point for data processing and the coinage of the term *data center*.

Computer-Communications Evolution

Originally, access to the computational capability of first-generation computers was through the use of punched cards. After an employee of the organization used a keypunch to create a deck of cards, that card deck was submitted to a window in the data center, typically labeled input/output (I/O) control. An employee behind the window would accept the card deck and complete a form that contained instructions for running the submitted job. The card deck and instructions would then be sent to a person in production control, who would schedule the job and turn it over to operations for execution at a predefined time. Once the job was completed, the card deck and any resulting output would be sent back to I/O control, enabling the job originator to return to the window in the data center to retrieve his or her card deck and the resulting output. With a little bit of luck, programmers might see the results of their efforts on the same day that they submitted their jobs.

Because the computer represented a considerable financial investment for most organizations, it was understandable that these organizations would be receptive to the possibility of extending their computers' accessibility. By the mid-1960s, several computer manufacturers had added remote access capabilities to one or more of their computers.

Remote Batch Transmission

One method of providing remote access was the installation of a batch terminal at a remote location. That terminal was connected via a telephone company–supplied analog leased line and a pair of modems to the computer in the corporate data center.

The first type of batch terminal developed to communicate with a data center computer contained a card reader, a printer, a serial communications adapter, and hard-wired logic in one common housing. The serial communications adapter converted the parallel bits of each internal byte read from the card reader into a serial data stream for transmission. Similarly, the adapter performed a reverse conversion process, converting a sequence of received serial bits into an appropriate number of parallel bits to represent a character internally within the batch terminal. Because the batch terminal was located remotely from the data center, it was often referred to as a *remote batch terminal*, while the process of transmitting data was referred to as *remote batch transmission*. In addition, the use of a remote terminal as a mechanism for grouping card decks of individual jobs, all executed at the remote data center, resulted in the term *remote job entry terminal* being used as a name for this device.

Figure 1.1 illustrates in schematic form the relationships between a batch terminal, transmission line, modems, and the data center computer. Because the transmission line connects a remote batch terminal in one geographic area to a computer located in a different geographic area, Figure 1.1 represents one of the earliest types of wide area data communications networks.

Paralleling the introduction of remote batch terminals was the development of a series of terminal devices, control units, and specialized communications equipment, which resulted in the rapid expansion of interactive computer applications. One of the most prominent collections of products was introduced by the IBM Corporation under the trade name *3270 Information Display System*.




Figure 1.1 Remote batch transmission. The transmission of data from a remote batch terminal represents one of the first examples of wide area data communications networks.

IBM 3270 Information Display System

The IBM 3270 Information Display System was a term originally used to describe a collection of products ranging from interactive terminals that communicate with a computer, referred to as *display stations*, through several types of control units and communications controllers. Later, through the introduction of additional communications products from IBM and numerous third-party vendors and the replacement of previously introduced products, the IBM 3270 Information Display System became more of a networking architecture and strategy rather than a simple collection of products.

First introduced in 1971, the IBM 3270 Information Display System was designed to extend the processing power of the data center computer to remote locations. Because the data center computer typically represented the organization's main computer, the term *mainframe* was coined to refer to a computer with a large processing capability. As the mainframe was primarily designed for data processing, its utilization for supporting communications degraded its performance.

Communications Controller

To offload communications functions from the mainframe, IBM and other computer manufacturers developed hardware to sample communications lines for incoming bits, group bits into bytes, and pass a group of bytes to the mainframe for processing. This hardware also performed a reverse function for data destined from the mainframe to remote devices. When first introduced, such hardware was designed using fixed logic circuitry, and the resulting device was referred to as a *communications controller*. Later, minicomputers were developed to execute communications programs, with the ability to change the functionality of communications support by the modification of software—a considerable enhancement to the capabilities of this series of products. Because both hard-wired communications controllers and programmed minicomputers performing communications offloaded communications processing from the mainframe, the term *front-end processor* evolved to refer to this category of communications equipment. Although most vendors refer to a minicomputer used to offload communications processing from the mainframe as a front-end processor, IBM has retained the term *communications controller*, even though their fixed logic hardware products were replaced over 20 years ago by programmable minicomputers.

Control Units

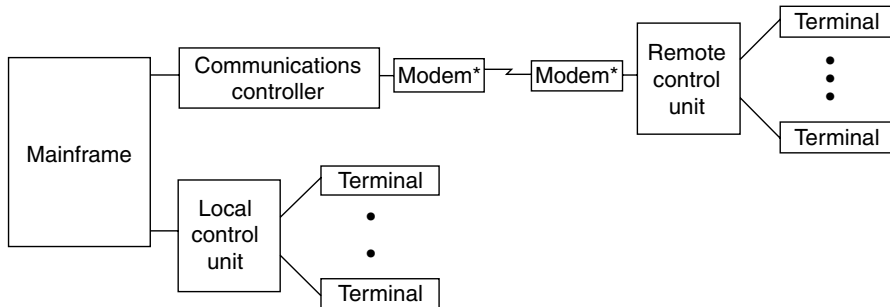
To reduce the number of controller ports required to support terminals, as well as the amount of cabling between controller ports and terminals, IBM developed *poll and select* software to support its 3270 Information Display System. This software enabled the communications controller to transmit messages from one port to one or more terminals in a predefined group of devices. To share the communications controller port, IBM developed a product called a *control unit*, which acts as an interface between the communications controller and a group of terminals.

In general terms, the communications controller transmits a message to the control unit. The control unit examines the terminal address and retransmits the message to the appropriate terminal. Thus, control units are devices that reduce the number of lines required to link display stations to mainframe computers. Both local and remote control units are available; the key differences between them are the method of attachment to the mainframe computer and the use of intermediate devices between the control unit and the mainframe.

Local control units are usually attached to a channel on the mainframe, whereas remote control units are connected to the mainframe's front-end processor, which is also known as a communications controller in the IBM environment. Because a local control unit is within a limited distance of the mainframe, no intermediate communications devices, such as modems or data service units, are required to connect a local control unit to the mainframe. In comparison, a remote control unit can be located in another building or in a different city; it normally requires the utilization of intermediate communications devices, such as a pair of modems or a pair of data service units, for communications to occur between the control unit and the communications controller. The relationship of local and remote control units to display stations, mainframes, and a communications controller is illustrated in Figure 1.2.

Network Construction

To provide batch and interactive access to the corporate mainframe from remote locations, organizations began to build sophisticated networks. At first, communications equipment such as modems and transmission lines was obtainable only from AT&T and other telephone companies. Beginning in 1974 in the United States with the well-known Carterphone decision, competitive non-telephone company sources for the supply of communications equipment became available. The divestiture of AT&T during the 1980s and



*Note: Modems replaced by data service units when a digital transmission facility used.

Figure 1.2 Relationship of 3270 information display products.

the emergence of many local and long-distance communications carriers paved the way for networking personnel to be able to select from among several or even hundreds of vendors for transmission lines and communications equipment.

As organizations began to link additional remote locations to their mainframes, the cost of providing communications began to escalate rapidly. This, in turn, provided the rationale for the development of a series of line-sharing products referred to as *multiplexers* and *concentrators*. Although most organizations operated separate data and voice networks, in the mid-1980s communications carriers began to make available for commercial use high-capacity circuits known as T1 in North America and E1 in Europe. Through the development of T1 and E1 multiplexers, voice, data, and video transmission can share the use of common high-speed circuits. Because the interconnection of corporate offices with communications equipment and facilities normally covers a wide geographical area outside the boundary of one metropolitan area, the resulting network is known as a *wide area network* (WAN).

Figure 1.3 shows an example of a wide area network spanning the continental United States. In this example, regional offices in San Francisco and New York are connected with the corporate headquarters, located in Atlanta, via T1 multiplexers and T1 transmission lines operating at 1.544 Mbps. Assuming that each T1 multiplexer is capable of supporting the direct attachment of a *private branch exchange* (PBX), both voice and data are carried by the T1 circuits between the two regional offices and corporate headquarters. The three T1 circuits can be considered the primary data highway, or *backbone*, of the corporate network.

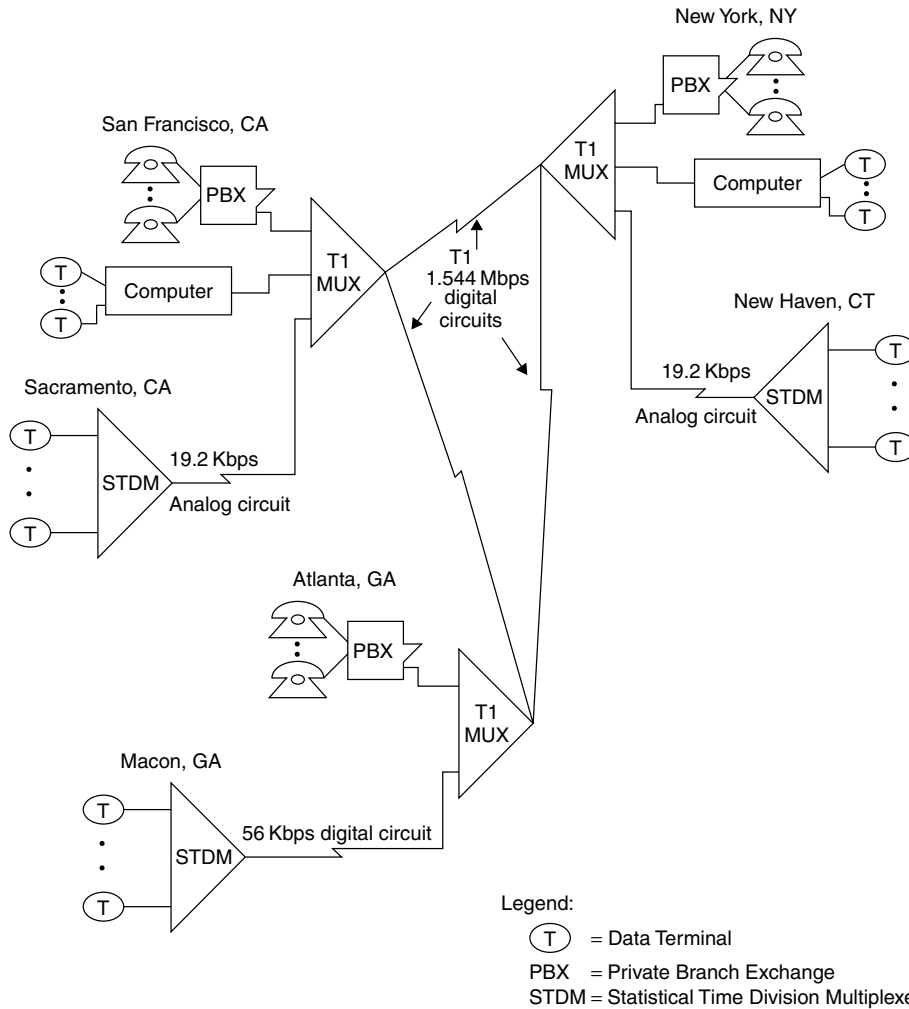


Figure 1.3 Wide area network example. A WAN uses telecommunications lines obtained from one or more communications carriers to connect geographically dispersed locations.

In addition to the three major corporate sites that require the ability to route voice calls and data between locations, let us assume that the corporation also has three smaller area offices located in Sacramento, California; Macon, Georgia; and New Haven, Connecticut. If these locations only require data terminals to access the corporate network for routing to the computers located

in San Francisco and New York, one possible mechanism to provide network support is obtained through the use of tail circuits. These tail circuits could be used to connect a *statistical time division multiplexer (STDM)* in each area office, each serving a group of data terminals to the nearest T1 multiplexer, using either analog or digital circuits. The T1 multiplexer would then be configured to route data terminal traffic over the corporate backbone portion of the network to its destination.

Network Characteristics

There are certain characteristics we can associate with wide area networks. First, the WAN is typically designed to connect two or more geographical areas. This connection is accomplished by the lease of transmission facilities from one or more communications vendors. Secondly, most WAN transmission occurs at or under a data rate of 1.544 Mbps or 2.048 Mbps, which are the operating rates of T1 and E1 transmission facilities.

A third characteristic of WANs concerns the regulation of the transmission facilities used for their construction. Most, if not all, transmission facilities marketed by communications carriers are subject to a degree of regulation at the federal, state, and possibly local government levels. Even though we now live in an era of deregulation, carriers must seek approval for many offerings before making new facilities available for use. In addition, although many of the regulatory controls governing the pricing of services were removed, the communications market is still not a truly free market. Thus, regulatory agencies at the federal, state, and local levels still maintain a degree of control over both the offering and pricing of new services and the pricing of existing services.

1.2 Local Area Networks

The origin of local area networks can be traced, in part, to IBM terminal equipment introduced in 1974. At that time, IBM introduced a series of terminal devices designed for use in transaction-processing applications for banking and retailing. What was unique about those terminals was their method of connection: a common cable that formed a loop provided a communications path within a localized geographical area. Unfortunately, limitations in the data transfer rate, incompatibility between individual IBM loop systems, and other problems precluded the widespread adoption of this method of networking. The economics of media sharing and the ability to provide common access

to a centralized resource were, however, key advantages, and they resulted in IBM and other vendors investigating the use of different techniques to provide a localized communications capability between different devices. In 1977, Datapoint Corporation began selling its Attached Resource Computer Network (ARCNet), considered by most people to be the first commercial local area networking product. Since then, hundreds of companies have developed local area networking products, and the installed base of terminal devices connected to such networks has increased exponentially. They now number in the hundreds of millions.

Comparison to WANs

Local area networks can be distinguished from wide area networks by geographic area of coverage, data transmission and error rates, ownership, government regulation, and data routing—and, in many instances, by the type of information transmitted over the network.

Geographic Area

The name of each network provides a general indication of the scope of the geographic area in which it can support the interconnection of devices. As its name implies, a LAN is a communications network that covers a relatively small local area. This area can range in scope from a department located on a portion of a floor in an office building, to the corporate staff located on several floors in the building, to several buildings on the campus of a university.

Regardless of the LAN's area of coverage, its geographic boundary will be restricted by the physical transmission limitations of the local area network. These limitations include the cable distance between devices connected to the LAN and the total length of the LAN cable. In comparison, a wide area network can provide communications support to an area ranging in size from a town or city to a state, country, or even a good portion of the entire world. Here, the major factor governing transmission is the availability of communications facilities at different geographic areas that can be interconnected to route data from one location to another.

To better grasp the differences between LANs and WANs, today we can view the LAN as being analogous to our local telephone company, while the WAN can be compared with the long-distance communications carrier. However, this may not be true in the future when local telephone companies obtain permission to offer long-distance service and long-distance communications

carriers obtain regulatory approval to offer local telephone service. However, for the present we will presume that telephone support in different cities is provided by the local telephone company in each city. Thus, for calls between cities, the local telephone companies must connect to the long-distance carrier. Similarly, we can have separate LANs in different cities or within different buildings in the same city; however, to interconnect those LANs we would normally require a wide area network.

Until the turn of the millennium these differences between LANs and WANs with the respect to geographic area of coverage were distinct and easy to recognize. However, within the past two years a new role has been proposed for Gigabit Ethernet that could enable this technology to function as a miniature WAN. As we will note later in this book, the ability to transmit Gigabit Ethernet over optical fiber makes it possible to transmit this technology over extended distances. In fact, by the time you read this book Gigabit Ethernet may provide a low-cost alternative to synchronous optical networking (SONET) transmission.

Data Transmission and Error Rates

Two additional areas that differentiate LANs from WANs and explain the physical limitation of the LAN geographic area of coverage are the data transmission rate and error rate for each type of network. Older LANs, such as the original version of Ethernet and Token-Ring, normally operate at a low megabit-per-second rate, typically ranging from 4 Mbps to 16 Mbps. More modern high-speed Ethernet networks, such as Fast Ethernet that operates at 100 Mbps, Gigabit Ethernet and 10 Gigabit Ethernet provide transmission rates of 1 Gbps and 10 Gbps, respectively. In comparison, the communications facilities used to construct a major portion of most WANs provide a data transmission rate at or under the T1 and E1 data rates of 1.544 Mbps and 2.048 Mbps.

Although some readers may have encountered references to WAN transmission rates of 10 and 40 Gbps in various trade literature, the use of optical carriers (OCs) at those data rates is primarily by communications carriers whose transmission facilities are shared by tens of thousands to millions of users. Thus, in considering the data transmission rate with respect to LANs and WANs on a non-communications carrier basis, we can say that LANs provide a transmission capacity up to 10 Gbps while WANs are limited to a low Mbps data rate.

Because LAN cabling is primarily within a building or over a small geographical area, it is relatively safe from natural phenomena, such as thunderstorms and lightning. This safety enables transmission at a relatively high data rate,

resulting in a relatively low error rate. In comparison, because wide area networks are based on the use of communications facilities that are much farther apart and always exposed to the elements, they have a much higher probability of being disturbed by changes in the weather, electronic emissions generated by equipment, or such unforeseen problems as construction workers accidentally causing harm to a communications cable. Because of these factors, the error rate on WANs is considerably higher than the rate experienced on LANs. On most WANs you can expect to experience an error rate between 1 in a million and 1 in 10 million (1×10^6 to 1×10^7) bits. In comparison, the error rate on a typical LAN may exceed that range by one or more orders of magnitude, resulting in an error rate from 1 in 10 million to 1 in 100 million bits.

Ownership

The construction of a wide area network requires the leasing of transmission facilities from one or more communications carriers. Although your organization can elect to purchase or lease communications equipment, the transmission facilities used to connect diverse geographical locations are owned by the communications carrier. In comparison, an organization that installs a local area network normally owns all of the components used to form the network, including the cabling used to form the transmission path between devices.

Regulation

Because wide area networks require transmission facilities that may cross local, state, and national boundaries, they may be subject to a number of governmental regulations at the local, state, and national levels. Most of those regulations govern the services that communications carriers can provide customers and the rates (*tariff*) they can charge for those services. In comparison, regulations affecting local area networks are primarily in the areas of building codes. Such codes regulate the type of wiring that can be installed in a building and whether the wiring must run in a conduit.

Data Routing and Topology

In a local area network data is routed along a path that defines the network. That path is normally a bus, ring, tree, or star structure, and data always flows on that structure. The topology of a wide area network can be much more complex. In fact, many wide area networks resemble a mesh structure,

including equipment to reroute data in the event of communications circuit failure or excessive traffic between two locations. Thus, the data flow on a wide area network can change, while the data flow on a local area network primarily follows a single basic route.

Type of Information Carried

The last major difference between local and wide area networks is the type of information carried by each network. Many wide area networks support the simultaneous transmission of voice, data, and video information. In comparison, most local area networks are currently limited to carrying data. In addition, although all wide area networks can be expanded to transport voice, data, and video, many local area networks are restricted by design to the transportation of data. An exception to the preceding is asynchronous transfer mode (ATM), which can provide both a local and wide area network transmission capability. Asynchronous transfer mode was designed to support voice, data, and video from end-to-end, enabling different types of data to be transported from one LAN to another via an ATM WAN. Unfortunately, the cost of ATM network adapters considerably exceeded the cost of other types of LAN equipment used to include different types of Ethernet adapters. As the base of Ethernet expanded, the cost associated with establishing an Ethernet infrastructure decreased, widening the price gap between ATM and Ethernet, making it difficult for asynchronous transfer mode to establish a viable market for local area networking. Today the vast majority of ATM equipment is used by communications carriers in the wide area network. Table 1.1 summarizes the similarities and differences between local and wide area networks.

Utilization Benefits

In its simplest form, a local area network is a cable that provides an electronic highway for the transportation of information to and from different devices connected to the network. Because a LAN provides the capability to route data between devices connected to a common network within a relatively limited distance, numerous benefits can accrue to users of the network. These can include the ability to share the use of peripheral devices, thus obtaining common access to data files and programs, the ability to communicate with other people on the LAN by electronic mail, and the ability to access the larger processing capability of mainframes through common gateways that link a local area network to larger computer systems. Here, the gateway can be directly cabled to the mainframe if it resides at the same location, or it may be connected remotely via a corporate wide area network.

TABLE 1.1 Comparing LANs and WANs

Characteristic	Local Area Network	Wide Area Network
Geographic area of coverage	Localized to a building, group of buildings, or campus	Can span an area ranging in size from a city to the globe
Data transmission rate	Typically 4 Mbps to 16 Mbps, with high-speed Ethernet and fiber optic-based networks operating at 100 Mbps and 1 and 10 Gbps	Normally operate at or below T1 and E1 transmission rates of 1.544 Mbps and 2.048 Mbps
Error rate	1 in 10^7 to 1 in 10^8	1 in 10^6 to 1 in 10^7
Ownership	Usually with the implementor	Communications carrier retains ownership of line facilities
Data routing	Normally follows fixed route	Switching capability of network allows dynamic alteration of data flow
Topology	Usually limited to bus, ring, tree, or star	Virtually unlimited design capability
Type of information carried	Primarily data	Voice, data, and video commonly integrated

Peripheral sharing allows network users to access color laser printers, CD-ROM jukebox systems, and other devices that may be needed only a small portion of the time a workstation is in operation. Thus, users of a LAN can obtain access to resources that would probably be too expensive to justify for each individual workstation user.

The ability to access data files and programs from multiple workstations can substantially reduce the cost of software. In addition, shared access to database information allows network users to obtain access to updated files on a real-time basis.

One popular type of application program used on LANs enables users to transfer messages electronically. Commonly referred to as *electronic mail* or *e-mail*, this type of application program can be used to supplement and, in many cases, eliminate the need for paper memoranda.

For organizations with a mainframe, a local area network gateway can provide a common method of access. Without the use of a LAN gateway, each

personal computer requiring access to a mainframe would require a separate method of access. This might increase both the complexity and the cost of providing access.

Perhaps the most popular evolving use of LANs is to provide a group of computer users with economical access to the Internet. Instead of having a large number of employees obtain individual modem dial-up or ISDN dial access accounts with an Internet service provider (ISP), it is often more economical to connect an organizational LAN to the Internet via a single connection to an ISP. In addition, the connection to the Internet will usually provide a higher transmission capability than obtainable on an individual user basis. Later in this book we will turn our attention to methods that can be used to connect organizational LANs to the Internet, as well as the use of different products to protect organizational computer facilities from Internet users that have no business accessing those facilities.

Technological Characteristics

Although a local area network is a limited distance transmission system, the variety of options available for constructing such networks is anything but limited. Many of the options available for the construction of local area networks are based on the technological characteristics that govern their operation. These characteristics include different topologies, signaling methods, transmission media, access methods used to transmit data on the network, and the hardware and software required to make the network operate.

Topology

The *topology* of a local area network is the structure or geometric layout of the cable used to connect stations on the network. Unlike conventional data communications networks, which can be configured in a variety of ways with the addition of hardware and software, most local area networks are designed to operate based on the interconnection of stations that follow a specific topology. The most common topologies used in LANs include the loop, bus, ring, star, and tree, as illustrated in Figure 1.4.

Loop As previously mentioned in this chapter, IBM introduced a series of transaction-processing terminals in 1974 that communicated through the use of a common controller on a cable formed into a loop. This type of topology is illustrated at the top of Figure 1.4.




Figure 1.4 Local area network topology. The five most common geometric layouts of LAN cabling form a loop, bus, ring, star, or tree structure.

Because the controller employed a poll-and-select access method, terminal devices connected to the loop required a minimum of intelligence. Although this reduced the cost of terminals connected to the loop, the controller lacked the intelligence to distribute the data flow evenly among terminals. A lengthy exchange between two terminal devices or between the controller and a terminal would thus tend to bog down this type of network structure. A second problem associated with this network structure was the centralized

placement of network control in the controller. If the controller failed, the entire network would become inoperative. Due to these problems, the use of loop systems is restricted to several niche areas, and they are essentially considered a derivative of a local area network.

Bus In a bus topology structure, a cable is usually laid out as one long branch, onto which other branches are used to connect each station on the network to the main data highway. Although this type of structure permits any station on the network to talk to any other station, rules are required for recovering from such situations as when two stations attempt to communicate at the same time. Later in this chapter, we will examine the relationships among the network topology, the method employed to access the network, and the transmission medium employed in building the network.

Ring In a ring topology, a single cable that forms the main data highway is shaped into a ring. As with the bus topology, branches are used to connect stations to one another via the ring. A ring topology can thus be considered to be a looped bus. Typically, the access method employed in a ring topology requires data to circulate around the ring, with a special set of rules governing when each station connected to the network can transmit data.

Star The fourth major local area network topology is the star structure, illustrated in the lower portion of Figure 1.4. In a star network, each station on the network is connected to a network controller. Then, access from any one station on the network to any other station can be accomplished through the network controller. Here, the network controller functions like a telephone switchboard, because access from one station to another station on the network can occur only through the central device. In fact, you can consider a telephone switchboard or PBX as representing a star-structured LAN whose trunks provide connections to the wide area network telephone infrastructure.

Tree A tree network structure represents a complex bus. In this topology, the common point of communications at the top of the structure is known as the *headend*. From the headend, feeder cables radiate outward to nodes, which in turn provide workstations with access to the network. There may also be a feeder cable route to additional nodes, from which workstations gain access to the network. One common example of a tree structure topology is the cable TV network many readers use on a daily basis. With the upgrade

of many cable TV systems to two-way amplifiers and the support of digital transmission, the local cable TV infrastructure can be considered to represent an evolving type of tree-structured local area network.

Mixed Topologies Some networks are a mixture of topologies. For example, as previously discussed, a tree structure can be viewed as a series of inter-connected buses. Another example of the mixture of topologies is a type of Ethernet known as 10BASE-T, which is described in detail in Chapter 3. That network can actually be considered a *star-bus* topology, because up to 16 or 24 devices known as *stations* are first connected to a common device known as a *hub*, which in turn can be connected to other hubs to expand the network.

Comparison of Topologies

Although there are close relationships among the topology of the network, its transmission media, and the method used to access the network, we can examine topology as a separate entity and make several generalized observations. First, in a star network, the failure of the network controller will render the entire network inoperative. This is because all data flow on the network must pass through the network controller. On the positive side, the star topology normally consists of telephone wires routed to a LAN switch. A local area network that can use in-place twisted-pair telephone wires in this way is simple to implement and usually very economical.

In a ring network, the failure of any node connected to the ring normally inhibits data flow around the ring. Due to the fact that data travels in a circular path on a ring network, any cable break has the same effect as the failure of the network controller in a star-structured network. Because each network station is connected to the next network station, it is usually easy to install the cable for a ring network. In comparison, a star network may require cabling each section to the network controller if existing telephone wires are not available, and this can result in the installation of very long cable runs.

In a bus-structured network, data is normally transmitted from a single station to all other stations located on the network, with a destination address included within each transmitted data block. As part of the access protocol, only the station with the destination address in the transmitted data block will respond to the data. This transmission concept means that a break in the bus affects only network stations on one side of the break that wish to communicate with stations on the other side of the break. Thus, unless a network station functioning as the primary network storage device becomes inoperative, a failure in a bus-structured network is usually less serious than a

failure in a ring network. However, some local area networks, such as Token-Ring and FDDI, were designed to overcome the effect of certain types of cable failures. Token-Ring networks include a backup path which, when manually placed into operation, may be able to overcome the effect of a cable failure between hubs (referred to as *multistation access units* or *MAUs*). In an FDDI network, a second ring can be activated automatically as part of a self-healing process to overcome the effect of a cable break.

A tree-structured network is similar to a star-structured network in that all signals flow through a common point. In the tree-structured network, the common signal point is the headend. Failure of the headend renders the network inoperative. This network structure requires the transmission of information over relatively long distances. For example, communications between two stations located at opposite ends of the network would require a signal to propagate twice the length of the longest network segment. Due to the propagation delay associated with the transmission of any signal, the use of a tree structure may result in a response time delay for transmissions between the nodes that are most distant from the headend.

Although the first type of Ethernet network was based on a bus-structured topology, other types of Ethernet networks incorporate the use of different topologies. Today you can select bus-based, star-bus, or tree-structured Ethernet networks. Thus, you can select a particular type of Ethernet network to meet a particular topology requirement.

Signaling Methods

The *signaling method* used by a local area network refers to both the way data is encoded for transmission and the frequency spectrum of the media. To a large degree, the signaling method is related to the use of the frequency spectrum of the media.

Broadband versus Baseband

Two signaling methods used by LANs are broadband and baseband. In *broadband signaling*, the bandwidth of the transmission medium is subdivided by frequency to form two or more subchannels, with each subchannel permitting data transfer to occur independently of data transfer on another subchannel. In *baseband signaling*, only one signal is transmitted on the medium at any point in time.

Broadband is more complex than baseband, because it requires information to be transmitted via the modulation of a carrier signal, thus requiring the use of special types of modems, as discussed later in this chapter.