


LOCAL AND METROPOLITAN AREA NETWORKS



Fourth Edition

William Stallings

EXHIBIT

Ex. 1011
(Part I)

THE WILLIAM STALLINGS BOOKS ON COMPUTER AND DATA COMMUNICATIONS TECHNOLOGY

DATA AND COMPUTER COMMUNICATIONS, THIRD EDITION

A comprehensive survey that has become the standard in the field, covering four main areas: (1) data communications, including transmission, media, signal encoding, link control, and multiplexing; (2) communication networks, including circuit- and packet-switched, local, packet radio, and satellite; (3) communications architecture, including the OSI model and related protocols; and (4) ISDN and broadband ISDN.

LOCAL AND METROPOLITAN AREA NETWORKS, FOURTH EDITION

An in-depth presentation of the technology and architecture of local and metropolitan area networks. Covers topology, transmission media, medium access control, standards, internetworking, and network management. Provides an up-to-date coverage of LAN/MAN standards.

ISDN AND BROADBAND ISDN: SECOND EDITION

An in-depth presentation of the technology and architecture of integrated services digital networks (ISDN). Covers the integrated digital network (IDN), ISDN services, architecture, signaling system no. 7 (SS7) and detailed coverage of the 1988 and 1990 CCITT standards. This new edition also provides coverage of frame relay and broadband ISDN topics including ATM and SONET.

COMPUTER ORGANIZATION AND ARCHITECTURE, THIRD EDITION

A unified view of this broad field. Covers fundamentals such as CPU, control unit, microprogramming, instruction set, I/O, and memory. Also covers advanced topics such as RISC, superscalar, and parallel organization.

BUSINESS DATA COMMUNICATIONS

A comprehensive presentation of data communications and telecommunications from a business perspective. Covers voice, data, image, and video communications and applications technology and includes a number of case studies.

THE WILLIAM STALLINGS BOOKS ON COMPUTER AND DATA COMMUNICATIONS TECHNOLOGY

OPERATING SYSTEMS

A state-of-the art survey of operating system principles. Covers fundamental technology as well as contemporary design issues, such as threads, real-time systems, multiprocessor scheduling, distributed systems, and security.

HANDBOOK OF COMPUTER-COMMUNICATIONS STANDARDS VOLUME 1 THE OPEN SYSTEMS INTERCONNECTION (OSI) MODEL AND OSI-RELATED STANDARDS, SECOND EDITION

A description of the master plan for all computer-communications standards: the OSI model. The book also provides a detailed presentation of OSI-related standards at all 7 layers, including HDLC, X.25, ISO internet, ISO transport, ISO session, ISO presentation, Abstract Syntax ONE (ASN.1), and common application service elements (CASE).

HANDBOOK OF COMPUTER-COMMUNICATIONS STANDARDS VOLUME 2 LOCAL AREA NETWORK STANDARDS, SECOND EDITION

A detailed examination of all current local network standards, including logical link control (LLC, IEEE 802.2), CSMA/CD (IEEE 802.3), token bus (IEEE 802.4), token ring (IEEE 802.5), and fiber distributed data interface (FDDI, ANS X3T9.5).

HANDBOOK OF COMPUTER-COMMUNICATIONS STANDARDS VOLUME 3 THE TCP/IP PROTOCOL SUITE, SECOND EDITION

A description of the protocol standards that are mandated on all DOD computer procurements and are becoming increasingly popular on commercial local network products, including TCP, IP, FTP, SMTP, and TELNET. The network management standards, SNMP and CMOT, are also presented.

Local and Metropolitan Area Networks

Local and Metropolitan Area Networks

FOURTH EDITION

William Stallings

MACMILLAN PUBLISHING COMPANY

New York

MAXWELL MACMILLAN CANADA

Toronto

MAXWELL MACMILLAN INTERNATIONAL

New York Oxford Singapore Sydney

Editor: John Griffin
Production Supervisor: John Travis
Production Manager: Roger Vergnes
Text Designer: Natasha Sylvester
Cover Designer: Robert Vega

This book was set in Palatino by Compset, Inc., and printed and bound by Book Press. The cover was printed by Lehigh Press.

Copyright © 1993 by Macmillan Publishing Company, a division of Macmillan, Inc.

Printed in the United States of America

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without permission in writing from the publisher.

Earlier editions entitled *Local Networks*, copyright © 1984, 1987 and 1990 by Macmillan Publishing Company.

Macmillan Publishing Company is part
of the Maxwell Communication Group of Companies.

Macmillan Publishing Company
866 Third Avenue, New York, New York 10022

Maxwell Macmillan Canada, Inc.
1200 Eglinton Avenue East
Suite 200
Don Mills, Ontario M3C 3N1

Library of Congress Cataloging-in-Publication Data

Stallings, William.

Local and metropolitan area networks / William Stallings. — 4th ed.

p. cm.

Rev. ed. of: *Local networks*. 3rd ed. © 1990.

Includes bibliographical references and index.

ISBN 0-02-415465-2

1. Local area networks (Computer networks) 2. Metropolitan area networks (Computer networks) I. Stallings, William. Local networks. II. Title.

TK5105.7.S77 1993

004.6—dc20

92-16096
CIP

Printing: 2 3 4 5 6 7 8 Year: 3 4 5 6 7 8 9 0 1

To my wife, Tricia

Preface

The local area network (LAN) has come to play a central role in information distribution and office functioning within businesses and other organizations. The major factor driving the widespread introduction of the LAN has been the proliferation of small computer systems, especially personal computers, but also including workstations and mini-computers.

With the dropping price of LAN hardware and software, LANs have become more numerous and larger, and they have taken on more and more functions within the organization. The upshot is that the LAN, once installed, quickly becomes almost as essential as the telephone system. At the same time, there is a proliferation of LAN types and options and a need to interconnect a number of LANs at the same site and with LANs at other sites. This has led to the development of LANs of higher and higher data rates and the relatively recent introduction of the metropolitan area network (MAN).

Objectives

This book focuses on the broad and evolving field of local and metropolitan area networks. The aim of the text is to provide a reasoned balance among breadth, depth, and timeliness. The book emphasizes topics of fundamental importance concerning the technology and architecture of these networks. Certain key related areas, such as perfor-

mance, internetworking, and network management are also treated in some detail.

The book explores the key topics in the field in the following general categories:

- *Technology and architecture:* There is a small collection of ingredients that serves to characterize and differentiate local and metropolitan area networks, including transmission medium, topology, communication protocols, switching technique, and hardware/software interface.
- *Network type:* It is convenient to classify the networks covered in this book into three types, based partly on technology and partly on application: These are local area network (LAN), metropolitan area network (MAN), and digital switch/digital private branch exchange (PBX).
- *Design approaches:* The book examines alternative design choices and assesses their relative merits.

Intended Audience

This book is intended for a broad range of readers interested in local networks:

- *Students and professionals in computer science and data communications:* The book is intended as both a textbook for study and a basic reference volume for this exciting area within the broader fields of computer science and data communications.
- *Local network designers and implementors:* The book discusses the critical design issues and illustrates alternative approaches to meeting user requirements.
- *Local network customers and system managers:* The book alerts the reader to some of the key issues and tradeoffs, and what to look for in the way of network services and performance.

The book is intended for both an academic and a professional audience. As a textbook, it is intended as a one-semester course. It covers much of the material in the Computer Communication Networks course of The joint ACM/IEEE Computing Curricula 1991.

The book also serves as a basic reference volume and is suitable for self-study. For the reader with little or no background in data communications, a brief primer is included.

Plan of the Text

The book is organized to clarify both the unifying and the differentiating concepts that underlie this field. The organization of the chapters is as follows:

1. *Introduction*: The chapter defines the term *local network* and looks at some of the applications and benefits.
2. *Topics in Data Communications and Computer Networking*: This necessarily brief survey explains the relevant concepts used throughout the book.
3. *Overview of LAN/MAN Technology*: Introduces the key elements of transmission medium and topology. A classification of networks into LANs, MANs, and WANs is developed.
4. *Topologies and Transmission Media for LANs and MANs*: Examines the design issues relating to the implementation of LANs and MANs, with emphasis on the topology and transmission medium alternatives.
5. *Local Area Network Architecture*: Describes the logical link control and medium access control architecture of LANs. LAN standards are also described.
6. *Metropolitan Area Network Architecture*: Treats the two most important MANs: FDDI and IEEE 802.6.
7. *Circuit-Switched Local Networks*: Networks in this category constitute the major alternative to LANs for meeting general local interconnection needs. The category includes the data-only digital switch and the voice/data digital private branch exchange (PBX).
8. *The Network Interface*: The nature of the interface between an attached device and a LAN or MAN is an important design issue. This chapter explores some alternatives.
9. *LAN/MAN Performance*: This chapter gives some insight into the performance problems and the differences in performance of various LANs and MANs.
10. *Internetworking*: In the majority of cases, LANs will be connected in some fashion to other networks, either by means of other LANs, by way of MANs, or using wide-area networks. The key alternatives of bridge and router are explored.
11. *Network Management*: Network management tools and systems are indispensable for LANs and MANs. This chapter explores the types of systems that are available and examines the standards developed for general network management and LAN/MAN management.

In addition, the book includes an extensive glossary, a list of frequently-used acronyms, and a bibliography. Each chapter includes problems and suggestions for further reading.

Throughout, there is a heavy emphasis on standards, including standards based on the Open Systems Interconnection (OSI) model and specific LAN and MAN standards, such as IEEE 802 and FDDI. This emphasis reflects the growing importance of such standards in defining the available products and future research directions in this field.

Related Materials

The author has produced other material that may be of interest to students and professionals. *Advances in Local and Metropolitan Area Network Technology* (1993, IEEE Computer Society Press, 10662 Los Vaqueros Circle, P.O. Box 3014, Los Alamitos, CA 90720, telephone 714-821-8380) is a companion to this text, and follows the same topical organization. It contains reprints of many of the key references used herein.

A set of videotape courses specifically designed for use with this book is available from The Media Group, Boston University, 565 Commonwealth Avenue, Boston, MA 02215; telephone (617) 353-3227.

Data and Computer Communications, Third Edition (Macmillan, 1991) covers fundamental concepts in the areas of data transmission, communication networks, and computer-communications protocols. *ISDN and Broadband ISDN, Second Edition* (Macmillan, 1992) covers the concepts and technology of integrated services digital networks (ISDN) and broadband ISDN, which are all-digital networks gradually being introduced to replace existing wide-area networks. *Networking Standards* (Addison-Wesley, 1993) covers the leading-edge standards that are defining the networks and distributed applications recently introduced or currently under development.

The Fourth Edition

I began work on the first edition of this book in 1982. At the time of its publication, it was the only book-length technical treatment of LANs (and remains the only textbook on the subject). Little did I anticipate that it would still be going strong over a decade later. To paraphrase a recent Oscar-winner, you like this book! You really like it! Any author is bound to feel a sense of pride and satisfaction on being asked to produce a fourth edition of a book that covers such a fast-moving field as this one. The book has withstood the test of time, and its success confirms that the basic organization and emphasis of the book is sound. However, because the field is fast-moving, each new edition requires a major revision to keep up.

This edition is no exception. The revision in this case even extends to the title, which now includes the phrase Metropolitan Area Networks. The inclusion of MANs is dictated by three developments:

1. The fiber distributed data interface (FDDI), which is generally referred to as a local area network (LAN), is finding increasing application as a backbone MAN, thanks to the increased demand for this service and the maturing of bridge and router technologies.
2. FDDI-II has been adopted. This revision of FDDI addresses some of the integrated-voice data requirements of a MAN.

3. After an almost uncountable number of false starts, the IEEE 802.6 committee has finally settled on a MAN standard, and that standard has received broad industry and customer support.

So the inclusion of MANs is a major new addition to the book. Another significant change in this edition is the revision of the chapter on internetworking. The material on bridge routing standards has been expanded. The spanning tree and source routing approaches, introduced in the third edition, receive expanded coverage. In addition, a new bridge standard, SRT, is introduced. In the area of routing, the new routing protocol standards, ES-IS and IS-IS are examined.

A final major change in this edition is the expansion of the coverage on network management to an entire chapter. The chapter covers the basis technology of network management systems and local network management. In addition, the ISO network management standards and the IEEE 802 LAN/MAN management standards are covered.

In addition to these major changes, there have been expansions and updates in every chapter. To give some feel for the overall scope of this revision, approximately 33% of the tables, 30% of the figures, and 24% of the references in this edition are new. All in all, this fourth edition constitutes a major revision. I have tried in a balanced manner to provide a comprehensive survey of the technology and architecture of local and metropolitan area networks.

Acknowledgment

My association with Macmillan's college division now stretches back over more than a decade. I have always had the strong and enthusiastic support of the division's staff and am grateful for all the support and encouragement I have received over the years. Two people in particular I would like to thank.

In a changing world, it is remarkable that the production editor for every one of my dozen books, going back to the first edition of this book, has been the same man: John Travis. Over the years, John has caught many errors, both editorial and—more important—technical, and he has managed to bring every single one of these books out on time. Quite an achievement.

My current, and I hope permanent, editor is John Griffin. His feel for both the technical and marketing side of the business has helped direct my writing into the most fruitful channels.

Of course, I have dealt with many other people in the College Division over the years. The names have changed from time to time, but the supportive atmosphere and the professionalism have not.

W. S.

Contents

Preface	vii
<i>CHAPTER 1</i>	
Introduction	1
1.1 LANs, MANs, and WANs	1
1.2 Benefits and Pitfalls	6
1.3 Applications	7
1.4 Information Distribution	16
1.5 Outline of the Book	18
1.6 Recommended Reading	21
1.7 Problems	21
<i>CHAPTER 2</i>	
Topics in Data Communications and Computer Networking	23
2.1 Data Communications Concepts	23
2.2 Communication Switching Techniques	37
2.3 Computer Networking	48
2.4 Recommended Reading	67
2.5 Problems	68
Appendix 2A: The Cyclic Redundancy Check	69

CHAPTER 3**Overview of LAN/MAN Technology** 73

- 3.1 Topologies 73
- 3.2 Transmission Media 78
- 3.3 Relationship Between Medium and Topology 90
- 3.4 Classes of Networks 95
- 3.5 Recommended Reading 100
- 3.6 Problems 100

CHAPTER 4**Topologies and Transmission Media for LANs and MANs** 103

- 4.1 Metallic Media: Bus/Tree Topology 103
- 4.2 Metallic Media: Star Topology 120
- 4.3 Metallic Media: Ring Topology 123
- 4.4 Optical Fiber Star 133
- 4.5 Optical Fiber Ring 138
- 4.6 Optical Fiber Bus 139
- 4.7 Recommended Reading 142
- 4.8 Problems 142
- Appendix 4A: Characteristic Impedance** 143
- Appendix 4B: Decibels** 146
- Appendix 4C: Scrambling and Descrambling** 147

CHAPTER 5**Local Area Network Architecture** 151

- 5.1 LAN Protocols 151
- 5.2 Link Layer Protocol for LANs 157
- 5.3 Medium Access Control—Bus/Tree 170
- 5.4 Medium Access Control—Ring 193
- 5.5 Recommended Reading 205
- 5.6 Problems 206
- Appendix 5A: IEEE 802 Standards** 207
- Appendix 5B: Service Primitives and Parameters** 215

CHAPTER 6**Metropolitan Area Network Architecture** 219

- 6.1 FDDI 219
- 6.2 FDDI-II 241
- 6.3 IEEE 802.6 247
- 6.4 Recommended Reading 272
- 6.5 Problems 273

*CHAPTER 7***Circuit-Switched Local Networks** 275

- 7.1 Star Topology Networks 275
- 7.2 Digital Switching Concepts 277
- 7.3 Digital Data Switching Devices 291
- 7.4 The Digital Private Branch Exchange 297
- 7.5 Digital PBX Versus LAN 313
- 7.6 Recommended Reading 315
- 7.7 Problems 315

*CHAPTER 8***The Network Interface** 319

- 8.1 The Requirement 319
- 8.2 Packet-Switched Interfacing 322
- 8.3 The Device/NIU Interface 327
- 8.4 Terminal Handling for LANs 340
- 8.5 Circuit-Switched Networks 353
- 8.6 Analog Devices 354
- 8.7 Recommended Reading 354
- 8.8 Problems 354

*CHAPTER 9***LAN/MAN Performance** 357

- 9.1 LAN/MAN Performance Considerations 358
- 9.2 LAN Performance 368
- 9.3 MAN Performance 395
- 9.4 End-to-End Performance 403
- 9.5 Recommended Reading 407
- 9.6 Problems 409

*CHAPTER 10***Internetworking** 411

- 10.1 Bridges 412
- 10.2 Routing with Bridges 417
- 10.3 Routers 450
- 10.4 Routing with Routers 464
- 10.5 Recommended Reading 474
- 10.6 Problems 474

CHAPTER 11

Network Management	477
11.1 Network Management Requirements	478
11.2 Network Management Systems	483
11.3 OSI Network Management	485
11.4 LAN-Specific Network Management	496
11.5 IEEE 802 LAN/MAN Management	502
11.6 FDDI Management	510
11.7 Recommended Reading	514
11.8 Problems	514
Glossary	515
References	523
Index	539

Local and Metropolitan Area Networks

CHAPTER 1

Introduction

1.1

LANs, MANs, AND WANs

For businesses, government agencies, universities, and other organizations, data communications networks have become indispensable. Of most importance are networks that interconnect equipment within a single building or a group of buildings. For want of a better term, we will refer to such networks as *local networks*. In fact, this book is concerned with three types of local networks: local area networks (LANs), metropolitan area networks (MANs), and circuit-switching local networks. Before defining these terms, we need to understand the trends responsible for the importance of local networks.

Of most importance is the dramatic and continuing decrease in computer hardware costs, accompanied by an increase in computer hardware capability. Today's microprocessors have speeds, instruction sets, and memory capacities comparable to the most powerful minicomputers of a few years ago. This trend has spawned a number of changes in the way information is collected, processed, and used in organizations. There is increasing use of small, single-function systems, such as word processors and small business computers, and of general-purpose microcomputers, such as personal computers and Unix-based multiuser workstations. These small, dispersed systems are more accessible to the

user, more responsive, and easier to use than large central time-sharing systems.

All of these factors lead to an increased number of systems at a single site: office building, factory, operations center, and so on. At the same time there is likely to be a desire to interconnect these systems for a variety of reasons, including:

- To share and exchange data between systems
- To share expensive resources

The ability to exchange data is a compelling reason for interconnection. Individual users of computer systems do not work in isolation and will want to retain some of the benefits provided by a central system. These include the ability to exchange messages with other users, the ability to access data from several sources in the preparation of a document or for an analysis, and the opportunity for multiple users to share information in a common file.

To appreciate the second reason, consider that although the cost of data processing hardware has dropped, the cost of essential electromechanical equipment, such as bulk storage and line printers, remains high. In the past, with a centralized data processing facility, these devices could be attached directly to the central host computer. With the dispersal of computer power, these devices must somehow be shared.

A Definition of Local Networks

We will elaborate on these and other reasons later in this chapter. For now, the discussion above should be enough to motivate the following definition of a *local network*:

A local network is a communications network that provides interconnection of a variety of data communicating devices within a small area.

There are three elements of significance in this definition. First, a local network is a communications network. That is, it is a facility for moving bits of data from one attached device to another. The application-level software and protocols that are required for attached devices to function cooperatively are beyond the scope of this book. As a corollary to this definition, note that a collection of devices interconnected by individual point-to-point links is not included in the definition or in this book.

Second, we interpret the phrase *data communicating devices* broadly, to include any device that communicates over a transmission medium. Examples:

- Computers
- Terminals

- Peripheral devices
- Sensors (temperature, humidity, security alarm sensors)
- Telephones
- Television transmitters and receivers
- Facsimile

Of course, not all types of local networks are capable of handling all of these devices.

Third, the geographic scope of a local network is small. The most common occurrence is a network that is confined to a single building. Networks that span several buildings, such as on a college campus or military base, are also common. A borderline case is a network with a radius of a few tens of kilometers. With appropriate technology, such a system will behave like a local network.

Another element that could be added to the definition is that a local network is generally privately owned rather than a public or commercially available utility. Indeed, typically, a single organization will own both the network and the attached devices.

Some of the typical characteristics of local networks are:

- High data rates (0.1 to 100 Mbps)
- Short distances (0.1 to 25 km)
- Low error rate (10^{-8} to 10^{-11})

The first two parameters serve to differentiate local networks from two cousins: multiprocessor systems and wide-area networks.

Other distinctions can be drawn between local networks and their two cousins, and these have a significant impact on design and operation. Local networks generally experience significantly fewer data transmission errors and significantly lower communications costs than those of long-haul networks. Cost-performance trade-offs are thus significantly different. Also, because local networks are generally owned by the same organization as the attached devices, it is possible to achieve greater integration between the network and the devices; this topic is explored in Chapter 8.

A distinction between local networks and multiprocessor systems is the degree of coupling. Multiprocessor systems are tightly coupled, usually have some central control, and completely integrate the communications function. Local networks tend to exhibit the opposite characteristics.

Types of Local Networks

There are two basic types of local networks: those based on circuit switching and those based on a technology referred to as packet broadcasting (Figure 1.1). We will define the terms *circuit switching* and *packet*




FIGURE 1.1 Transmission Methods for Local Networks

broadcasting in more detail in Chapter 2. For now, a brief definition of each should suffice:

- *Packet broadcasting:* Devices share a communications network in which a transmission from any one device is heard by all other devices. Data to be transmitted are broken up into small blocks, called *packets*. Packets include both user data and control information that indicate the destination of the data. Each packet is sent onto a network and may be received by all other devices on the network.
- *Circuit switching.* The network consists of a central switch to which all devices attach. Two devices communicate by setting up a circuit through the switch. The circuit consists of a path and dedicated resources for transferring data between the two devices through the switch.

The key to packet broadcasting is the use of a transmission medium shared by a number of devices. An early example of the use of a shared transmission medium is the multidrop line. The multidrop line, how-

ever, is used to permit communication between one primary station (usually a host computer) and a number of secondary stations (usually terminals). Communication on the multidrop line is controlled by the primary, and secondary-to-secondary exchange is generally not allowed. For a local network, peer communication among a number of cooperating devices is required. This type of local network is generally referred to as a **local area network (LAN)** and has the following key characteristics:

- A transmission medium is shared among the attached devices.
- Transmission is in the form of packets.
- A transmission from any one station is received by all other stations (hence the term *packet broadcasting*).
- There is no master station; rather, all of the stations cooperate to assure orderly use of the transmission medium.

In recent years, a new type of network, referred to as a **metropolitan area network (MAN)**, has been developed. A metropolitan area network shares the characteristics listed above with the LAN; the difference is that the MAN covers larger distances and, generally, operates at higher data rates.

The most familiar example of a circuit-switching local network is the **private branch exchange (PBX)**. The PBX was originally developed to provide an on-premise telephone exchange system. The voice PBX provides a point of interconnection for extension telephones within the office and a trunk connection to the nearest central office telephone exchange. Calls within the office are made through the PBX; calls outside the office are directed by the PBX to the public telephone network or a leased line.

With the advent of digital technology, the **digital PBX** has appeared on the scene and now dominates the PBX market. The digital PBX handles all signals internally as digital signals but still uses circuit-switching technology. The digital PBX is suited to handle both voice and data connections.

A final example of a local network that employs circuit switching is the **digital data switch**. The digital data switch is designed specifically to deal with data rather than voice. The main difference between the digital data switch and the digital PBX is that the former does not contain many of the call processing features normally found in the digital PBX, such as call forwarding and camp-on.

All the above networks can be distinguished from **wide-area networks (WANs)**. As the name implies, WANs are networks that cover substantial distances. Public telephone networks and packet-switching networks are examples of WANs.

The focus of this book is on LANs and MANs, with a chapter devoted to circuit-switching local networks.

1.2

BENEFITS AND PITFALLS

Table 1.1 lists some of the major benefits of a local network. Whether these are realized or not, of course, depends on the skill and wisdom of those involved in selecting and managing the local network.

One of the most important potential benefits of a local network relates to system evolution. In a nonnetworked installation such as a time-sharing system, all data processing power is in one or a few systems. In order to upgrade hardware, existing applications software must be either converted to new hardware or reprogrammed, with the risk of error in either case. Even adding new applications on the same hardware, or enhancing those that exist, involves the risk of introducing errors and reducing the performance of the entire system. With a local network it is possible to gradually replace applications or systems, avoiding the "all-or-nothing" approach. Another facet of this capability is that old equipment can be left in the system to run a single application if the cost of moving that application to a new machine is not justified.

A local network tends to improve the reliability, availability, and survivability of a data processing facility (see Section 12.2). With multiple interconnected systems, the loss of any one system should have minimal impact. Further, key systems can be made redundant so that other systems can quickly take up the load after a failure.

We have already mentioned resource sharing. This includes not only expensive peripheral devices, but data. Data may be housed and con-

TABLE 1.1 Benefits and Pitfalls of Local Networks

Potential Benefits

System evolution: incremental changes with contained impact
 Reliability/availability/survivability: multiple interconnected systems disperse functions and provide backup capability
 Resource sharing: expensive peripherals, hosts, data
 Multivendor support: customer not locked in to a single vendor
 Improved response/performance
 User needs single terminal to access multiple systems
 Flexibility of equipment location
 Integration of data processing and office automation

Potential Pitfalls

Interoperability is not guaranteed: software, data
 A distributed database raises problems of integrity, security/privacy
 Creeping escalation: more equipment will be procured than is actually needed
 Loss of control: more difficult to manage and enforce standards

trolled from a specific facility but, via the network, may be available to many users.

A local network provides at least the potential of connecting devices from multiple vendors, thus giving the customer greater flexibility and bargaining power. However, a local network will provide only a rather low or primitive level of interconnection. For the network to function properly, higher levels of networking software must be supplied within the attached devices (see Section 2.3 and Chapter 8).

These are, in most cases, the most significant benefits of a local network. Several others are also listed in Table 1.1.

Alas, there are also some pitfalls, or at least potential pitfalls. As we mentioned, a local network does not guarantee that two devices can be used cooperatively, a concept known as *interoperability*. For example, two word processors from different vendors can be attached to a local network and can perhaps exchange data. But they probably will use different file formats and different control characters, so that it is not possible, directly, to take a file from one and begin editing it on the other. Some sort of format-conversion software is needed.

With a local network, it is likely that data will be distributed or, at least, that access to data may come from multiple sources. This raises questions of integrity (e.g., two users trying to update the database simultaneously) and security and privacy.

Another pitfall might be referred to as "creeping escalation." With the dispersal of computer equipment and the ease of incrementally adding equipment, it becomes easier for managers of suborganizations to justify equipment procurement for their department. Although each procurement may be individually justifiable, the totality of procurements throughout an organization may well exceed the total requirements.

There is also a loss of control problem. The prime virtue of networking—distributed systems—is also its prime danger. It is difficult to manage this resource, to enforce standards for software and data, and to control the information available through a network.

1.3

APPLICATIONS

The range of applications for local networks is wide, as indicated by the broad definition given above. Table 1.2 lists some of the potential applications. Again, we emphasize that not all local networks are capable of supporting all applications. To give some feeling for the use of local networks, we discuss in this section five rather different types of applications.

TABLE 1.2 Local Network Applications

Data Processing	Energy management
Data entry	Heating
Transaction processing	Ventilation
File transfer	Air conditioning
Inquiry/response	Process control
Batch/RJE	Fire and security
Office automation	Sensors/alarms
Document/word processing	Cameras and monitors
Electronic mail	Telephones
Intelligent copying/facsimile	Teleconferencing
Factory automation	Television
CAD/CAM	Off-the-air
Inventory control/order entry/shipping	Video presentations
Monitor and control of factory floor equipment	

Personal Computer Networks

We start at one extreme, a system designed to support microcomputers, such as personal computers. With the relatively low cost of such systems, individual managers within organizations are independently procuring personal computers for stand-alone applications, such as spreadsheet and project management tools. Today's personal computers put processor, file storage, high-level languages, and problem-solving tools in an inexpensive, "user-friendly" package. The reasons for acquiring such systems are compelling.

But a collection of stand-alone processors will not meet all of an organization's needs; central processing facilities are still required. Some programs, such as econometric forecasting models, are too big to run on a small computer. Corporate-wide data files, such as accounting and payroll, require a centralized facility but should be accessible to a number of users. In addition, there are other kinds of files that, although specialized, must be shared by a number of users. Further, there are sound reasons for connecting individual intelligent workstations not only to a central facility but to each other as well. Members of a project or organizational team need to share work and information. By far the most efficient way to do so is electronically.

Figure 1.2 is an example of a local network of personal computers for a hypothetical engineering group or department. The figure shows four types of users who have personal computers, each equipped with particular applications.

Each type of user is provided with electronic mail and word processing to improve the efficiency of creating and distributing messages, memos, and reports. Managers are also given a set of program and bud-

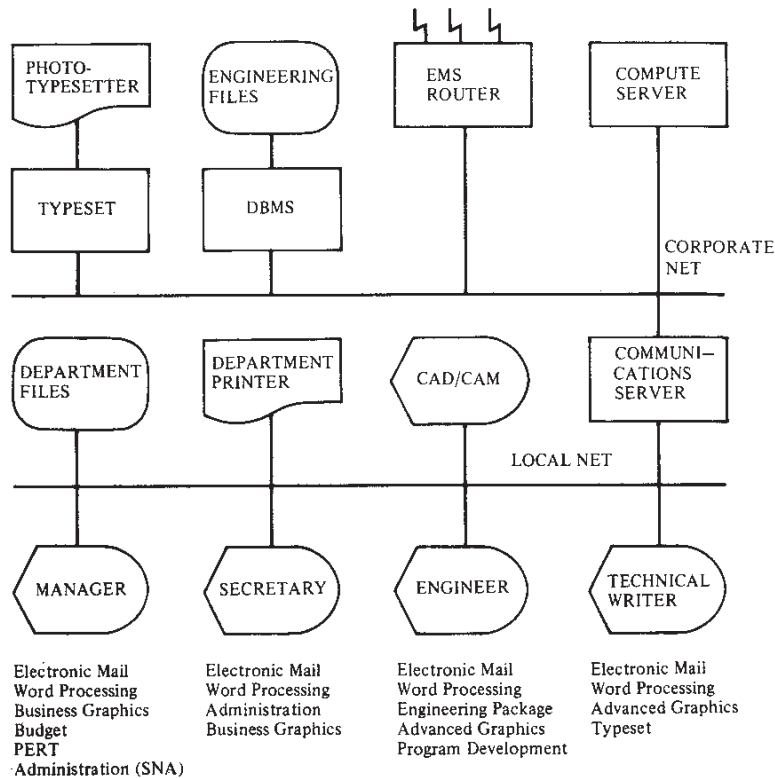


FIGURE 1.2 Personal Computers in Support of a Working Team

get management tools. With the amount of automation that personal computers supply, the role of secretaries becomes less that of a typist and more that of an administrative assistant. Tools such as electronic calendar and graphics support become valuable for these workers. In the same fashion, engineers and technical writers can be supplied with tailored systems.

Certain expensive resources, such as a disk and printer, can be shared by all users of the departmental local network. In addition, the network can tie into larger corporate network facilities. For example, the corporation may have a building-wide local network (see Office Automation below) and a long-haul corporate-wide network using, for example, IBM's SNA. A communications server can provide controlled access to these resources.

A key requirement for the success of such a network is low cost. The cost of attachment to the network for each device should be on the order

of one to a few hundred dollars; otherwise, the attachment cost will approach the cost of the attached device. However, the capacity and data rate need not be high, so this is a realizable goal. For example, see [THUR85].

Computer Room Networks

At the other extreme from a personal computer local network is one designed for use in a computer room containing large, expensive mainframe computers. This type of network is likely to find application at very large data processing sites. Typically, these sites will be large companies or research installations with large data processing budgets. Because of the size involved, a small difference in productivity can mean millions of dollars. Further, although such networks are few in number, the collective cost of the equipment they support is very high. Consequently, this type of application deserves a close look.

Consider a site that uses a dedicated mainframe computer. This implies a fairly large application or set of applications. As the load at the site grows, the existing model may be replaced by a more powerful one, perhaps a multiprocessor system. At some sites, a single-system replacement will not be able to keep up; equipment growth rates will be exceeded by demand growth rates. The facility will eventually require multiple independent computers. Again, there are compelling reasons for interconnecting these systems. The cost of system interrupt is high, so it should be possible, easily and quickly, to shift applications to backup systems. It must be possible to test new procedures and applications without degrading the production system. Large bulk storage files must be accessible from more than one computer. Load leveling should be possible to maximize utilization.

An example of this type of installation is the one at the National Center for Atmospheric Research (NCAR), shown in Figure 1.3. This entails storage of massive amounts of data and the use of huge number-crunching simulation and analysis programs. There is also an extensive on-site graphics facility.

Initially, the NCAR facility consisted of a single mainframe run in batch mode. When it became clear that additional batch machines were needed, NCAR investigated the requirements for a new configuration to meet their needs. The result was four objectives:

1. Provide front-end processors to remove job and file preparation tasks from the batch computers
2. Provide an efficient method for interactive processing
3. Design a system architecture that would allow different services for special needs and purposes

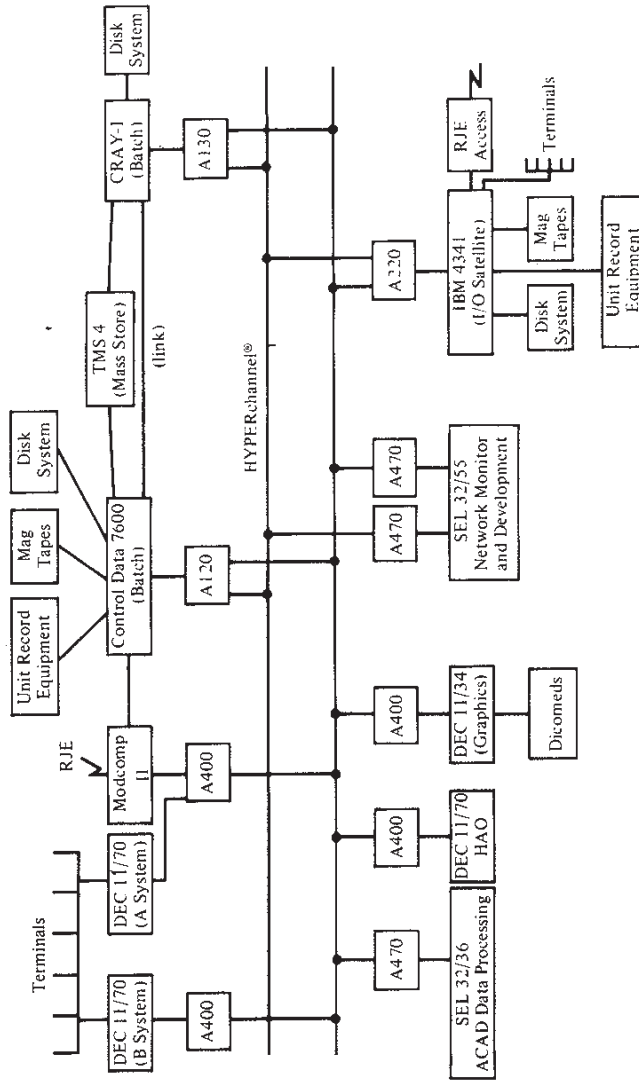


FIGURE 1.3 A Computer Room Network

4. Provide a system that would allow configuration flexibility without excessive modification to existing resources

The result of this study was a plan that called for the procurement of front-end processors, special-purpose computers, and bulk storage systems. A network was needed that met two requirements:

- Easy addition and subtraction of equipment
- High sustained data transfer speeds

It can be seen that some key requirements for computer-room networks are the opposite of those for personal computer local networks. High data rates are required to keep up with the work, which typically involves the transfer of large blocks of data. The electronics for achieving high speeds are expensive, on the order of tens of thousands of dollars per attachment. Fortunately, given the much higher cost of attached devices, such costs are reasonable.

Office Automation

Most local network applications will fall between these two extremes. Moderate data rates and moderate attachment costs are requirements. In some cases, the local network will support one or a few types of devices and rather homogeneous traffic. Others will support a wide variety of device and traffic types.

A good generic example of the latter is an office automation system, which can be defined as the incorporation of appropriate technology to help people manage information.

The key motivation for the move to office automation is productivity. As the percentage of white-collar workers has increased, the information and paper-work volume has grown. In most installations, secretarial and other support functions are heavily labor intensive. Increased labor costs combined with low productivity and increasing work load have caused employers to seek effective ways of increasing their rather low capital investment in this type of work.

At the same time, principals (managers, skilled "information workers") are faced with their own productivity bind. Work needs to be done faster with less wait time and waste time between segments of a job. This requires better access to information and better communication and coordination with others.

Table 1.3 lists elements of a hypothetical integrated office automation system. A study of this list gives some idea of the range of devices and information types that are part of the system. For this system to work and be truly effective, a local network is needed that can support the

TABLE 1.3 Elements of an Integrated Office Automation System

Basic IOAS Components	Optional IOAS Components
<p>Action elements</p> <ul style="list-style-type: none"> Word management (keying and editing) Terminal-oriented computer-based message system Automated file indexing Electronic filing and retrieval Off-line connection to computer-operated micrographics (for system purging) <p>Control elements</p> <ul style="list-style-type: none"> Electronic calendar Electronic tickler file <p>Inquiry elements</p> <ul style="list-style-type: none"> Automated file searching and retrieval Directory of users (names, addresses, telephone numbers, etc.) Capability for open-loop computer-aided retrieval (CAR) of micrographics Capability for input/output control of physical files <p>Extended-Application IOAS Components</p> <p>Action elements</p> <ul style="list-style-type: none"> Automated departmental billing for IOAS usage Individual applications Personal computing (permits individual to program) Unit applications Departmental applications Divisional applications Regional applications Line-of-business applications Functional applications (mathematical formulas) <p>Control elements</p> <ul style="list-style-type: none"> System usage monitoring (departmental level) Specialized applications (as above) <p>Inquiry elements: specialized applications (as above)</p>	<p>Action elements</p> <ul style="list-style-type: none"> Interconnection to other terminal-oriented, computer-based message systems Interconnection to public teletypewriter systems OCR (optical character recognition) input Digitized, hard-copy input (temporary; for incoming mail) Store-and-forward fax Soft-copy fax Interconnection to external fax devices and networks Audio output electronic mail (digital-to-audio conversion) Business graphics (black-and-white) Electronic calculator Sorting capabilities Photocomposer output On-line output of computer-operated micrographics (COM) Computer teleconferencing <p>Control elements</p> <ul style="list-style-type: none"> COM format previewing Project management and control Management of multiauthored document preparation <p>Inquiry elements</p> <ul style="list-style-type: none"> Soft-copy CAR Electronic publishing (manuals, price lists, news, etc.) Interconnection to other internal systems and data bases Interconnection to external research data base services

Source: [BARC81].

various devices and transmit the various types of information. A discussion of the use of local networks to tie together office automation equipment such as this can be found in [STAL90c].

Factory Local Networks

The factory environment is increasingly being dominated by automated equipment: programmable controllers, automated materials handling devices, time and attendance stations, machine vision devices, and various forms of robots. To manage the production or manufacturing process, it is essential to tie this equipment together. And, indeed, the very nature of the equipment facilitates this. Microprocessor devices have the potential to collect information from the shop floor and accept commands. With the proper use of the information and commands, it is possible to improve the manufacturing process and to provide detailed machine control.

The more that a factory is automated, the greater is the need for communications. Only by interconnecting all of the devices and by providing mechanisms for their cooperation can the automated factory be made to work. The means of interconnection is the factory local area network [SCHO84, MCGA85, HALL85].

To get some feeling for the requirements for a factory local network, consider the requirements developed by General Motors [STAL90a]. GM's specification of a communications network is driven by the sophisticated communications strategy it has evolved to meet its requirements. These requirements reflect those that obtain in other factory and robotics environments. Among the key areas are the following:

- Work force involvement has proven to be a valuable tool for GM's quality and cost-improvement effort. In an attempt to provide facts about the state of the business, employees are told GM's competitive position in relation to quality and costs. This information is communicated by video setups at numerous locations in the plant complex.
- An indirect effect on manufacturing costs has been the escalating cost of utilities. To try to control this area, GM measures usage of water, gas, pressurized air, steam, electricity, and other resources—often by means of computers and programmable controllers.
- GM is investigating and, in some cases, implementing asynchronous machining and assembly systems that are much more flexible than the traditional systems of the past. To facilitate flexibility, the communication requirements increase an order of magnitude.

- To protect its large investment in facilities, GM uses closed-circuit TV surveillance and computerized monitoring systems to warn of fires or other dangers.
- Accounting systems, personnel systems, material and inventory control systems, warranty systems, and others use large mainframe computers with remote terminals located throughout the manufacturing facility.
- The nature of process-control and factory environments dictates that communications be extremely reliable and that the maximum time required to transmit critical control signals and alarms be bounded and known.

To innerconnect all of the equipment in a facility, a local network is needed. The requirements listed above dictate the following characteristics of the local network:

- High capacity
- Ability to handle a variety of data traffic
- Large geographic extent
- High reliability
- Ability to specify and control transmission delays

Integrated Voice and Data Local Networks

In virtually all offices today, the telephone system is separate from any local network that might be used to interconnect data processing devices. With the advent of digital voice technology, the capability now exists to integrate the telephone switching system of a building with the data processing equipment, providing a single local network for both.

Such integrated voice/data networks might simplify network management and control. It will also provide the required networking for the kinds of integrated voice and data devices to be expected in the future. An example is an executive voice/data workstation that provides verbal message storage, voice annotation of text, and automated dialing.

Summary

This section has only scratched the surface of possible applications of local networks. This book focuses on the common principles underlying the design and implementation of all local networks, and so will not pursue the topic of specific applications. Nevertheless, in the course of the book, the reader will gain an appreciation of the variety of uses of local networks.

1.4

INFORMATION DISTRIBUTION

In determining the requirements for local networking, it is important to examine the traffic patterns that are reasonable to expect. Figure 1.4 illustrates the distribution of nonvoice information that has been consistently reported in a number of studies. About half of the information generated within a small unit of an organization (e.g., a department) remains within that unit. Typically only summary-type information or consolidated data are disseminated beyond the basic unit of an organization. Another 25% is normally shared with peer departments within a somewhat larger grouping (e.g., a division) and the immediate superior of the department. In a typical office layout, this would translate to a radius of about 600 feet. Another 15% goes elsewhere within the organization, such as to other departments within other divisions, central staff organizations, and top management. Finally, only about 10% of the total generated information is distributed beyond the confines of a single building or cluster of buildings. Example destinations include remote corporate headquarters, customers, suppliers, and government agencies.

Another way of looking at local network requirements is to consider the kinds of data processing equipment to be supported. In rough terms, we can group this equipment into three categories:

1. *Personal computers and terminals*: the workhorse in most office environments is the microcomputer, including personal computers and workstations. Additionally, when shared systems are present

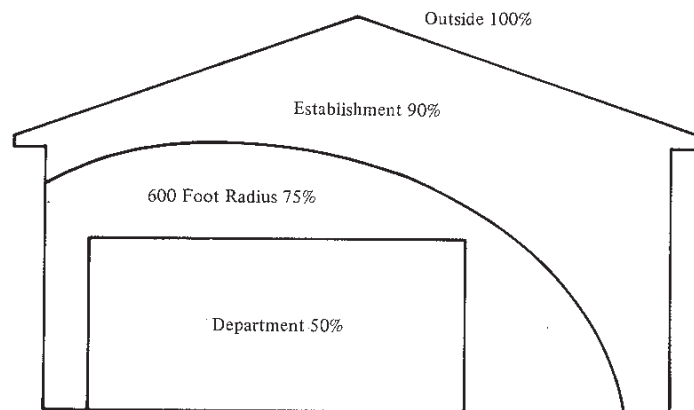


FIGURE 1.4 Information Distribution

in an organization, terminals are also to be found. Most of this equipment is found at the departmental level, used by individual professionals and secretarial personnel. When used for network applications, the load generated tends to be rather modest.

2. *Minicomputers*: minicomputers may function as servers within a department or be shared by users in a number of departments. In many organizations, a number of commonly used applications will be provided on time-shared minicomputers. Because of this shared use, these machines may generate more substantial traffic than microcomputers.
3. *Mainframes*: for large database and scientific applications, the mainframe is still the machine of choice. When the machines are networked, bulk data transfers dictate that a high-capacity network be used.

Figure 1.5 illustrates the performance spectrum involved. Larger, more expensive machines tend to require a higher data rate on the local network to support them. The higher the data rate, the greater the cost of the network.

The requirements indicated by Figures 1.4 and 1.5 suggest that a single local network will not, in many cases, be the most cost-effective solution. A single network would have to be rather high speed to support the aggregate demand. However, the cost of attachment to a local net-




FIGURE 1.5 Office Network Performance Spectrum

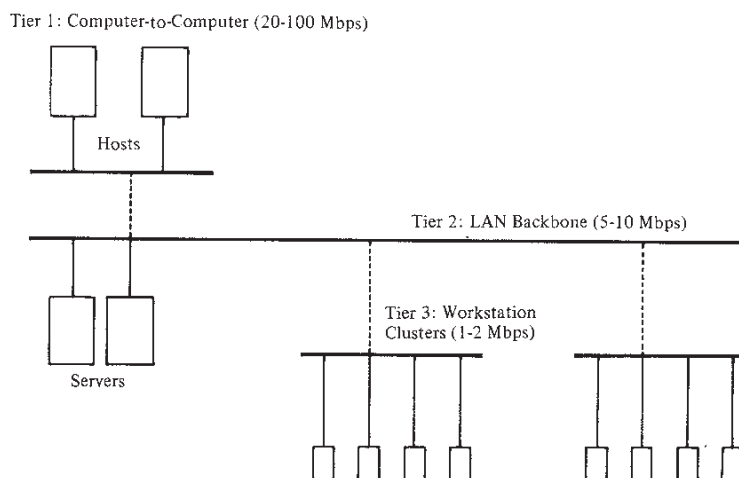


FIGURE 1.6 Tiered Local Networks

work tends to increase as a function of the network data rate. Accordingly, a high-speed local network would be very expensive for attachment of low-cost personal computers.

An alternative approach, which is becoming increasingly common, is to employ two or three tiers of local networks (Figure 1.6). Within a department, a low-cost, low-speed local network supports a cluster of microcomputers and terminals. These departmental local networks are then lashed together with a backbone local network of higher capacity. In addition, shared systems are also supported off of this backbone. If mainframes are also part of the office equipment suite, then a separate high-speed local network supports these devices and may be linked, as a whole, to the backbone local network to support a modest amount of traffic between the mainframes and other office equipment. We will see that local network standards and products address the need for all three types of local networks.

1.5

OUTLINE OF THE BOOK

This chapter, of course, serves as an introduction to the entire book. A brief synopsis of the remaining chapters follows.

Topics in Data Communications and Computer Networking

This book focuses on a specific aspect of data communications and computer networking. In order to provide context, and to make the book as self-contained as possible, Chapter 2 provides a basic overview of the entire field. The chapter begins with a look at some data communications concepts, including techniques for encoding analog and digital data for both analog and digital signaling, and the concept of multiplexing; the concepts of asynchronous and synchronous transmission are also discussed. The chapter then examines the properties of circuit switching and packet switching. Finally, communications architecture is discussed, using the Open Systems Interconnection (OSI) model as a basis for discussion.

LAN/MAN Technology

The essential technology underlying all forms of local networks comprises topology, transmission medium, and medium access control technique. Chapter 3 provides an overview of the first two of these elements. Four topologies are in common use: star, ring, bus, tree. The most common transmission media for local networking are twisted pair (unshielded and shielded), coaxial cable (baseband and broadband), and optical fiber. These topologies and transmission media are discussed, and the most promising combinations are described. The chapter closes with a discussion of various types of local networks.

Topologies and Transmission Media for LANs and MANs

Chapter 4 is concerned with the topologies and transmission media used in LANs and MANs. The use of twisted pair and coaxial cable in bus/tree LANs is examined first, followed by a discussion of twisted pair for star and ring LANs. The remainder of the chapter examines the increasingly important use of optical fiber in LANs and MANs; star, ring, and bus topologies are covered.

Local Area Network Architecture

Chapter 5 focuses on the protocols needed for stations attached to a LAN to cooperate with each other in the exchange of packets. Specifically, the chapter deals with link control and medium access control protocols. The latter include token-passing and contention-based protocols, such as token ring, token bus, and CSMA/CD. An appendix to Chapter 5 summarizes the standards for LANs that have been issued by the IEEE 802 committee.

Metropolitan Area Network Architecture

Chapter 6 is devoted to a study of the medium access control and physical layer specifications for MANs. The chapter concentrates on the two standards that have been developed: the fiber-distributed data interface (FDDI) and the IEEE 802.6 MAN standard.

Circuit-Switched Local Networks

There is a class of local networks based on the use of circuit switching, including the digital data switch and digital private branch exchange (PBX). Circuit switching is achieved by the use of time-division switching techniques. Chapter 7 begins with an overview of time-division switching techniques, and then examines their application in digital data switches and digital PBXs.

The Network Interface

A local network is a communications facility that supports a number of attached devices. Each device attaches to the network via a *network interface*. Chapter 8 examines the logic required at this interface. A number of issues, including the use of host-to-front-end protocols, is discussed. The differences in handling terminals and computers are also described.

Network Performance

In a LAN or MAN, the data rate, length, and medium access control technique of the network are the key factors in determining the effective capacity of the network. Chapter 9 examines performance of LANs and MANs, and introduces a key parameter, a , that provides a concise but powerful means of characterizing network performance. The issue of end-to-end performance is also considered. Chapter 10 looks at the rather different issues involved in assessing performance on circuit-switched local networks.

Internetworking

The increasing deployment of local networks has led to an increased need to interconnect local networks with each other and with wide-area networks. Chapter 10 focuses on the two most important devices used in internetworking involving local networks: bridges and routers. In both cases, there are two types of protocols involved: protocols for forwarding packets and protocols for exchanging routing information.

Network Management

The final chapter looks at the important issue of network management. Following a general discussion of network management requirements and systems, the OSI-based standards for network management are introduced. The remainder of the chapter looks at network management functions and services that are specific to LANs and MANs, including a discussion of standards in this area.

1.6

RECOMMENDED READING

[SLON91], [MART89], and [NAUG91] are book-length treatments of LANs. [KESS92] covers MANs. [STAL93a] is a collection of reprints of key articles on LANs and MANs. [STAL90b] and [STAL90c] provide a more detailed discussion of applications of local networks.

1.7

PROBLEMS

- 1.1 A computer network is an interconnected set of computers and other devices (terminals, printers, etc.) that can communicate and cooperate with each other to perform certain applications. A subset of a computer network is a communications network (sometimes called a subnetwork) that provides the necessary functions for transferring data between network devices. List functions and capabilities that should be part of the subnetwork and those that should be part of the computer network outside the subnetwork.
- 1.2 On what grounds should a collection of devices connected by point-to-point links be excluded from the definition of local network?
- 1.3 An alternative to a local network for meeting local requirements for data processing and computer applications is a centralized time-sharing system plus a large number of terminals dispersed throughout the local area. What are the major benefits and pitfalls of this approach compared to a local network?
- 1.4 What are the key factors that determine the response time and throughput performance of a local network? Of a centralized system?
- 1.5 In what ways is the human-machine interface of a local network likely to differ from that of a centralized system for:
 - Application users?
 - System operator/managers?

CHAPTER 2

Topics in Data Communications and Computer Networking

The purpose of this chapter is to make this book self-contained for the reader with little or no background in data communications. For the reader with greater interest, references for further study are supplied at the end of the chapter.

2.1

DATA COMMUNICATIONS CONCEPTS

Analog and Digital Data Communications

The terms *analog* and *digital* correspond, roughly, to continuous and discrete, respectively. These two terms are used frequently in data communications in at least three contexts:

- Data
- Signaling
- Transmission

Very briefly, we define *data* as entities that convey meaning. A useful distinction is that data have to do with the form of something; *information* has to do with the content or interpretation of those data. *Signals* are electric or electromagnetic encoding of data. *Signaling* is the act of propagating the signal along some suitable medium. Finally, *transmission* is the communication of data by the propagation and processing of

signals. In what follows, we try to make these abstract concepts clear by discussing the terms *analog* and *digital* in these three contexts.

The concepts of analog and digital data are simple enough. *Analog data* take on continuous values on some interval. For example, voice and video are continuously varying patterns of intensity. Most data collected by sensors, such as temperature and pressure, are continuous-valued. *Digital data* take on discrete values; examples are text and integers.

In a communications system, data are propagated from one point to another by means of electric signals. An *analog signal* is a continuously varying electromagnetic wave that may be transmitted over a variety of media, depending on frequency; examples are wire media, such as twisted pair and coaxial cable, fiber optic cable, and atmosphere or space propagation. A *digital signal* is a sequence of voltage pulses that may be transmitted over a wire medium; for example, a constant positive voltage level may represent binary 1 and a constant negative voltage level may represent binary 0.

The principal advantages of digital signaling are that it is generally cheaper than analog signaling and is less susceptible to noise interference. The principal disadvantage is that digital signals suffer more from attenuation than do analog signals. Figure 2.1 shows a sequence of voltage pulses, generated by a source using two voltage levels, and the received voltage some distance down a conducting medium. Because of the attenuation or reduction of signal strength at higher frequencies, the pulses become rounded and smaller. It should be clear that this attenuation can rather quickly lead to the loss of the information contained in the propagated signal.

Both analog and digital data can be represented, and hence propagated, by either analog or digital signals. This is illustrated in Figure 2.2. Generally, analog data are a function of time and occupy a limited frequency spectrum. Such data can be directly represented by an electromagnetic signal occupying the same spectrum. The best example of this is voice data. As sound waves, voice data have frequency components in the range 20 Hz to 20 kHz. However, most of the speech energy is in a much narrower range. The standard spectrum of voice signals is 300 to 3400 Hz, and this is quite adequate to propagate speech intelligibly and clearly. The telephone instrument does just that. For all sound input



FIGURE 2.1 Attenuation of Digital Signals

Analog Signals – Represent data with continuously varying electromagnetic wave

Analog Data



Digital Data



Digital Signals – Represent data with sequence of voltage pulses

Analog Data



Digital Data

Direct representation or coded

FIGURE 2.2 Analog and Digital Signaling for Analog and Digital Data

in the range of 300 to 3400 Hz, an electromagnetic signal with the same frequency-amplitude pattern is produced. The process is performed in reverse to convert the electromagnetic energy back into sound.

Digital data can also be represented by analog signals by use of a *modem* (modulator/demodulator). The modem converts a series of binary (two-valued) voltage pulses into an analog signal by modulating a *carrier frequency*. The resulting signal occupies a certain spectrum of frequency centered about the carrier and may be propagated across a medium suitable for that carrier. The most common modems represent digital data in the voice spectrum and hence allow those data to be propagated over ordinary voice-grade telephone lines. At the other end of the line, a modem demodulates the signal to recover the original data. Various modulation techniques are discussed below.

In an operation very similar to that performed by a modem, analog data can be represented by digital signals. The device that performs this function for voice data is a *codec* (coder-decoder). In essence, the codec takes an analog signal that directly represents the voice data and approximates that signal by a bit stream. At the other end of a line, the bit stream is used to reconstruct the analog data.

Finally, digital data can be represented directly, in binary form, by two voltage levels. To improve propagation characteristics, however, the binary data are often encoded, as explained below.

TABLE 2.1 Analog and Digital Transmission

(a) Treatment of Signals		
	Analog Transmission	Digital Transmission
Analog Signal	Is propagated through amplifiers; same treatment for both analog and digital data	Assumes digital data; at propagation points, data in signal are recovered and new analog signal is generated
Digital Signal	Not used	Repeaters retransmit new signal; same treatment for both analog and digital data
(b) Possible Combinations		
	Analog Transmission	Digital Transmission
Analog Signal	Analog signal	Digital signal
Digital Signal	Analog signal	Digital signal Analog signal

A final distinction remains to be made. Both analog and digital signals may be transmitted on suitable transmission media. The way these signals are treated is a function of the transmission system. Table 2.1 summarizes the methods of data transmission. Analog transmission is a means of transmitting analog signals without regard to their content; the signals may represent analog data (e.g., voice) or digital data (e.g., data that pass through a modem). In either case, the analog signal will attenuate after a certain distance. To achieve longer distances, the analog transmission system includes amplifiers that boost the energy in the signal. Unfortunately, the amplifier also boosts the noise components. With amplifiers cascaded to achieve long distances, the signal becomes more and more distorted. For analog data, such as voice, quite a bit of distortion can be tolerated and the data remain intelligible. However, for digital data, cascaded amplifiers will introduce errors.

Digital transmission, in contrast, is concerned with the content of the signal. We have mentioned that a digital signal can be transmitted only a limited distance before attenuation endangers the integrity of the data. To achieve greater distances, repeaters are used. A repeater receives the digital signal, recovers the pattern of 1's and 0's, and retransmits a new signal. Thus the attenuation is overcome.

The same technique may be used with an analog signal if it is assumed that the signal carries digital data. At appropriately spaced points, the transmission system has retransmission devices rather than amplifiers. The retransmission device recovers the digital data from the analog signal and generates a new, clean analog signal. Thus noise is not cumulative.

For long-haul communications, digital signaling is not as versatile and practical as analog signaling. For example, digital signaling is im-

possible for satellite and microwave systems. However, digital transmission is superior to analog, both in terms of cost and quality, and wide-area communications systems are gradually converting to digital transmission for both voice and digital data.

We will see that in local networks the trade-offs do not always lead to the same solutions as for wide-area communications. It is still true, within the local context, that digital techniques tend to be cheaper because of the declining cost of digital circuitry. However, the limited distances of local networks limit the severity of the noise and attenuation problems, and the cost and quality of analog techniques approach those of digital. Consequently, there is a secure place for analog signaling and analog transmission in local networks.

Data Encoding Techniques

As we have pointed out, data, either analog or digital, must be converted into a signal for purposes of transmission.

In the case of **digital data**, different signal elements are used to represent binary 1 and binary 0. The mapping from binary digits to signal elements is the *encoding scheme* used for transmission. To understand the significance of the encoding scheme, consider that there are two important tasks in interpreting signals (analog or digital) that carry digital data at the receiver. First, the receiver must know when a bit begins and ends, so that the receiver may sample the incoming signal once per bit time. Second, the receiver must recognize the value of each bit. A number of factors determine how successful the receiver will be in interpreting the incoming signal. For example, the greater the strength of the signal, the more it will withstand attenuation and the more it will stand out from any noise that is present. Also, the higher the data rate, the more difficult the receiver's task is, since each bit occupies a smaller amount of time: the receiver must be more careful about sampling properly and will have less time to make decisions. Finally, the encoding scheme will affect receiver performance. We will describe a number of different encoding techniques for converting digital data to both analog and digital signals.

In the case of **analog data**, the encoding scheme will also affect transmission performance. In this case, we are concerned about the quality, or fidelity, of the transmission. That is, we would like the received data to be as close as possible to the transmitted data. For the purposes of this text, we are concerned about the encoding of analog data in digital form, and techniques for this encoding are presented below.

Digital Data, Analog Signals. The basis for analog signaling is a continuous constant-frequency signal known as the *carrier signal*. Digital data are encoded by modulating one of the three characteristics of the

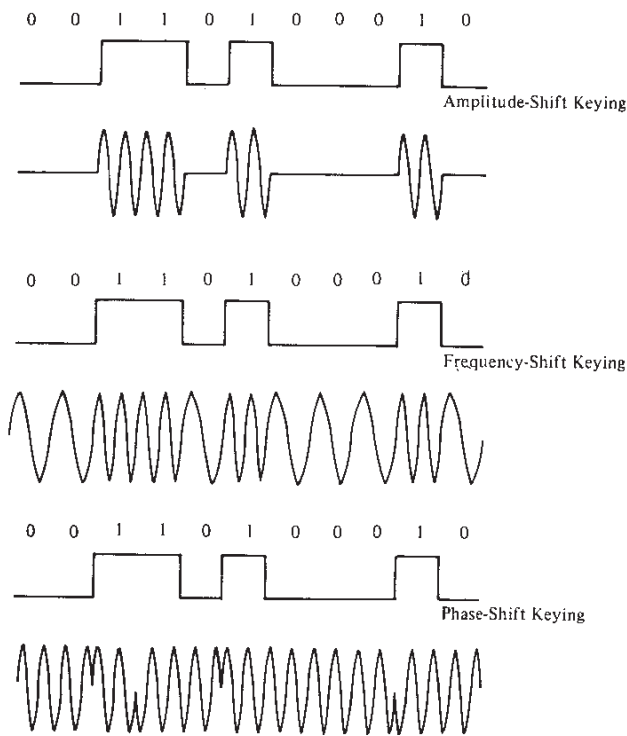


FIGURE 2.3 Modulation of Analog Signals for Digital Data

carrier: amplitude, frequency, or phase, or some combination of these. Figure 2.3 illustrates the three basic forms of modulation of analog signals for digital data:

- Amplitude-shift keying (ASK)
- Frequency-shift keying (FSK)
- Phase-shift keying (PSK)

In all these cases, the resulting signal contains a range of frequencies on both sides of the carrier frequency. That range is referred to as the *bandwidth* of the signal.

In ASK, the two binary values are represented by two different amplitudes of the carrier frequency. In some cases, one of the amplitudes is zero; that is, one binary digit is represented by the presence, at constant amplitude, of the carrier, and the other is represented by the absence of the carrier. ASK is susceptible to sudden gain changes and is a rather inefficient modulation technique. On voice-grade lines, it is typically used up to only 1200 bps.

In FSK, the two binary values are represented by two different frequencies near the carrier frequency. This scheme is less susceptible to error than ASK. On voice-grade lines, it is typically used up to 1200 bps. It is also commonly used for high-frequency (3 to 30 MHz) radio transmission. It can also be used at even higher frequencies on local networks that use coaxial cable.

Figure 2.4 shows an example of the use of FSK for full-duplex operation over a voice-grade line. *Full duplex* means that data can be transmitted in both directions at the same time. To accomplish this, one bandwidth is used for sending, another for receiving. The figure is a specification for the Bell System 108 series modems. In one direction (transmit or receive), the modem passes frequencies in the range 300 to 1700 Hz. The two frequencies used to represent 1 and 0 are centered on 1170 Hz, with a shift of 100 Hz on either side. Similarly, for the other direction (receive or transmit) the modem passes 1700 to 3000 Hz and uses a center frequency of 2125 Hz. The shaded area around each pair of frequencies indicates the actual bandwidth of each signal. Note that there is little overlap and thus little interference.

In PSK, the phase of the carrier signal is shifted to represent data. Figure 2.3 shows an example of a two-phase system. In this system, a 0 is represented by sending a signal burst of the same phase as the previous signal burst sent. A 1 is represented by sending a signal burst of opposite phase to the previous one. PSK can use more than two phase shifts. A four-phase system would encode 2 bits with each signal burst. The PSK technique is more noise resistant and efficient than FSK; on a voice-grade line, rates up to 9600 bps are achieved.

Finally, the techniques discussed above may be combined. A common combination is PSK and ASK, where some or all of the phase shifts may occur at one of two amplitudes.

Digital Data, Digital Signals. Although a common means of transmitting digital data is to pass them through a modem and transmit them as an analog signal, we will see that the transmission of digital data as




FIGURE 2.4 Full-Duplex FSK Transmission on a Voice-Grade Line

digital signals is the technique used in a number of local networks. The use of digital signals may be less expensive and, under some circumstances, provide better performance than analog signaling. In this subsection, we consider two families of coding techniques: NRZ codes and biphase codes.

With **Nonreturn-to-Zero (NRZ) codes**, two different voltage levels, one positive and one negative, are used as the signal elements for the two binary digits. The name refers to the fact that the voltage level never returns to zero, but is always positive or negative. NRZ is the most common and easiest way to transmit digital signals. However, we shall see that its use is not appropriate for local networks.

Figure 2.5a shows the use of a constant negative voltage to represent binary 1 and a constant positive voltage to represent binary 0. This code is known as **NRZ-L (NRZ-level)**. This code is often used for very short connections, such as between a terminal and a modem or a terminal and a nearby computer.

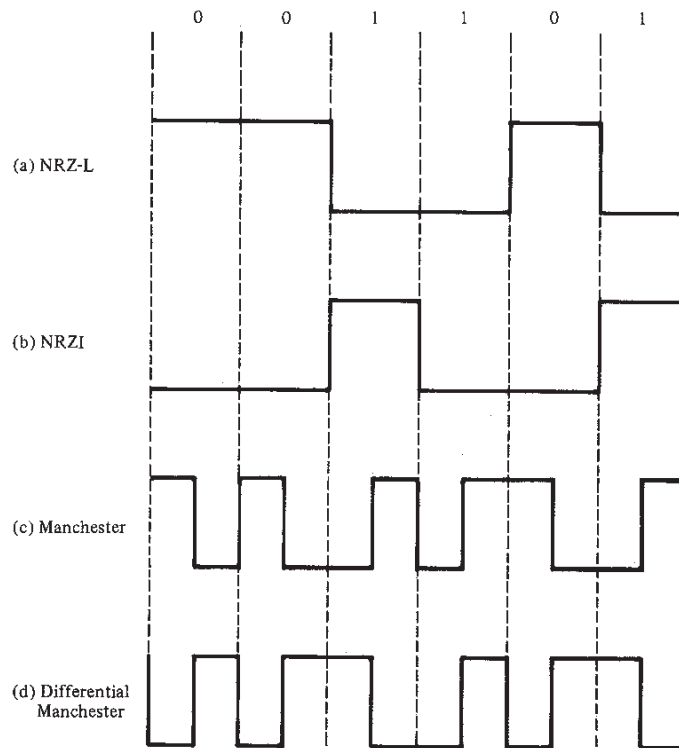


FIGURE 2.5 Digital Signal Encoding

A variation on NRZ is **NRZI** (NRZ, invert on ones). As with NRZ-L, NRZI maintains a constant voltage pulse for the duration of a bit time. The data themselves are encoded as the presence or absence of a signal transition at the beginning of the bit time. A transition (low-to-high or high-to-low) at the beginning of a bit time denotes a binary 1 for that bit time; no transition indicates a binary 0 (Figure 2-5b).

NRZI is an example of differential encoding. In differential encoding, the signal is decoded by comparing the polarity of adjacent signal elements rather than determining the absolute value of a signal element. One benefit of this scheme is that it may be more reliable to detect a transition in the presence of noise than to compare a value to a threshold. Another benefit is that with a complex transmission layout, it is easy to lose the sense of the polarity of the signal. For example, on a twisted-pair medium, if the leads from an attached device to the twisted pair are accidentally inverted, all 1's and 0's will be inverted. This cannot happen with differential encoding.

There are several disadvantages to NRZ transmission. It is difficult to determine where one bit ends and another begins. To picture the problem, consider that with a long string of 1's or 0's for NRZ-L, the output is a constant voltage over a long period of time. Under these circumstances, any drift between the timing of transmitter and receiver will result in the loss of synchronization between the two. Also, there is a direct-current (dc) component during each bit time that may accumulate if positive or negative pulses predominate. Thus, alternating-current (ac) coupling, which uses a transformer and provides excellent electrical isolation between data communicating devices and their environment, is not possible. Furthermore, the dc component can cause plating or other deterioration at attachment contacts.

There is a set of alternative coding techniques, grouped under the term **biphase codes**, which overcomes these problems. Two of these techniques, Manchester and Differential Manchester, are in common use for local networks. All of the biphase techniques require at least one transition per bit time and may have as many as two transitions. Thus, the maximum modulation rate is twice that for NRZ; this means that the bandwidth or transmission capacity required is correspondingly greater. To compensate for this, the biphase schemes have several advantages:

- *Synchronization*: Because there is a predictable transition during each bit time, the receiver can synchronize on that transition. For this reason, the biphase codes are known as self-clocking codes.
- *No dc component*: Because of the transition in each bit time, biphase codes have no dc component, yielding the benefits just described.
- *Error detection*: The absence of an expected transition can be used to detect errors. Noise on the line would have to invert both the

signal before and after the expected transition to cause an undetected error.

In the **Manchester** code (Figure 2.5c), there is a transition at the middle of each bit period. The mid-bit transition serves as a clock and also as data: a low-to-high transition represents a 1, and a high-to-low transition represents a 0. In **Differential Manchester** (Figure 2.5d), the mid-bit transition is used only to provide clocking. The encoding of a 0 is represented by the presence of a transition at the beginning of a bit period, and a 1 is represented by the absence of a transition at the beginning of a bit period. Differential Manchester exhibits the further advantage of being a differential encoding technique.

Analog Data, Digital Signals. The most common example of the use of digital signals to encode analog data is *pulse code modulation* (PCM), which is used to encode voice signals. This section describes PCM and then looks briefly at a similar, less used scheme, *delta modulation* (DM).


PCM is based on the sampling theorem, which states [JORD85]:

If a signal $f(t)$ is sampled at regular intervals of time and at a rate higher than twice the highest significant signal frequency, then the samples contain all the information of the original signal. The function $f(t)$ may be reconstructed from these samples by the use of a low-pass filter.

If voice data are limited to frequencies below 4000 Hz, a conservative procedure for intelligibility, then 8000 samples per second would be sufficient to completely characterize the voice signal. Note, however, that these are analog samples. To convert to digital, each of these analog samples must be assigned a binary code. Figure 2.6 shows an example in which each sample is approximated by being "quantized" into one of 16 different levels. Each sample can then be represented by 4 bits. Of course, it is now impossible to recover the original signal exactly. By using a 7-bit sample, which allows 128 quantizing levels, the quality of the recovered voice signal is comparable to that achieved via analog transmission. Note that this implies that a data rate of 8000 samples per second \times 7 bits per sample = 56 kbps is needed for a single voice signal.

Typically, the PCM scheme is refined using a technique known as *nonlinear encoding*, which means, in effect, that the 128 quantization levels are not equally spaced. The problem with equal spacing is that the mean absolute error for each sample is the same, regardless of signal level. Consequently, lower-amplitude values are relatively more distorted. By using a greater number of quantizing steps for signals of low amplitude, and a small number of quantizing steps for signals of large amplitude, a marked reduction in overall signal distortion is achieved.

PCM can, of course, be used for other than voice signals. For example, a color TV signal has a useful bandwidth of 4.6 MHz, and reason-



Digit	Binary equivalent	Pulse-code waveform
0	0000	
1	0001	
2	0010	
3	0011	
4	0100	
5	0101	
6	0110	
7	0111	
8	1000	
9	1001	
10	1010	
11	1011	
12	1100	
13	1101	
14	1110	
15	1111	

(b)
FIGURE 2.6 Pulse Code Modulation

able quality can be achieved with 10-bit samples, for a data rate of 92 Mbps.

With DM, a bit stream is produced by approximating the derivative of an analog signal rather than its amplitude. A 1 is generated if the current sample is greater in amplitude than the immediately preceding sample; a 0 is generated otherwise. For equal data rates, DM is comparable to PCM in terms of signal quality. Note that for equal data rates, DM requires a higher sampling rate: a 56-kbps voice signal is generated from 8000 PCM samples per second but 56,000 DM samples per second. In general, DM systems are less complex and less expensive than com-

parable PCM systems. A discussion of these and other encoding schemes can be found in [CROC83] and [JAYA84].

Multiplexing


In both local and long-haul communications, it is almost always the case that the capacity of the transmission medium exceeds that required for the transmission of a single signal. To make cost-effective use of the transmission system, it is desirable to use the medium efficiently by having it carry multiple signals simultaneously. This is referred to as *multiplexing*, and two techniques are in common use: frequency-division multiplexing (FDM) and time-division multiplexing (TDM).

FDM takes advantage of the fact that the useful bandwidth of the medium exceeds the required bandwidth of a given signal. A number of signals can be carried simultaneously if each signal is modulated onto a different carrier frequency, and the carrier frequencies are sufficiently separated so that the bandwidths of the signals do not overlap. A simple example of FDM is full-duplex FSK transmission (Figure 2.4). A general case of FDM is shown in Figure 2.7a. Six signal sources are fed into a multiplexer that modulates each signal onto a different frequency (f_1, \dots, f_6). Each signal requires a certain bandwidth centered around its carrier frequency, referred to as a *channel*. To prevent interference, the channels are separated by guard bands, which are unused portions of the spectrum.


An example is the multiplexing of voice signals. We mentioned that the useful spectrum for voice is 300 to 3400 Hz. Thus a bandwidth of 4 kHz is adequate to carry the voice signal and provide a guard band. For both North America (Bell System standard) and internationally [Consultative Committee on International Telegraphy and Telephony (CCITT) standard], a standard voice multiplexing scheme is twelve 4-kHz voice channels from 60 to 108 kHz. For higher-capacity links, both Bell and CCITT define larger groupings of 4-kHz channels.

TDM takes advantage of the fact that the achievable bit rate (sometimes, unfortunately, called bandwidth) of the medium exceeds the required data rate of a digital signal. Multiple digital signals can be carried on a single transmission path by interleaving portions of each signal in time. The interleaving can be at the bit level or in blocks of bytes or larger quantities. For example, the multiplexer in Figure 2.7b has six inputs that might each be, say, 9.6 kbps. A single line with a capacity of 57.6 kbps could accommodate all six sources. Analogously to FDM, the sequence of time slots dedicated to a particular source is called a *channel*. One cycle of time slots (one per source) is called a *frame*.

The TDM scheme depicted in Figure 2.7 is also known as *synchronous TDM*, referring to the fact that time slots are preassigned and fixed. Hence the timing of transmission from the various sources is synchro-



(a) Frequency-Division Multiplexing



(b) Time-Division Multiplexing

FIGURE 2.7 Multiplexing

nized. In contrast, asynchronous TDM allows time on the medium to be allocated dynamically. Examples of this will be discussed later. Unless otherwise noted, the term TDM will be used to mean synchronous TDM only.

One example of TDM is the standard scheme used for transmitting PCM voice data, known in Bell parlance as *T1 carrier*. Data are taken from each source, one sample (7 bits) at a time. An eighth bit is added for signaling and supervisory functions. For T1, 24 sources are multiplexed, so there are $8 \times 24 = 192$ bits of data and control signals per frame. One final bit is added for establishing and maintaining synchronization. Thus a frame consists of 193 bits and contains one 7-bit sample per source. Since sources must be sampled 8000 times per second, the required data rate is $8000 \times 193 = 1.544$ Mbps. As with voice FDM, higher data rates are defined for larger groupings.

TDM is not limited to digital signals. Analog signals can also be interleaved in time. Also, with analog signals, a combination of TDM and

FDM is possible. A transmission system can be frequency-divided into a number of channels, each of which is further divided via TDM. This technique is possible with broadband local networks, discussed in Chapter 4.

Asynchronous and Synchronous Transmission

A fundamental requirement of digital data communication (analog or digital signal) is that the receiver know the starting time and duration of each bit that it receives.

The earliest and simplest scheme for meeting this requirement is asynchronous transmission. In this scheme, data are transmitted one character (of 5 to 8 bits) at a time. Each character is preceded by a start code and followed by a stop code (Figure 2.8a). The *start code* has the encoding for 0 and a duration of 1 bit time; in other words, the start code is 1 bit with a value of 0. The *stop code* has a value of 1, and a minimum duration, depending on the system, of from 1 to 2 bit times. When there are no data to send, the transmitter sends a continuous stop code. The receiver identifies the beginning of a new character by the transition from 1 to 0. The receiver must have a fairly accurate idea of the duration of each bit in order to recover all the bits of the character. However, a small amount of drift (e.g., 1% per bit) will not matter since the receiver resynchronizes with each stop code. This means of communication is simple and cheap but requires an overhead of 2 to 3 bits per character. This technique is referred to as *asynchronous* because characters are sent independently from each other. Thus characters may be sent at a nonuniform rate.

A more efficient means of communication is synchronous transmission. In this mode, blocks of characters or bits are transmitted without start and stop codes, and the exact departure or arrival time of each bit is predictable. To prevent timing drift between transmitter and receiver,

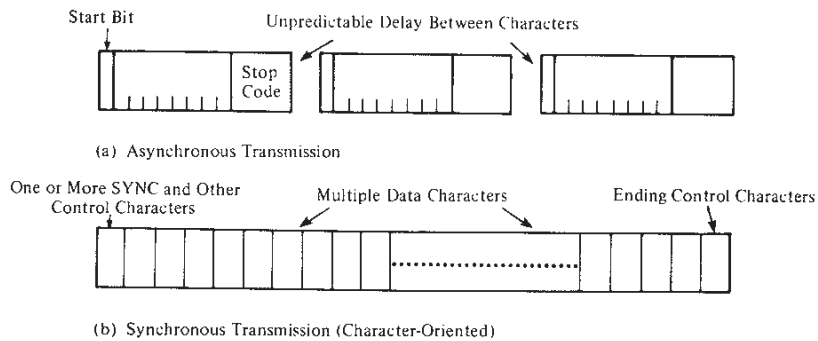


FIGURE 2.8 Asynchronous and Synchronous Transmission

their clocks must somehow be synchronized. One possibility is to provide a separate clock line between transmitter and receiver. Otherwise, the clocking information must be embedded in the data signal. For digital signals, this can be achieved with biphase encoding. For analog signals, a number of techniques can be used; the carrier frequency itself can be used to synchronize the receiver based on the phase of the carrier.

With synchronous transmission, there is another level of synchronization required, to allow the receiver to determine the beginning and end of a block of data. To achieve this, each block begins with a *preamble* bit pattern and ends with a *postamble* bit pattern. The data plus preamble and postamble are called a *frame*. The nature of the preamble and postamble depends on whether the block of data is character-oriented or bit-oriented.

With *character-oriented* schemes, each block is preceded by one or more synchronization characters (Figure 2.8b). The synchronization character, usually called *SYNC*, is chosen such that its bit pattern is significantly different from any of the regular characters being transmitted. The postamble is another unique character. The receiver thus is alerted to an incoming block of data by the *SYNC* characters and accepts data until the postamble character is seen. The receiver can then look for the next *SYNC* pattern.

Character-oriented schemes, such as IBM's *BISYNC*, are gradually being replaced by more efficient and flexible *bit-oriented schemes*, which treat the block of data as a bit stream rather than a character stream. The preamble-postamble principle is the same, with one difference. Since the data are assumed to be an arbitrary bit pattern, there is no assurance that the preamble or postamble pattern will not appear in the data. This event would destroy the higher-level synchronization.

For example, two common bit-oriented schemes, *HDLC* and *SDLC*, use the pattern 01111110 (called a *flag*) as both preamble and postamble. To avoid the appearance of this pattern in the data stream, the transmitter will always insert an extra 0 bit after each occurrence of five 1's in the data to be transmitted. When the receiver detects a sequence of five 1's, it examines the next bit. If the bit is 0, the receiver deletes it. This procedure is known as *bit stuffing*. *HDLC* is examined in more detail in Section 2.3.

2.2

COMMUNICATION SWITCHING TECHNIQUES

So far we have discussed how data can be encoded and transmitted over a communication link. In its simplest form, data communication takes place between two devices that are directly connected by some form of

transmission medium (many of these media are described in Chapter 3). Often, however, it is impractical for two devices to be directly connected. This is so for one (or both) of the following contingencies:

- The devices are very far apart. It would be inordinately expensive, for example, to string a dedicated link between two devices thousands of miles apart.
- There is a set of devices, each of which may require a link to many of the others at various times. Examples are all of the telephones in the world and all of the terminals and computers owned by a single organization. Except for the case of a very few devices, it is impractical to provide a dedicated wire between each pair of devices.

The solution to this problem is to attach each device to a communication network. Communication is achieved by transmitting data from source to destination through a network of intermediate nodes. These nodes are not concerned with the content of the data; rather, their purpose is to provide a switching facility that will move the data from node to node until they reach their destination. Figure 2.9 illustrates the situation. We have a collection of devices that wish to communicate; we will refer to them generically as *stations*. The stations may be computers, terminals, telephones, or other communicating devices. We also have a collection of devices whose purpose is to provide communications, which we will refer to as *nodes*. The nodes are connected to each other in some fashion by transmission links. Each station attaches to a node. The collection of nodes is referred to as a *communications network*. If the




FIGURE 2.9 Generic Switching Network




FIGURE 2.10 Spectrum of Switching Techniques [PRYC91]

attached devices are computers and terminals, then the collection of nodes plus stations is referred to as a *computer network*.

Figure 2.10 describes a spectrum of switching techniques available to transport information across a network. The two extreme ends of the spectrum represent the two traditional switching techniques: circuit switching and packet switching; the remaining techniques are of more recent vintage. In general, the techniques toward the left end of the line provide transmission with little or no variability and with minimal processing demands on attached stations, while techniques toward the right end provide increased flexibility to handle varying bit rates and unpredictable traffic at the expense of increasing processing complexity.

We begin with a detailed look at the two most common switching techniques, and then examine briefly the more advanced techniques.

Circuit Switching

Communication via circuit switching implies that there is a dedicated communication path between two stations. That path is a connected sequence of links between nodes. On each physical link, a channel is dedicated to the connection. The most common example of circuit switching is the telephone network.

Communication via circuit switching involves three phases, which can be explained with reference to Figure 2.9.

1. *Circuit establishment*: Before any data can be transmitted, an end-to-end (station-to-station) circuit must be established. For example, station A sends a request to node 4 requesting a connection to station E. Typically, the circuit from A to 4 is a dedicated line, so that part of the connection already exists. Node 4 must find the next leg in a route leading to node 6. Based on routing information and measures of availability and perhaps cost, node 4 selects the circuit to node 5, allocates a free channel (using TDM or FDM) on

that circuit, and sends a message requesting connection to E. So far, a dedicated path has been established from A through 4 to 5. Since a number of stations may attach to 4, it must be able to establish internal paths from multiple stations to multiple nodes. How this is done is explained in Chapter 7. The remainder of the process proceeds similarly. Node 5 dedicates a channel to node 6 and internally ties that channel to the channel from node 4. Node 6 completes the connection to E. In completing the connection, a test is made to determine if E is busy or is prepared to accept the connection.

2. *Data transfer*: Signals can now be transmitted from A through the network to E. The data may be digital (e.g., terminal to host) or analog (e.g., voice). The signaling and transmission may each be either digital or analog. In any case, the path is: A-4 circuit, internal switching through 4, 4-5 channel, internal switching through 5, 5-6 channel, internal switching through 6, 6-E circuit. Generally, the connection is full duplex, and data may be transmitted in both directions.
3. *Circuit disconnect*: After some period of data transfer, the connection is terminated, usually by the action of one of the two stations. Signals must be propagated to 4, 5, and 6 to deallocate the dedicated resources.

Note that the connection path is established before data transmission begins. Thus channel capacity must be available and reserved between each pair of nodes in the path, and each node must have internal switching capacity to handle the connection. The switches must have the intelligence to make these allocations and to devise a route through the network.

Circuit switching can be rather inefficient. Channel capacity is dedicated for the duration of a connection, even if no data are being transferred. For a voice connection, utilization may be rather high, but it still does not approach 100%. For a terminal-to-computer connection, the capacity may be idle during most of the time of the connection. In terms of performance, there is a delay prior to data transfer for call establishment. However, once the circuit is established, the network is effectively transparent to the users. Data are transmitted at a fixed data rate with no delay other than the propagation delay through the transmission links. The delay at each node is negligible.

Packet Switching

Long-haul circuit-switching telecommunications networks were originally designed to handle voice traffic, and the majority of traffic on these networks continues to be voice. A key characteristic of circuit-switching

networks is that resources within the network are dedicated to a particular call. For voice connections, the resulting circuit will enjoy a high percentage of utilization since, most of the time, one party or the other is talking. However, as the circuit-switching network began to be used increasingly for data connections, two shortcomings became apparent:

1. In a typical terminal-to-host data connection, much of the time the line is idle. Thus, with data connections, a circuit-switching approach is inefficient.
2. In a circuit-switching network, the connection provides for transmission at constant data rate. Thus each of the two devices that are connected must transmit and receive at the same data rate as the other, which limits the utility of the network in interconnecting a variety of host computers and terminals.

To understand how packet switching addresses these problems, let us briefly summarize packet-switching operation. Data are transmitted in blocks, called *packets*. A typical upper bound on packet length is 1000 octets (bytes). If a source has a longer message to send, the message is broken up into a series of packets (Figure 2.11). Each packet consists of a portion of the data (or all of the data for a short message) that a station wants to transmit, plus a packet header that contains control information. The control information, at a minimum, includes the information that the network requires in order to be able to route the packet through the network and deliver it to the intended destination. At each node en route, the packet is received, stored briefly, and passed on to the next node.

Figure 2.12 illustrates the basic operation. A transmitting computer or other device sends a message as a sequence of packets (a). Each

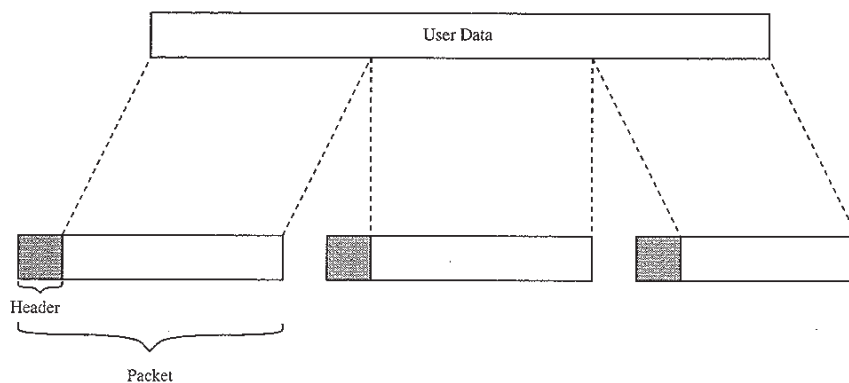


FIGURE 2.11 The Use of Packets




FIGURE 2.12 Packet Switching (Datagram Approach)

packet includes control information indicating the destination station (computer, terminal, etc.). The packets are initially sent to the node to which the sending station attaches. As each packet arrives at this node, it stores the packet briefly, determines the next leg of the route, and queues the packet to go out on that link. Each packet is transmitted to the next node (b) when the link is available. All of the packets eventually work their way through the network and are delivered to the intended destination.

The packet-switching approach has a number of advantages over circuit switching:

1. Line efficiency is greater, since a single node-to-node link can be dynamically shared by many packets over time. The packets are

queued up and transmitted as rapidly as possible over the link. By contrast, with circuit switching, time on a node-to-node link is preallocated using synchronous time-division multiplexing. Much of the time, such a link may be idle because a portion of its time is dedicated to a connection that is idle.

2. A packet-switching network can carry out data-rate conversion. Two stations of different data rates can exchange packets, since each connects to its node at its proper data rate.
3. When traffic becomes heavy on a circuit-switching network, some calls are blocked; that is, the network refuses to accept additional connection requests until the load on the network decreases. On a packet-switching network, packets are still accepted, but delivery delay increases.
4. Priorities can be used. Thus, if a node has a number of packets queued for transmission, it can transmit the higher-priority packets first. These packets will therefore experience less delay than lower-priority packets.

Let us now consider the operation of a packet-switching network. Consider that a station has a message to send through a packet-switching network that is of greater length than the maximum packet size. It therefore breaks up the message into packets and sends these packets, one at a time, to the network. A question arises as to how the network will handle this stream of packets as it attempts to route them through the network and deliver them to the intended destination. There are two approaches that are used in contemporary networks: datagram and virtual circuit.

In the **datagram** approach, each packet is treated independently, with no reference to packets that have gone before. This approach is illustrated in Figure 2.12. Each node chooses the next node on a packet's path, taking into account information received from neighboring nodes on traffic, line failures, and so on. So the packets, each with the same destination address, may not all follow the same route (c), and they may arrive out of sequence at the exit point. In this example, the exit node restores the packets to their original order before delivering them to the destination. In some datagram networks, it is up to the destination rather than the exit node to do the reordering. Also, it is possible for a packet to be destroyed in the network. For example, if a packet-switching node crashes momentarily, all of its queued packets may be lost. Again, it is up to either the exit node or the destination to detect the loss of a packet and to decide how to recover it. In this technique, each packet, treated independently, is referred to as a *datagram*.

In the **virtual circuit** approach, a preplanned route is established before any packets are sent; this route serves to support a logical connection between the end systems. Once the route is established, all of the




FIGURE 2.13 Packet Switching (Virtual Circuit Approach)

packets between a pair of communicating parties follow this same route through the network, as illustrated in Figure 2.13. Because the route is fixed for the duration of the logical connection, it is somewhat similar to a circuit in a circuit-switching network and is referred to as a *virtual circuit*. Each packet now contains a virtual circuit identifier as well as data. Each node on the preestablished route knows where to direct such packets; no routing decisions are required. At any time, each station can have more than one virtual circuit to any other station and can have virtual circuits to more than one station.

So the main characteristic of the virtual circuit technique is that a route between stations is set up prior to data transfer. Note that this setup does not mean that the route is a dedicated path, as in circuit switching. A packet is still buffered at each node and queued for output

over a line. The difference from the datagram approach is that, with virtual circuits, the node need not make a routing decision for each packet. It is made only once for all packets using that virtual circuit.

If two stations wish to exchange data over an extended period of time, there are certain advantages to virtual circuits. First, the network may provide services related to the virtual circuit, including sequencing, error control, and flow control. *Sequencing* is provided since all packets follow the same route; and therefore they arrive in the original order. *Error control* is a service assuring not only that packets arrive in proper sequence, but that all packets arrive correctly. For example, if a packet in a sequence from node 4 to node 6 fails to arrive at node 6, or arrives with an error, node 6 can request a retransmission of that packet from node 4. *Flow control* is a technique for assuring that a sender does not overwhelm a receiver with data. For example, if station E is buffering data from station A and perceives that it is about to run out of buffer space, it can request, via the virtual circuit facility, that station A suspend transmission until further notice. Another advantage is that packets should transit the network more rapidly with a virtual circuit; it is not necessary to make a routing decision for each packet at each node.

One advantage of the datagram approach is that the call setup phase is avoided. Thus, if a station wishes to send only one or a few packets, datagram delivery will be quicker. Another advantage of the datagram service is that because it is more primitive it is more flexible. For example, if congestion develops in one part of the network, incoming datagrams can be routed away from the congestion. With the use of virtual circuits, packets follow a predefined route, and thus it is more difficult for the network to adapt to congestion. A third advantage is that datagram delivery is inherently more reliable. With the use of virtual circuits, if a node fails, all virtual circuits that pass through that node are lost. With datagram delivery, if a node fails, subsequent packets may find an alternate route that bypasses that node.

Table 2.2 summarizes the main features of circuit switching and the two forms of packet switching that we have discussed.

Multirate Circuit Switching

One of the drawbacks of circuit switching is its inflexibility with respect to data rate. If a station attaches to an ordinary circuit-switching network, it is committed to operating at a particular data rate. This data rate must be used regardless of the application, whether it is digitized voice or some data application. Thus, an application with a low data rate requirement would make inefficient use of the network link.

To overcome this inflexibility, an enhanced service, known as multirate circuit switching, was developed. This technique combines circuit switching with multiplexing. The station attaches to the network by

TABLE 2.2 Comparison of Communication Switching Techniques

Circuit Switching	Datagram Packet Switching	Virtual Circuit Packet Switching
Dedicated transmission path	No dedicated path	No dedicated path
Continuous transmission of data	Transmission of packets	Transmission of packets
Fast enough for interactive	Fast enough for interactive	Fast enough for interactive
Messages are not stored	Packets may be stored until delivered	Packets stored until delivered
Path is established for entire conversation	Route established for each packet	Route established for entire conversation
Call setup delay; negligible transmission delay	Packet transmission delay	Call setup delay; packet transmission delay
Busy signal if called party busy	Sender may be notified if packet not delivered	Sender notified of connection denial
Overload may block call setup; no delay for established calls	Overload increases packet delay	Overload may block call setup; increases packet delay
Electromechanical or computerized switching nodes	Small switching nodes	Small switching nodes
User responsible for message-loss protection	Network may be responsible for individual packets	Network may be responsible for packet sequences
Usually no speed or code conversion	Speed and code conversion	Speed and code conversion

TABLE 2.2 (Cont.)

Circuit Switching	Datagram Packet Switching	Virtual Circuit Packet Switching
Fixed bandwidth transmission	Dynamic use of bandwidth	Dynamic use of bandwidth
No overhead bits after call setup	Overhead bits in each packet	Overhead bits in each packet

means of a single physical link. That link is used to carry multiple fixed-data-rate channels between the station and a network node. The traffic on each channel can be switched independently through the network to various destinations.

For this technique, it is possible to develop a scheme in which all of the available channels operate at the same data rate, or a scheme that uses various data rates. For example, integrated services digital network (ISDN) is a standardized digital telecommunications specification. It defines a variety of station-network interfaces, all of which employ multirate circuit switching. The simplest ISDN interface consists of two 64-kbps channels and one 16-kbps channel.

Although this technique is more flexible than simple circuit switching, the same fundamental limitation exists. The user now has the choice of a number of data rates, but each rate remains fixed and the likelihood of inefficient use of a particular channel remains.

Frame Relay

Packet switching was developed at a time when digital long-distance transmission facilities exhibited a relatively high error rate compared to today's facilities. As a result, there is a considerable amount of overhead built into packet-switching schemes to compensate for errors. The overhead includes additional bits added to each packet to enhance redundancy, and additional processing at the end stations and the intermediate network nodes to detect and recover from errors.

With modern, high-speed telecommunications systems, this overhead is unnecessary and counterproductive. It is unnecessary because the rate of errors has been dramatically lowered and any remaining errors can easily be caught by logic in the end systems that operates above the level of the packet-switching logic. It is counterproductive because the overhead involved soaks up a significant fraction of the high capacity provided by the network.

To take advantage of the high data rates and low error rates of contemporary networking facilities, frame relay was developed. Whereas the original packet-switching networks were designed with a data rate to the end user of about 64 kbps, frame relay networks are designed to operate at user data rates of up to 2 Mbps. The key to achieving these high data rates is to strip out most of the overhead involved with error control.

Cell Relay

Cell relay, also known as asynchronous transfer mode, is in a sense a culmination of all of the developments in circuit switching and packet switching over the past 20 years. One useful way to view cell relay is as an evolution from frame relay. The most obvious difference between cell relay and frame relay is that frame relay uses variable-length packets and cell relay uses fixed-length packets, called cells. As with frame relay, cell relay provides minimum overhead for error control, depending on the inherent reliability of the transmission system and on higher layers of logic to catch and correct remaining errors. By using a fixed packet length, the processing overhead is reduced even further for cell relay compared to frame relay. The result is that cell relay is designed to work in the range of 10's and 100's of Mbps, compared to the 2 Mbps of frame relay.

Another way to view cell relay is as an evolution from multirate circuit switching. With multirate circuit switching, only fixed-data-rate channels are available to the end system. Cell relay allows the definition of virtual channels with data rates that are dynamically defined at the time that the virtual channel is created. By using small, fixed-size cells, cell relay is so efficient that it can offer a constant-data-rate channel even though it is using a packet-switching technique. Thus cell relay extends multirate circuit switching to allow multiple channels with the data rate of each channel dynamically set on demand.

2.3

COMPUTER NETWORKING

Communications Architecture

In Chapter 1 we discussed some of the motivations for and benefits of local networking. Many of these factors apply equally well to computer networks in general, whether local or long-haul. Indeed, the move to distributed nonlocal computer networks predates the coming of local networks.

When work is done that involves more than one computer, additional elements are needed: the hardware and software to support the communication between or among the systems. Communications hardware is reasonably standard and generally presents few problems. However, when communication is desired among heterogeneous (different vendors, different models of the same vendor) machines, the software development effort can be a nightmare. Different vendors use different data formats and data exchange conventions. Even within one vendor's product line, different model computers may communicate in unique ways.

As the use of computer communications and computer networking proliferates, a one-at-a-time special-purpose approach to communications software development is too costly to be acceptable. The only alternative is for computer vendors to adopt and implement a common set of conventions. For this to happen, a set of international or at least national standards must be promulgated by appropriate organizations. Such standards would have two effects:

1. Vendors feel encouraged to implement the standards because of an expectation that, because of wide usage of the standards, their products would be less marketable without them.
2. Customers are in a position to require that the standards be implemented by any vendor wishing to propose equipment to them.

It should become clear from the ensuing discussion that no single standard will suffice. The task of communication in a truly cooperative way between applications on different computers is too complex to be handled as a unit. The problem must be decomposed into manageable parts. Hence before one can develop standards, there should be a structure or *architecture* that defines the communications tasks.

This line of reasoning led the International Organization for Standardization (ISO) in 1977 to establish a subcommittee to develop such an architecture. The result was the *Open Systems Interconnection* (OSI) model, which is a framework for defining standards for linking heterogeneous computers. OSI provides the basis for connecting open systems for distributed applications processing. The term *open* denotes the ability of any two systems conforming to the reference model and the associated standards to connect.

Before introducing the OSI model, we consider a simpler architecture that clarifies some of the key concepts involved.

A Three-Layer Model

In very general terms, communications can be said to involve three agents: applications, computers, and networks. The applications that

we are concerned with here are distributed applications that involve the exchange of data between two computer systems. These applications and others execute on computers that can often support multiple simultaneous applications. Computers are connected to networks and the data to be exchanged are transferred by the network from one computer to another. Thus the transfer of data from one application to another involves first getting the data to the computer in which the application resides and then getting them to the intended application within the computer.

With these concepts in mind, it appears natural to organize the communication task into three relatively independent layers:

- Network access layer
- Transport layer
- Application layer

The *network access layer* is concerned with the exchange of data between a computer and the network to which it is attached. The sending computer must provide the network with the address of the destination computer, so that the network may route the data to the appropriate destination. The sending computer may wish to invoke certain services, such as priority, that might be provided by the network. The specific software used at this layer depends on the type of network to be used; different standards have been developed for circuit switching, packet switching, local area networks, and others. Thus it makes sense to separate those functions having to do with network access into a separate layer. By doing this, the remainder of the communications software, above the network access layer, need not be concerned about the specifics of the network to be used. The same higher-layer software should function properly regardless of the particular network to which the computer is attached.

Regardless of the nature of the applications that are exchanging data, there is usually a requirement that data be exchanged reliably. That is, we would like to be assured that all of the data arrive at the destination application and that the data arrive in the same order in which they were sent. The mechanisms for providing reliability are essentially independent of the nature of the applications. Thus it makes sense to collect those mechanisms in a common layer shared by all applications, referred to as the *transport layer*.

Finally, the *application layer* contains the logic needed to support the various user applications. For each different type of application, such as file transfer, a separate module is needed that is particular to that application.

Figures 2.14 and 2.15 illustrate this simple architecture. Figure 2.14 shows three computers connected to a network. Each computer contains software at the network access and transport layers, and software at the

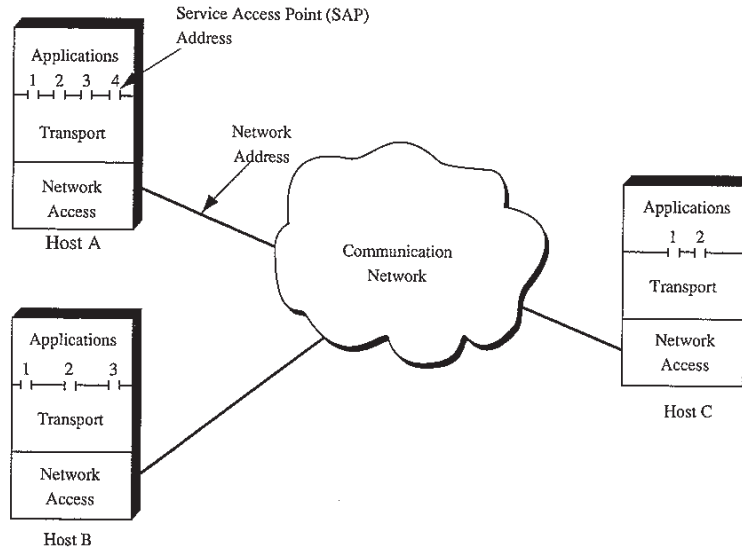


FIGURE 2.14 Communications Architectures and Networks

application layer for one or more applications. For successful communication, every entity in the overall system must have a unique address. Actually, two levels of addressing are needed. Each computer on the network must have a unique network address to allow the network to deliver data to the proper computer. Each application on a computer must have an address that is unique within that computer to allow the transport layer to deliver data to the proper application. These latter addresses are known as *service access points (SAPs)*, connoting the fact

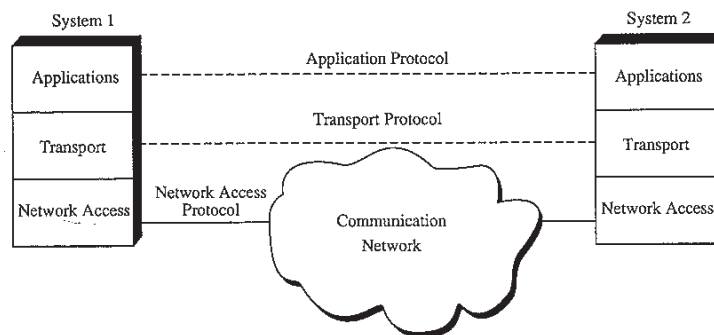


FIGURE 2.15 Protocols in a Simplified Architecture

that each application is individually accessing the services of the transport layer.

Figure 2.15 indicates the way in which modules at the same level on different computers communicate with each other: by means of a protocol. A *protocol* is the set of rules or conventions governing the way in which two entities cooperate to exchange data. A *protocol specification* details the control functions that may be performed, the formats and control codes used to communicate those functions, and the procedures that the two entities must follow.

Let us trace a simple operation. Suppose that an application, associated with SAP 1 at computer A, wishes to send a message to another application, associated with SAP 2 at computer B. The application at computer A hands the message over to its transport layer with instructions to send it to SAP 2 on computer B. The transport layer hands the message over to the network access layer, which instructs the network to send the message to computer B. Note that the network need not be told the identity of the destination service access point. All that it needs to know is that the data are intended for computer B.

To control this operation, control information, as well as user data, must be transmitted, as suggested in Figure 2.16. Let us say that the sending application generates a block of data and passes this to the transport layer. The transport layer may break this block into two smaller pieces to make it more manageable. To each of these pieces the transport layer appends a transport header, containing protocol control information. The combination of data from the next higher layer and control information is known as a *protocol data unit (PDU)*; in this case, it is referred to as a transport protocol data unit. The header in each

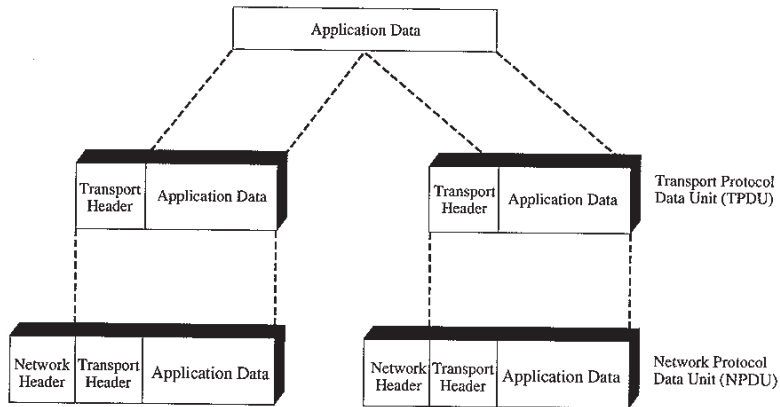


FIGURE 2.16 Protocol Data Units

transport PDU contains control information to be used by the peer transport protocol at computer B. Examples of items that may be stored in this header include:

- *Destination SAP:* When the destination transport layer receives the transport protocol data unit, it must know to whom the data are to be delivered.
- *Sequence number:* Since the transport protocol is sending a sequence of protocol data units, it numbers them sequentially, so that if they arrive out of order, the destination transport entity may reorder them.
- *Error-detection code:* The sending transport entity may calculate and insert an error-detecting code, so that the receiver can determine if an error has occurred and discard the protocol data unit.

The next step is for the transport layer to hand each protocol data unit over to the network layer, with instructions to transmit it to the destination computer. To satisfy this request, the network access protocol must present the data to the network with a request for transmission. As before, this operation requires the use of control information. In this case, the network access protocol appends a network access header to the data it receives from the transport layer, creating a network access PDU. Examples of the items that may be stored in the header include:

- *Destination computer address:* The network must know to whom (which computer on the network) the data are to be delivered.
- *Facilities requests:* The network access protocol might want the network to make use of certain facilities, such as priority.

Figure 2.17 puts all of these concepts together, showing the interaction between modules to transfer one block of data. Let us say that the file transfer module in computer X is transferring a file one record at a time to computer Y. Each record is handed over to the transport layer module. We can picture this action as being in the form of a command or procedure call, A-SEND (application-send). The arguments of this procedure call include the destination computer address, the destination service access point, and the record. The transport layer appends the destination service access point and other control information to the record to create a transport PDU, which is then handed down to the network access layer in a T-SEND command. In this case, the arguments for the command are the destination computer address and the transport protocol data unit. The network access layer uses this information to construct a network PDU. Suppose the network is an X.25 packet-switching network. In this case, the network protocol data unit is an X.25 data packet. The transport protocol data unit is the data field of the

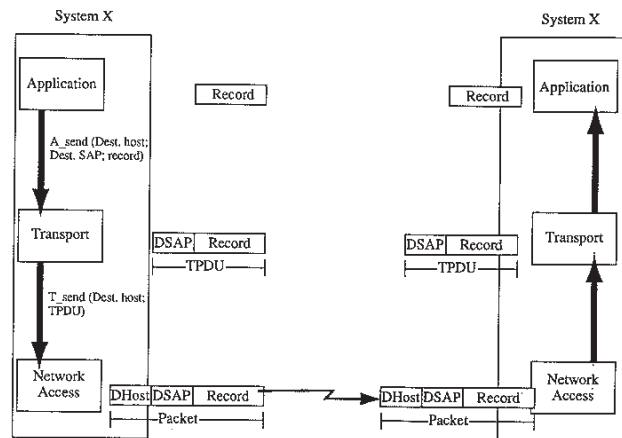


FIGURE 2.17 Operation of a Communication Architecture

packet, and the packet header includes the virtual circuit number for a virtual circuit connecting X and Y.

The network accepts the data packet from X and delivers it to Y. The network access module in Y receives the packet, strips off the packet header, and transfers the enclosed transport protocol data unit to X's transport layer module. The transport layer examines the transport protocol data unit header and, on the basis of the SAP field in the header, delivers the enclosed record to the appropriate application, in this case the file transfer module in Y.

The Concept of Open Systems

Open Systems Interconnection is based on the concept of cooperating distributed applications. In the OSI model, a system consists of a computer, all of its software, and any peripheral devices attached to it, including terminals. A distributed application is an activity that involves the exchange of information between two open systems. Examples of such activities include:

- A user at a terminal on one computer is logged onto an application such as transaction processing on another computer.
- A file management program on one computer transfers a file to a file management program on another computer.
- A user sends an electronic mail message to a user on another computer.
- A process control program sends a control signal to a robot.

OSI is concerned with the exchange of information between open systems and not with the internal functioning of each individual system. Specifically, it is concerned with the capability of systems to cooperate in the exchange of information and in the accomplishment of tasks.

The objective of the OSI effort is to define a set of standards that will enable open systems located anywhere in the world to cooperate by being interconnected through some standardized communications facility and by executing standardized OSI protocols.

An open system may be implemented in any way provided that it conforms to a minimal set of standards allowing communication to be achieved with other open systems. An open system consists of a number of applications, an operating system, and system software such as a data base management system and a terminal handling package. It also includes the communications software that turns a closed system into an open system. Different manufacturers will implement open systems in different ways, in order to achieve a product identity, which will increase their market share or create a new market. However, virtually all manufacturers are now committed to providing communications software that behaves in conformance with OSI in order to provide their customers with the ability to communicate with other open systems.

The OSI Model

A widely accepted structuring technique, and the one chosen by ISO, is layering. The communications functions are partitioned into a hierarchical set of layers. Each layer performs a related subset of the functions required to communicate with another system. It relies on the next lower layer to perform more primitive functions and to conceal the details of those functions. It provides services to the next higher layer. Ideally, the layers should be defined so that changes in one layer do not require changes in the other layers. Thus we have decomposed one problem into a number of more manageable subproblems.

The task of ISO was to define a set of layers and the services performed by each layer. The partitioning should group functions logically, and should have enough layers to make each layer manageably small, but should not have so many layers that the processing overhead imposed by the collection of layers is burdensome. The resulting OSI architecture has seven layers, which are listed with a brief definition in Table 2.3.

Table 2.3 defines, in general terms, the functions that must be performed in a system for it to communicate. Of course, it takes two to communicate, so the same set of layered functions must exist in two systems. Communication is achieved by having the corresponding (peer) layers in two systems communicate. The peer layers communicate

TABLE 2.3 The OSI Layers

Layer	Definition
1. Physical	Concerned with transmission of unstructured bit stream over physical link; involves such parameters as signal voltage swing and bit duration; deals with the mechanical, electrical, and procedural characteristics to establish, maintain, and deactivate the physical link (RS-232-C, RS-449, X.21)
2. Data link	Provides for the reliable transfer of data across the physical link; sends blocks of data (frames) with the necessary synchronization, error control, and flow control (HDLC, SDLC, BiSync)
3. Network	Provides upper layers with independence from the data transmission and switching technologies used to connect systems; responsible for establishing, maintaining, and terminating connections (X.25, layer 3)
4. Transport	Provides reliable, transparent transfer of data between end points; provides end-to-end error recovery and flow control
5. Session	Provides the control structure for communication between applications; establishes, manages, and terminates connections (sessions) between cooperating applications
6. Presentation	Performs generally useful transformations on data to provide a standardized application interface and to provide common communications services; examples: encryption, text compression, reformatting
7. Application	Provides services to the users of the OSI environment; examples: transaction server, file transfer protocol, network management

by means of a set of rules, or conventions, known as a protocol. The key elements of a protocol are:

- *Syntax*: The form in which information is exchanged (format, coding)
- *Semantics*: The interpretation of control information for coordination and error handling
- *Timing*: The sequence in which control events occur

Figure 2.18 illustrates the OSI architecture. Each computer contains the seven layers. Communication is between applications in the two computers, labeled application X and application Y in the figure. If application X wishes to send a message to application Y, it invokes the application layer (layer 7). Layer 7 establishes a peer relationship with layer 7 of the target computer, using a layer 7 protocol (application pro-

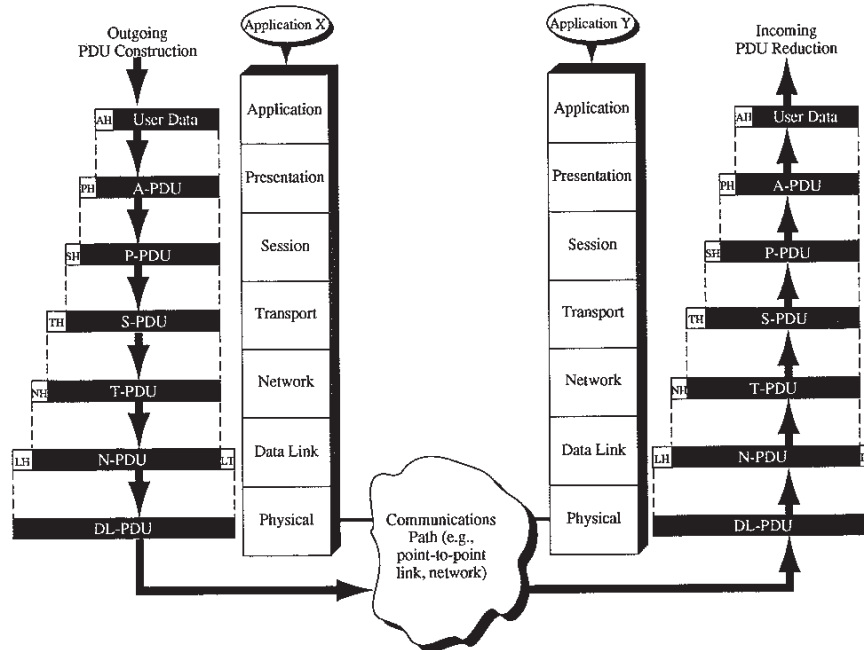


FIGURE 2.18 The OSI Environment

to col). This protocol requires services from layer 6, so the two layer 6 entities use a protocol of their own, and so on down to the physical layer, which actually transmits bits over a transmission medium.

The figure also illustrates the way in which the protocols at each layer are realized. When application X has a message to send to application Y, it transfers those data to an application layer module. That module appends an application header to the data; the header contains the control information needed by the peer layer on the other side. The original data plus the header, referred to as an application protocol data unit (PDU), is passed as a unit to layer 6. The presentation module treats the whole unit as data and appends its own header. This process continues down through layer 2, which generally adds both a header and a trailer. This layer-2 protocol data unit, usually called a *frame*, is then transmitted by the physical layer onto the transmission medium. When the frame is received by the target computer, the reverse process occurs. As we ascend the layers, each layer strips off the outermost header, acts on the protocol information contained therein, and passes the remainder up to the next layer.

We have already seen several examples of the use of control information in headers and trailers. With synchronous communication, a

preamble and postamble are added to each block of data. For packet-switching networks, each packet includes not only data but also (at least) an address.

Note that there is no direct communication between peer layers except at the physical layer. Even at that layer, the OSI model does not stipulate that two systems be directly connected. For example, a packet-switching or circuit-switching network may be used to provide the communications link. This point should become clearer below, when we discuss the network layer.

The attractiveness of the OSI approach is that it promises to solve the heterogeneous computer communications problem. Two systems, no matter how different, can communicate effectively if they have the following in common:

- They implement the same set of communications functions.
- These functions are organized into the same set of layers. Peer layers must provide the same functions, but note that it is not necessary that they provide them in the same way.
- Peer layers must share a common protocol.

To assure the above, standards are needed. Standards must define the functions and services to be provided by a layer (but not how it is to be done—that may differ from system to system). Standards must also define the protocols between peer layers (each protocol must be identical for the two peer layers). The OSI model, by defining a seven-layer architecture, provides a framework for defining these standards.

Protocols

In this section we discuss briefly each of the layers and, where appropriate, give examples of standards for protocols at those layers. Table 2.4 shows the relationship to the OSI model of some of the most important standards. Remember that the OSI layers are not standards; they merely provide a framework for standards.

The Consultative Committee on International Telegraphy and Telephony (CCITT) has developed standards for connecting *data terminal equipment* (DTE) to a packet-switching network that provides *data circuit-terminating equipment* (DCE). These terms correspond to the stations and nodes of Figure 2.9. The standard, X.25, specifically addresses layer 3 and subsumes standards for layers 2 and 1. (Observers are fond of saying that X.25 is an interface, not a protocol. This point is discussed under Network Layer below.) Layer 2 is referred to as LAP-B (Link Access Protocol—Balanced) and is almost identical with ISO's HDLC (High-Level Data Link Control) and ANSI's ADCCP (Advanced Data Communication Control Procedures).

TABLE 2.4 Some Well-Known Layers

OSI	CCITT	ISO	DOD	IEEE 802	ANS X3T9.5
7. Application					
6. Presentation		Various	Various		
5. Session		Session			
4. Transport		Transport (TP)	TCP		
3. Network	X.25	Internet Sublayer	IP		
2. Link	LAP-B			Logical link control Medium access control	Data link Physical
1. Physical	X.21			Physical	

ISO has issued standards for layers 4 and 5 and is in the process of issuing a variety of standards that cover layers 6 and 7. ISO has also developed a sublayer of layer 3 that deals with internetworking, which involves communication across multiple networks.

An internetworking protocol, called IP, has been developed by the Department of Defense (DOD) for its own needs, plus a Transmission Control Protocol (TCP). TCP subsumes all the functions of layer 4 plus some of layer 5. DOD intends to mandate these standards for its procurements. In addition, DOD has issued various standards at the upper layers [STAL86b]. The mismatch with the ISO protocols is, unfortunately, unresolved.

For the type of local network that we refer to as a *local area network* (LAN), the Institute of Electrical and Electronics Engineers (IEEE), through its 802 committee, has developed a three-layer architecture that corresponds to layers 1 and 2 of the OSI model. A number of standards have been developed by the committee for these layers. Similarly, a subcommittee responsible to the American National Standards Institute (ANSI), known as ANS X3T9.5, has developed standards for the type of local network we refer to as a *high-speed local network* (HSLN). These standards, one per layer, correspond nicely to layers 1 and 2 of the OSI model.

This variety may be disheartening, given the alleged benefit of standards, which is to put everyone on the same road. There is certainly room for pessimism. The DOD-ISO disparity makes a uniform federal government position unlikely. For LANs, the 802 committee has produced a number of options and alternatives at each layer.

However, the picture is not as bleak as Table 2.4 makes it seem. With the exception of local networks, which must be treated separately, stan-

dards have settled out quite well for layers 1 through 3. Above that, there is considerable cooperation among the various groups, so that uniform or nearly uniform standards are possible in many cases.

Physical Layer. The *physical layer* covers the physical interface between devices and the rules by which bits are passed from one to another. The physical layer has four important characteristics [BERT80, MCCL83]:

- Mechanical
- Electrical
- Functional
- Procedural

The most common standard in use today is RS-232-C. A typical use of RS-232-C is to connect a digital device to a modem, which in turn connects to a voice-grade telephone line. We will refer to this standard in describing these four characteristics.

The *mechanical characteristics* pertain to the point of demarcation. Typically, this is a pluggable connector. RS-232-C specifies a 25-pin connector, so that up to 25 separate wires are used to connect the two devices.

The *electrical characteristics* have to do with the voltage levels and timing of voltage changes. These characteristics determine the data rates and distances that can be achieved.

Functional characteristics specify the functions that are performed by assigning meaning to various signals. For RS-232-C, and for most other physical layer standards, this is done by specifying the function of each of the pins in the connector. For example, pin CA (Request to Send) is used for the device to signal the modem that it has data to send and that a carrier should be established for modulation. Pin CF (Received Line Signal Detector or Carrier Detect) is used for the modem to alert the device that a carrier is present on the line.

Procedural characteristics specify the sequence of events for transmitting data, based on the functional characteristics. For RS-232-C, the use of the various pins is defined. For example, when a device asserts Request to Send, the modem will assert Clear to Send if it is ready to transmit data. The device can then send data from pin BA (Transmitted Data) over that line to the corresponding pin on the modem.

The physical layer differs from the other OSI layers in that it cannot rely on a lower layer to transmit its PDUs. Rather, it must make use of a transmission medium whose characteristics are not part of the OSI model. There is no physical layer PDU structure as such; no header of protocol control information is used. The PDU simply consists of a block or stream of bits.

Data Link Layer. The data link layer must deal with both the requirements of the communications facility and the requirements of the user.

Whereas the physical layer provides only a raw-bit-stream service, the data link layer attempts to make the physical link reliable and provides the means to activate, maintain, and deactivate the link. The principal service provided by the data link layer to higher layers is that of error detection and control. Thus, with a fully functional data link layer protocol, the next higher layer may assume error-free transmission over the link.

In this subsection we will spend some time defining HDLC, which is a synchronous bit-oriented protocol. We do so for two reasons:

1. HDLC is the ancestor of the link layer protocol standard for LANs (IEEE 802).
2. Many of the concepts concerning protocols are illustrated.

HDLC, and bit-oriented protocols in general, are intended to provide the following capabilities [CARL80]:

- *Code-independent operation (transparency):* The protocol and the data it carries are independent.
- *Adaptability to various applications, configurations, and uses in a consistent manner:* For example, point-to-point, multidrop, and loop configurations should be supported.
- *Both two-way alternate and two-way simultaneous (full-duplex) data transfer.*
- *High efficiency:* The protocol should have a minimum of overhead bits. Also, it should work efficiently over links with long propagation delays and links with high data rates.
- *High reliability:* Data should not be lost, duplicated, or garbled.


With these requirements in mind, we turn to a description of HDLC.

Three modes of operation are defined: The *normal response mode* (NRM), *asynchronous response mode* (ARM), and *asynchronous balanced mode* (ABM). Both NRM and ARM can be used in point-to-point or multipoint configurations. For each there is one *primary station* and one or more *secondary stations*. The primary station is responsible for initializing the link, controlling the flow of data to and from secondary stations, recovering from errors, and logically disconnecting secondary stations. In NRM, a secondary station may transmit only in response to a poll from the primary; in ARM, the secondary may initiate a transmission without a poll. NRM is ideally suited for a multidrop line consisting of a host computer and a number of terminals. ARM may be needed for certain kinds of loop configurations.

ABM is used on point-to-point links only, and each station assumes the role of both primary and secondary. ABM is more efficient for point-to-point lines since there is no polling overhead and both stations may initiate transmissions.

Data are transmitted in frames that consist of six fields (Figure 2.19).

Frame Structure:



Control Field Structure

	1	2	3	4	5	6	7	8
Information	0	N(S)			P/F	N(R)		
Supervisory	1	0	TYPE		P/F	N(R)		
Unnumbered	1	1	TYPE		P/E	MODIFIER		

FIGURE 2.19 The HDLC Frame Structure

- FLAG: Used for synchronization, this field indicates the start and end of a frame. The flag pattern, 01111110, is avoided in the data by bit stuffing.
- ADDRESS: This field identifies the secondary station for this transmission.
- CONTROL: This field identifies the function and purpose of the frame. It is described below.
- DATA: This field contains the data to be transmitted.
- CRC: This is a frame check sequence field. It uses a 16-bit *cyclic redundancy check* (CRC). The CRC field is a function of the contents of the address, control, and data fields. It is generated by the sender and again by the receiver. If the receiver's result differs from the CRC field, a transmission error has occurred (see Appendix 2A).

Three types of frames are used, each with a different control-field format. Information frames carry the data. Supervisory frames provide basic link control functions, and unnumbered frames provide supplemental link control functions.

The P/F (poll/final) bit is used by a primary station to solicit a response. More than one frame may be sent in response, with the P/F bit set to indicate the last frame. The P/F may be used with supervisory and unnumbered frames to force a response.

The N(S) and N(R) fields in the information frame provide an efficient technique for both flow control and error control. A station numbers the frames that it sends sequentially modulo 8, using the N(S) field. When a station receives a valid information frame, it acknowledges that frame with its own information frame by setting the N(R) field to the number of the next frame it expects to receive. This is known as a *piggybacked acknowledgment*, since the acknowledgment rides back on an information frame. Acknowledgments can also be sent on a supervisory frame. This scheme accomplishes three important functions.

1. *Flow control*: Once a station has sent seven frames, it can send no more until the first frame is acknowledged.
2. *Error control*: If a frame is received in error, a station can send a NAK (negative acknowledgment) via a supervisory frame to specify which frame was received in error. This is done in one of two ways. In the *go-back-n protocol*, the sending station retransmits the NAK'ed frame and all subsequent frames that have already been sent. In the *selective repeat technique*, the sending station retransmits only the frame in error.
3. *Pipelining*: More than one frame may be in transit at a time; this allows more efficient use of links with high propagation delay, such as satellite links.

The N(S)/N(R) technique is known as a *sliding-window protocol* because the sending station maintains a window of messages to be sent that gradually moves forward with transmission and acknowledgment. The process is depicted in Figure 2.20.

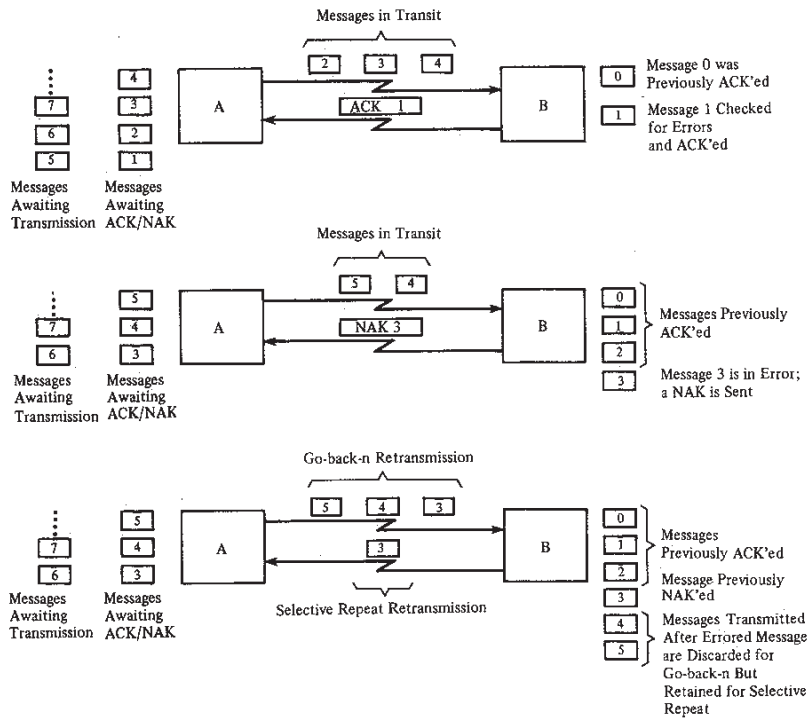


FIGURE 2.20 The Sliding-Window Technique

There are four types of supervisory frames:

1. *Receive Ready (RR)*: used to acknowledge correct receipt of frames up through $N(R)-1$. Alternatively, this is a poll command instructing secondary to begin transmission with sequence number $N(R)$.
2. *Receive Not Ready (RNR)*: used to indicate a temporary busy condition. $N(R)$ is used for a possibly redundant acknowledgment.
3. *Reject (REJ)*: used to indicate an error in frame $N(R)$ and to request retransmission of that and all subsequent frames.
4. *Selective Reject (SREJ)*: used to request retransmission of a single frame.

The unnumbered frames have no sequence number and are used for a number of special purposes, such as to initialize a station, set the mode, disconnect a station, and reject a command.

Network Layer. The network layer provides for the transfer of information between end systems across some sort of communications network. It relieves higher layers of the need to know anything about the underlying data transmission and switching technologies used to connect systems. At this layer, the computer system engages in a dialogue with the network to specify the destination address and to request certain network facilities, such as priority.

There is a spectrum of possibilities for intervening communications facilities to be managed by the network layer. At one extreme, there is a direct point-to-point link between stations. In this case, there may be no need for a network layer because the data link layer can perform the necessary function of managing the link.

Next, the systems could be connected across a single network, such as a circuit-switching or packet-switching network. Figure 2.21 shows how the presence of a network is accommodated by the OSI architecture. The lower three layers are concerned with attaching to and communicating with the network. The packets that are created by the end system pass through one or more network nodes that act as relays between the two end systems. The network nodes implement layers, 1, 2, and 3 of the architecture. In the figure, two end systems are connected through a single network node. Layer 3 in the node performs a switching and routing function. Within the node, there are two data link layers and two physical layers, corresponding to the links to the two end systems. Each data link (and physical) layer operates independently to provide service to the network layer over its respective link.

Note that the layer 1 and 2 protocols are local and they support the exchange of information between an end system and a network node. The upper four layers are end-to-end protocols between the attached end systems. Layer 3 has characteristics of both. The layer 3 protocol is

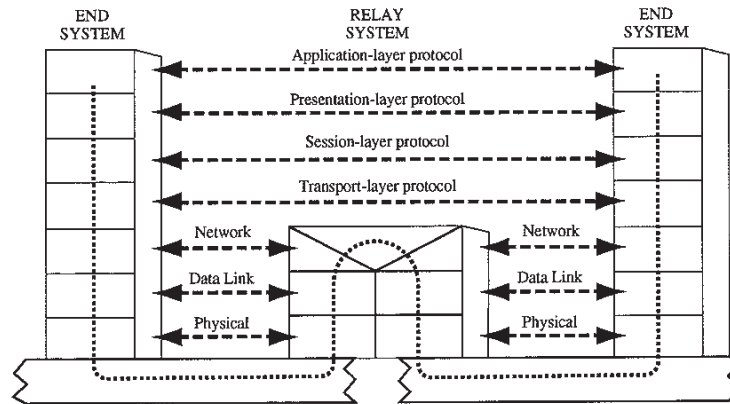


FIGURE 2.21 Communication Across a Network

local in the sense of interfacing to the network and requesting network services. It is end-to-end in the sense that it provides an address for transferring data to the other end system. The X.25 standard is a specification of the lowest three layers of OSI for interfacing an end system to a packet-switching network. A brief description is given in Chapter 8.

At the other extreme, two end systems might wish to communicate but are not even connected to the same network. Rather, they are connected to networks that, directly or indirectly, are connected to each other. This case requires the use of some sort of internetworking technique; we explore this approach in Chapter 10.

Transport Layer. The transport layer provides a reliable mechanism for the exchange of data between computers. It ensures that data are delivered error-free, in sequence, with no losses or duplications. The transport layer may also be concerned with optimizing the use of network services and providing a requested quality of service. For example, the session layer may specify acceptable error rates, maximum delay, priority, and security features.

The mechanisms used by the transport protocol to provide reliability are very similar to those used by data link control protocols such as HDLC: the use of sequence numbers, error-detecting codes, and retransmission after timeout. The reason for this apparent duplication of effort is that the data link layer deals with only a single, direct link, whereas the transport layer deals with a chain of network nodes and links. Although each link in that chain is reliable because of the use of HDLC, a node along that chain may fail at a critical time. Such a failure will affect data delivery, and it is the transport protocol that addresses this problem.

The size and complexity of a transport protocol depends on the type of service it can get from layer 3. For a reliable layer 3 with a virtual circuit capability, a minimal layer 4 is required. If layer 3 is unreliable and/or supports only datagrams, then the layer 4 protocol should include extensive error detection and recovery. Accordingly, ISO has defined five classes of transport protocol, each oriented toward a different underlying network layer service.

Session Layer. The session layer provides the mechanism for controlling the dialogue between the two end systems. In many cases, there will be little or no need for session-layer services, but for some applications, such services are used. The key services provided by the session layer include:

- *Dialogue discipline:* this can be two-way simultaneous (full-duplex) or two-way alternate (half-duplex).
- *Grouping:* the flow of data can be marked to define groups of data. For example, if a retail store is transmitting sales data to a regional office, the data can be marked to indicate the end of the sales data for each department. This would signal the host computer to finalize running totals for that department and start new running counts for the next department.
- *Recovery:* the session layer can provide a checkpointing mechanism, so that if a failure of some sort occurs between checkpoints, the session entity can retransmit all data since the last checkpoint.

ISO has issued a standard for the session layer that includes as options services such as those described above.

Presentation Layer. The presentation layer defines the format of the data to be exchanged between applications, and offers application programs a set of data transformation services. For example, data compression or data encryption could occur at this level.

Application Layer. The application layer provides a means for application programs to access the OSI environment. This layer contains management functions and generally useful mechanisms to support distributed applications. In addition, general-purpose applications such as file transfer, electronic mail, and terminal access to remote computers are considered to reside at this layer.

Perspective on the OSI Model

Figure 2.22 provides a useful perspective on the OSI architecture. The annotation suggests viewing the seven layers in three parts. The lower three layers contain the logic for a computer to interact with a network.

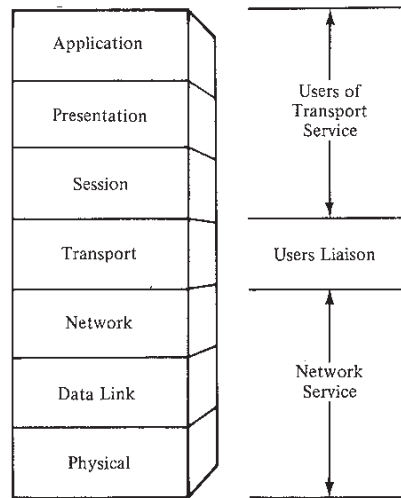


FIGURE 2.22 A Perspective on the OSI Architecture

The host is attached physically to the network, uses a data link protocol to reliably communicate with the network, and uses a network protocol to request data exchange with another device on the network and to request network services. The X.25 standard for packet-switching networks encompasses these three layers. Continuing from this perspective, the transport layer provides a reliable end-to-end service regardless of the intervening network facility; in effect, it is the user's liaison to the communications facility. Finally, the upper three layers, taken together, are involved in the exchange of data between end users, making use of a transport service for reliable data transfer.

2.4

RECOMMENDED READING

[STAL91] covers all of the topics in this chapter. [MART88] also provides a readable but less technical treatment of these topics. A thorough treatment of both analog and digital communications can be found in [COUC90]. Two books by Freeman also provide excellent coverage: [FREE91] concentrates on issues involved with the transmission of data; [FREE89] looks at design issues for communication systems, particularly circuit-switching systems. Another in-depth treatment is offered in the

three-volume [BELL90]. [MCNA88] is a popular and well-respected treatment of the topics in Section 2.1, focusing on digital data communications.

A thorough discussion of the OSI model can be found in [TANE88], which averages about one chapter per layer. [STAL93] covers the standards at each layer of the OSI model, emphasizing the more recent, leading-edge standards. [STAL92] contains reprints of key articles covering OSI and the standards at each layer.

2.5

PROBLEMS

- 2.1 Write a program to do bit stuffing.
- 2.2 A user may wish to use a character-oriented synchronous transmission protocol to send arbitrary bit streams. How can the protocol ensure that none of its control characters (e.g., SYNC) appear in the character stream? Write a program to do this.
- 2.3 Write a program that implements the sliding window technique for (1) selective repeat and (2) go-back-n.
- 2.4 Consider a transmission link between stations A and B with a probability of error in a frame of p .
 - a. Assume a selective repeat protocol and assume that station A is sending data and station B is sending acknowledgments only (RR, SREJ) and that it individually acknowledges each frame. Assume that acknowledgments are never lost. What is the mean number of transmissions required per frame?
 - b. Now assume a go-back-n protocol and that the link is such that A will transmit three additional frames before receiving RR or REJ for each frame. Also assume that acknowledgments are never lost. What is the mean number of transmissions required per frame?
- 2.5 Are the modem and the codec functional inverses (i.e., could an inverted modem function as a codec, and vice versa)?
- 2.6 List the major disadvantages with the layered approach to protocols.
- 2.7 Compare bit-oriented and character-oriented data link protocols in terms of advantages and disadvantages.
- 2.8 Among the principles used by ISO to define the OSI layers were:
 - The number of layers should be small enough to avoid unwieldy design and implementation, but large enough so that separate layers handle functions that are different in process or technology.
 - Layer boundaries should be chosen to minimize the number and size of interactions across boundaries.

Based on these principles, design an architecture with eight layers and make a case for it. Design one with six layers and make a case for that.

- 2.9 Another form of digital encoding of digital data is known as delay modulation or *Miller coding*. In this scheme, a logic 1 is represented by a midbit transition (in either direction). A logic 0 is represented by a transition at the end of the bit period if the next bit is 0, and is represented by the absence of a transition if the next bit is a 1. Draw a Miller code waveform for the bit stream of Figure 2.5. Why might this technique be preferable to NRZ? To Manchester?
- 2.10 What is the percentage of overhead in a T1 carrier (percentage of bits that are not user data)?
- 2.11 Define the following parameters for a switching network:
- N = number of hops between two given stations
 - L = message length, in bits
 - B = data rate, in bps, on all links
 - P = packet size, in bits
 - H = overhead (header) bits per packet
 - S = call setup time (circuit switching or virtual circuit) in seconds
 - D = propagation delay per hop in seconds
- a. For $N = 4$, $L = 3200$, $B = 9600$, $P = 1024$, $H = 16$, $S = 0.2$, $D = 0.001$, compute the end-to-end delay for circuit switching, message switching, virtual circuit packet switching, and datagram packet switching. Assume that there are no acknowledgments.
- b. Derive general expressions for the four techniques, taken two at a time (six expressions in all), showing the conditions under which the delays are equal.
- 2.12 What value of P , as a function of N , B , and H , results in minimum end-to-end delay on a datagram network? Assume that L is much larger than P , and D is zero.
- 2.13 Two stations communicate via a 1-Mbps satellite link. The satellite serves merely to retransmit data received from one station to the other, with negligible delay. The up-and-down propagation delay for a synchronous orbit is 270 ms. Using HDLC frames of length 1024 bits, what is the maximum possible data throughput (not counting overhead bits)?

APPENDIX 2A: THE CYCLIC REDUNDANCY CHECK

In HDLC and other data link control protocols, an error-detection technique is required so that the receiver can detect any bit errors in received

frames and request that the sender retransmit those frames. This technique requires the addition of a **frame check sequence (FCS)**, or **error-detecting code**, to each frame. On transmission, a calculation is performed on the bits of the frame to be transmitted; the result is inserted as an additional field in the frame. On reception, the same calculation is performed on the received bits and the calculated result is compared to the value stored in the incoming frame. If there is a discrepancy, the receiver assumes that an error has occurred.

One of the most common, and one of the most powerful, of the error-detecting codes is the cyclic redundancy check (CRC). For this technique, the message to be transmitted is treated as one long binary number. This number is divided by a unique prime binary number (a number divisible only by itself and 1), and the remainder is attached to the frame to be transmitted. When the frame is received, the receiver performs the same division, using the same divisor, and compares the calculated remainder with the remainder received in the frame. The most commonly used divisors are a 17-bit divisor, which produces a 16-bit remainder, and a 33-bit divisor, which produces a 32-bit remainder.

The measure of effectiveness of any error-detecting code is the percentage of errors it detects. It can be shown that all the following errors are indivisible by a prime divisor and hence are detectable [STAL88a]:

- All single-bit errors
- All double-bit errors, as long as the divisor has at least three 1's
- Any odd number of errors, as long as the divisor contains a factor of 11

TABLE 2.5 Effectiveness of the Cyclic Redundancy Check (CRC)

Type of Error	16-bit CRC	32-bit CRC
	Probability of Detection	Probability of Detection
Single bit errors	1.0	1.0
Two bits in error (separate or not)	1.0	1.0
Odd number of bits in error	1.0	1.0
Error burst of length less than the length of the CRC (16 or 32 bits, respectively)	1.0	1.0
Error burst of length equal to the length of the CRC	$1 - \frac{1}{2^{15}}$	$1 - \frac{1}{2^{31}}$
Error burst of length greater than the length of the CRC	$1 - \frac{1}{2^{16}}$	$1 - \frac{1}{2^{32}}$

- Any burst error for which the length of the burst is less than the length of the divisor polynomial; that is, less than or equal to the length of the FCS
- Most larger burst errors

These results are summarized in Table 2.5. As you can see, this is a very powerful means of error detection and requires very little overhead. As an example, if a 16-bit FCS is used with frames of 1000 bits, then the overhead is only 1.6%. With a 32-bit FCS, the overhead is 3.2%.

CHAPTER 3

Overview of LAN/MAN Technology

The principal technology ingredients that determine the nature of a LAN or MAN are:

- Topology
- Transmission medium
- Medium access control technique

Together, they in large measure determine the type of data that may be transmitted, the speed and efficiency of communications, and even the kinds of applications that a network may support.

This chapter surveys the topologies and transmission media that, within the state of the art, are appropriate for LANs and MANs. The issue of access control is also briefly raised. With this survey as background, three classes of local networks are defined. The discussion is brief, with the objective of providing a context for the material in Chapters 4 through 7.

3.1

TOPOLOGIES

The term *topology*, in the context of a communications network, refers to the way in which the end points or stations of the network are interconnected. A topology is defined by the layout of communications links

and switching elements, and it determines the data paths that may be used between any pair of stations.

To begin the discussion of topology, consider the question of why a communications network is needed at all. According to our definition in Chapter 1, the local network provides a means for interconnecting devices in a small area. Why not provide a direct connection between any pair of devices that need to communicate? Then no intermediate network of communications devices is required.

The problem with this approach is illustrated in Figure 3.1. Each device has a direct, dedicated link, called a *point-to-point link*, with each other device. If there are N devices, then $N(N - 1)$ links are required, and each device requires $(N - 1)$ input/output (I/O) ports. Thus, the cost of the system, in terms of cable installation and I/O hardware, grows with the square of the number of devices.

The infeasibility of this approach, sometimes known as the *mesh topology*, was recognized early for wide-area communications. The solution, as shown in Figure 2.9, was to introduce a network of switching nodes with the ability to route messages, creating logical links and eliminating the need for so many direct physical connections. In this approach, each device or station connects directly to a communication network node and communicates to other stations via the network.

This approach—the use of a collection of switching nodes—is not generally used for local networks. Because the distances involved are small, the expense of the switching nodes can be avoided. Topologies have been developed that require no or only one intermediate switching node and yet avoid the problems of the mesh topology.




FIGURE 3.1 The Problem with Direct Connection or Mesh Topology

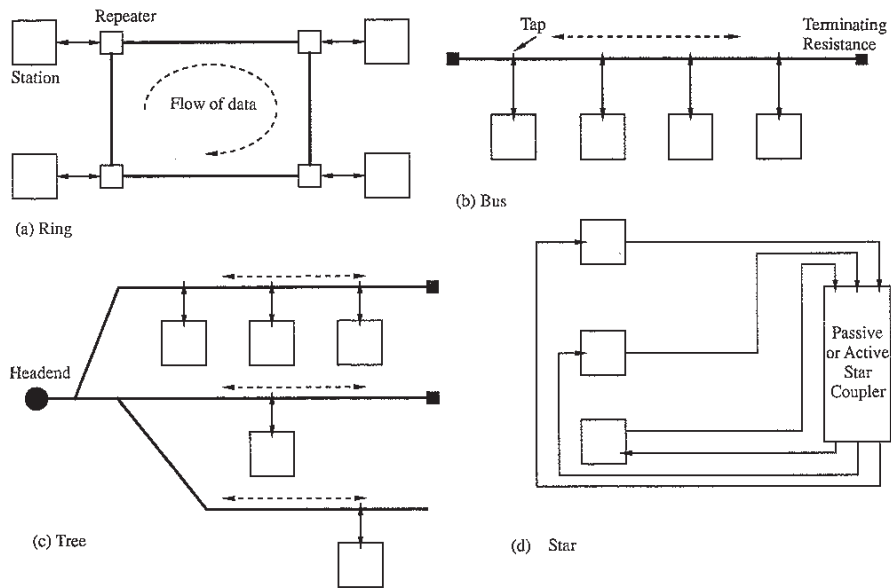


FIGURE 3.2 Local Network Topologies

Four simple topologies are described below: bus, tree, ring, and star (Figure 3.2). These are commonly used, as is, to construct LANs and MANs. They can also be used as building blocks for networks with more complex topologies. These refinements are discussed in later chapters.

Ring Topology

In the *ring topology*, the network consists of a set of *repeaters* joined by point-to-point links in a closed loop. Hence each repeater participates in two links. The repeater is a comparatively simple device, capable of receiving data on one link and transmitting it, bit by bit, on the other link as fast as it is received, with no buffering at the repeater. The links are unidirectional; that is, data are transmitted in one direction only and all are oriented in the same way. Thus data circulate around the ring in one direction (clockwise or counterclockwise).

Each station attaches to the network at a repeater. Data are transmitted in packets. So, for example, if station X wishes to transmit a message to station Y, it breaks the message up into packets. Each packet contains a portion of the data plus some control information, including Y's address. The packets are inserted into the ring one at a time and circulate through the other repeaters. Station Y recognizes its address and copies the packets as they go by.

Since multiple devices share the ring, control is needed to determine at what time each station may insert packets. This is almost always done with some form of distributed control. Each station contains access logic that controls transmission and reception; various techniques are explored in Chapter 5.

Bus and Tree Topologies

With the *bus topology*, the communications network is simply the transmission medium—no switches and no repeaters. All stations attach, through appropriate hardware interfacing, directly to a linear transmission medium, or *bus*. A transmission from any station propagates the length of the medium and can be received by all other stations.

The tree topology is a generalization of the bus topology. The transmission medium is a branching cable with no closed loops. The tree layout begins at a point known as the *headend*. One or more cables start at the headend, and each of these may have branches. The branches in turn may have additional branches to allow quite complex layouts. Again, a transmission from any station propagates throughout the medium and can be received by all other stations. For both bus and tree topologies, the medium is referred to as *multipoint*.

Because all nodes on a bus or tree share a common transmission link, only one station can transmit at a time. Some form of access control is required to determine which station may transmit next. Again, we examine this topic in Chapter 5.

As with the ring, packet transmission is typically used for communication. A station wishing to transmit breaks its message into packets and sends these one at a time. For each packet that a station wishes to transmit, it waits for its next turn and then transmits the packet. The intended destination station will recognize its address as the packets go by and copy them. There are no intermediate nodes and no switching or repeating is involved.

Star Topology

In the star topology, each station is directly connected to a common central switch. One example of the use of this topology is the case in which the central switch uses circuit-switching technology. The digital data switch and digital private branch exchange are examples of this approach.

The star topology is also employed for implementing a packet broadcasting local area network. In this case, each station attaches to a central node, referred to as the *star coupler*, via two point-to-point links, one for transmission in each direction. A transmission from any one station en-

ters the central node and is retransmitted on all of the outgoing links. Thus, although the arrangement is physically a star, it is logically a bus: a transmission from any station is received by all other stations, and only one station at a time may successfully transmit. Thus, the medium access control techniques used for the packet star topology are the same as for bus and tree.

There are two ways of implementing the star coupler. In the case of the **passive star coupler**, there is an electromagnetic linkage in the coupler, so that any incoming transmission is physically passed to all of the outgoing links. In the case of optical fiber, this coupling is achieved by fusing together a number of fibers, so that incoming light is automatically split among all of the outgoing fibers. In the case of coaxial cable or twisted pair, transformer coupling is used to split the incoming signal.

The other type of star coupler is the **active star coupler**. In this case, there is digital logic in the central node that acts as a repeater. As bits arrive on any input line, they are automatically regenerated and repeated on all outgoing lines. If multiple input signals arrive simultaneously, a collision signal is transmitted on all outgoing lines.

Choice of Topology

The choice of topology depends on a variety of factors, including reliability, expandability, and performance. This choice is part of the overall task of designing a local network. As the text proceeds, the trade-offs between the various approaches should become clear. A few general observations follow.

The bus/tree topology appears to be the most flexible one. It is able to handle a wide range of devices, in terms of number of devices, data rates, and data types. High bandwidth is achievable. Because the medium is passive, it would appear at first blush to be highly reliable. As we shall see, this is not necessarily the case. In particular, a break in the cable can disable a large part or all of the network.

Very high-speed links (e.g., optical fiber) can be used between the repeaters of a ring. Hence, the ring has the potential of providing the best throughput of any topology. There are practical limitations, in terms of numbers of devices and variety of data types. Finally, the reliability problem is obvious: a single link or repeater failure could disable the entire network.


The star topology, using circuit switching, readily integrates voice with data traffic. It lends itself well to low-data-rate (≤ 64 kbps) devices. The star topology is good for terminal-intensive requirements because of the minimal processing burden that it imposes on the attached devices.

3.2


TRANSMISSION MEDIA

The *transmission medium* is the physical path between transmitter and receiver in a communications network. Figure 3.3 shows the basic elements of a transmission system. The most common configuration is a point-to-point link between two transmitting/receiving devices, which, through appropriate interfaces, insert analog or digital signals onto the medium. One or more intermediate devices may be used to compensate for attenuation or other transmission impairments. Point-to-point links are used in the ring topology to connect adjacent repeaters, and in the star topology to connect devices to the central switch. Point-to-point links may also be used to connect two local networks in different buildings; we elaborate on this point below. Multipoint links are used to connect multiple devices, as in the bus and tree topologies. Devices attach to the medium at various points; again, repeaters (digital signals) or amplifiers (analog signals) may be used to extend the length of the medium.

Transmission media may be classified as guided or unguided. In both cases, communication is in the form of electromagnetic waves. With *guided media*, the waves are guided along a physical path. Examples of guided media are twisted pair, coaxial cable, and optical fiber, all of



(a) Point-to-point



(b) Multipoint

FIGURE 3.3 Simplified Transmission System Block Diagram

which are used in local networks. The atmosphere and outer space are examples of *unguided media*, which provide a means for transmitting electromagnetic waves but do not guide them. Various forms of transmission through the atmosphere are employed for building-to-building connections.

In this section, we describe these media using the following characteristics:

- *Physical description*: the nature of the transmission medium
- *Transmission characteristics*: include whether analog or digital signaling is used, modulation technique, capacity, and the frequency range over which transmission occurs
- *Connectivity*: point-to-point or multipoint
- *Geographic scope*: the maximum distance between points on the network; whether suitable for intrabuilding, interbuilding, and/or intracity use
- *Noise immunity*: resistance of medium to contamination of the transmitted data
- *Relative cost*: based on cost of components, installation, and maintenance

Twisted Pair

By far the most common transmission medium, for both analog and digital data, is *twisted pair*. The wiring within a building to connect the telephones is twisted pair, as are the local loops that connect all of the phones in a limited geographic area to a central exchange.

Physical Description. A twisted pair consists of two insulated wires arranged in a regular spiral pattern. The wires are copper or steel coated with copper. The copper provides conductivity; steel may be used for strength. A wire pair acts as a single communication link. Typically, a number of these pairs are bundled together into a cable by wrapping them in a tough protective sheath. Over longer distances, cables may contain hundreds of pairs. The twisting of the individual pairs minimizes electromagnetic interference between the pairs. The wires in a pair have thicknesses of from 0.016 to 0.036 inch.

Transmission Characteristics. Wire pairs may be used to transmit both analog and digital signals. For analog signals, amplifiers are required about every 5 to 6 km. For digital signals, repeaters are used every 2 or 3 km.

The most common use of wire pair is for analog transmission of voice. Although frequency components of speech may be found between 20 Hz and 20 kHz, a much narrower bandwidth is required for intelligible

speech reproduction [FREE91]. The standard bandwidth of a full-duplex voice channel is 300 to 3400 Hz. Multiple voice channels can be multiplexed, using FDM, on a single wire pair. A bandwidth of 4 kHz per channel provides adequate separation between channels. Twisted pair has a capacity of up to 24 voice channels using a bandwidth of up to 268 kHz.

Digital data may be transmitted over an analog voice channel using a modem. With a current modem design, speeds of up to 19.2 kbps using phase-shift keying (PSK) are practical. On a 24-channel wire pair, the aggregate data rate is 230 kbps.

It is also possible to use digital or baseband signaling on a wire pair. Bell offers a T1 circuit using twisted pair that handles 24 PCM voice channels, for an aggregate data rate of 1.544 Mbps. Higher data rates, depending on distance, are possible. A data rate of 4 Mbps represents a reasonable upper limit.

Connectivity. Twisted pair can be used for point-to-point and multipoint applications. As a multipoint medium, twisted pair is a less expensive, lower-performance alternative to coax cable but supports fewer stations. Point-to-point usage is far more common.

Geographic Scope. Twisted pair can easily provide point-to-point data transmission to a range of 15 km or more. Twisted pair for local networks is typically used within a single building or just a few buildings.

Noise Immunity. Compared to other guided media, twisted pair is limited in distance, bandwidth, and data rate. The medium is quite susceptible to interference and noise because of its easy coupling with electromagnetic fields. For example, a wire run parallel to an ac power line will pick up 60-Hz energy. Signals on adjacent pairs of cables may interfere with each other, a phenomenon known as *cross-talk*.

Several measures can be taken to reduce impairments. Shielding the wire with metallic braid or sheathing reduces interference. The twisting of the wire reduces low-frequency interference, and the use of different twist lengths in adjacent pairs reduces cross-talk. These measures are effective for wavelengths much greater than the twist length of the cable. Noise immunity can be as high or higher than for coaxial cable for low-frequency transmission. However, above 10 to 100 kHz, coaxial cable is typically superior.

Cost. Twisted pair is less expensive than either coaxial cable or fiber in terms of cost per foot. However, because of its connectivity limitations, installation costs may approach that of other media.

Coaxial Cable

The most versatile transmission medium is *coaxial cable*. In this section we discuss two types of coaxial cable currently in use for LAN applications: 75-ohm cable, which is the standard used in *community antenna television* (CATV) systems, and 50-ohm cable. As Table 3.1 illustrates, 50-ohm cable is used only for digital signaling, called *baseband*; 75-ohm cable is used for analog signaling with FDM, called *broadband*, and for high-speed digital signaling and analog signaling in which no FDM is possible. The latter is sometimes referred to as *single-channel broadband*.

Physical Description. Coaxial cable, like twisted pair, consists of two conductors, but it is constructed differently to permit it to operate over a wider range of frequencies. It consists of a hollow outer cylindrical conductor that surrounds a single inner wire conductor. The inner conductor can be either solid or stranded; the outer conductor can be either solid or braided. The inner conductor is held in place by either regularly spaced insulating rings or a solid dielectric material. The outer conductor is covered with a jacket or shield. A single coaxial cable has a diameter of from 0.4 to about 1 inch.

Transmission Characteristics. The 50-ohm cable is used exclusively for digital transmission. Manchester encoding is typically used. Data rates of up to 10 Mbps can be achieved.

CATV cable is used for both analog and digital signaling. For analog signaling, frequencies up to 300 to 400 MHz are possible. Analog data, such as video and audio, can be handled on CATV cable in much the same way as free-space radio and TV broadcasting. TV channels are each allocated 6 MHz of bandwidth; each radio channel requires much less. Hence a large number of channels can be carried on the cable using FDM.

When FDM is used, the CATV cable is referred to as *broadband*. The frequency spectrum of the cable is divided into channels, each of which carries analog signals. In addition to the analog data referred to above, digital data may also be carried in a channel. Various modulation schemes have been used for digital data, including ASK, FSK, and PSK. The efficiency of the modem will determine the bandwidth needed to support a given data rate. A good rule of thumb [STAH82] is to assume 1 Hz per bps for rates of 5 Mbps and above and 2 Hz per bps for lower rates. For example, a 5-Mbps data rate can be achieved in a 6-MHz TV channel, whereas a 4.8-kbps modem might use about 10 kHz. With current technology, a data rate of about 20 Mbps is achievable; at this rate, the bandwidth efficiency may exceed 1 bps/Hz.

To achieve data rates above 20 Mbps, two approaches have been taken. Both require that the entire bandwidth of the 75-ohm cable be

dedicated to this data transfer; no FDM is employed. One approach is to use digital signaling on the cable, as is done for the 50-ohm cable. A data rate of 50 Mbps has been achieved with this scheme. An alternative is to use a simple PSK system; using a 150-MHz carrier, a data rate of 50 Mbps has also been achieved. Much lower data rates are achieved using FSK.

Connectivity. Coaxial cable is applicable to point-to-point and to multipoint configurations. Baseband 50-ohm cable can support on the order of 100 devices per segment, with larger systems possible by linking segments with repeaters. Broadband 75-ohm cable can support thousands of devices. The use of 75-ohm cable at high data rates (50 Mbps) introduces technical problems that limit the number of devices to 20 to 30.

Geographic Scope. Maximum distances in a typical baseband cable are limited to a few kilometers. Broadband networks can span ranges of tens of kilometers. The difference has to do with the relative signal integrity of analog and digital signals. The types of electromagnetic noise usually encountered in industrial and urban areas are of relatively low frequencies, where most of the energy in digital signals resides. Analog signals may be placed on a carrier of sufficiently high frequency to avoid the main components of noise.

High-speed transmission (50 Mbps), digital or analog, is limited to about 1 km. Because of the high data rate, the physical distance between signals on the bus is very small. Hence very little attenuation or noise can be tolerated before the data are lost.

Noise Immunity. Noise immunity for coaxial cable depends on the application and implementation. In general, it is superior to that of twisted pair for higher frequencies.

Cost. The cost of installed coaxial cable falls between that of twisted pair and optical fiber.

Optical Fiber Cable

One of the most significant technological breakthroughs in information transmission has been the development of practical fiber optic communications systems. Optical fiber already enjoys considerable use in long-distance telecommunications, and its use in military applications is growing. The continuing improvements in performance and decline in prices, together with the inherent advantages of optical fiber, have made it increasingly attractive for local area networking. The following characteristics distinguish optical fiber from twisted pair or coaxial cable:

- *Greater capacity:* The potential bandwidth, and hence data rate, of optical fiber is immense; data rates of 2 Gbps over tens of kilometers have been demonstrated. Compare this to the practical maximum of hundreds of Mbps over about 1 km for coaxial cable and just a few Mbps over 1 km for twisted pair.
- *Smaller size and lighter weight:* Optical fibers are considerably thinner than coaxial cable or bundled twisted-pair cable—at least an order of magnitude thinner for comparable information transmission capacity. For cramped conduits in buildings and underground along public rights of way, the advantage of small size is considerable. The corresponding reduction in weight reduces structural support requirements.
- *Lower attenuation:* Attenuation is significantly lower for optical fiber than for coaxial cable or twisted pair and is constant over a wide range.
- *Electromagnetic isolation:* Optical fiber systems are not affected by external electromagnetic fields. Thus the system is not vulnerable to interference, impulse noise, or cross-talk. By the same token, fibers do not radiate energy, causing little interference with other equipment and providing a high degree of security from eavesdropping. In addition, fiber is inherently difficult to tap.

Physical Description. An optical fiber is a thin (2 to 125 μm), flexible medium capable of conducting an optical ray. Various glasses and plastics can be used to make optical fibers [JORD85]. The lowest losses have been obtained using fibers of ultrapure fused silica. Ultrapure fiber is difficult to manufacture; higher-loss multicomponent glass fibers are more economical and still provide good performance. Plastic fiber is even less costly and can be used for short-haul links, for which moderately high losses are acceptable.

An optical fiber cable has a cylindrical shape and consists of three concentric sections: the core, the cladding, and the jacket. The *core* is the innermost section, and consists of one or more very thin strands, or fibers, made of glass or plastic. Each fiber is surrounded by its own *cladding*, a glass or plastic coating that has optical properties different from those of the core. The outermost layer, surrounding one or a bundle of cladded fibers, is the *jacket*. The jacket is composed of plastic and other materials layered to protect against moisture, abrasion, crushing, and other environmental dangers.

Transmission Characteristics. Optical fiber transmits a signal-encoded beam of light by means of total internal reflection. Total internal reflection can occur in any transparent medium that has a higher index of refraction than the surrounding medium. In effect, the optical fiber acts

as a waveguide for frequencies in the range 10^{14} to 10^{15} Hz, which covers the visible spectrum and part of the infrared spectrum.

Figure 3.4 shows the principle of optical fiber transmission. Light from a source enters the cylindrical glass or plastic core. Rays at shallow angles are reflected and propagated along the fiber; other rays are absorbed by the surrounding material. This form of propagation is called multimode, referring to the variety of angles that will reflect. When the fiber core radius is reduced, fewer angles will reflect. By reducing the radius of the core to the order of a wavelength, only a single angle or mode can pass: the axial ray. This provides superior performance to multimode for the following reason. With multimode transmission, multiple propagation paths exist, each with a different path length and hence time to traverse the fiber. This causes signal elements to spread out in time and limits the rate at which data can be accurately received. Since there is a single transmission path with single-mode transmission, such distortion cannot occur. Finally, by varying the index of refraction of the core, a third type of transmission, known as multimode graded index, is possible. This type is intermediate between the other two in characteristics. The variable refraction has the effect of focusing the rays more efficiently than ordinary multimode, also known as multimode step index. Table 3.1 compares the three fiber transmission modes. As can be seen, tremendous capacities can be achieved, far exceeding those of coaxial cable or twisted pair.

Two different types of light source are used in fiber optic systems: the *light-emitting diode* (LED) and the *injection laser diode* (ILD). The LED is a solid-state device that emits light when a current is applied. The ILD is

TABLE 3.1 Comparison of Three Types of Optical Fibers

	Step-index Multimode	Graded-index Multimode	Single-mode
Light Source	LED or laser	LED or laser	laser
Bandwidth	wide (up to 200 MHz/ km)	very wide (200 MHz to 3 GHz/km)	extremely wide (3 GHz to 50 GHz/km)
Splicing	difficult	difficult	difficult
Typical Application	computer data links	moderate-length telephone lines	telecommunication long lines
Cost	least expensive	more expensive	most expensive
Core Diameter (μm)	50 to 125	50 to 125	2 to 8
Cladding Diameter (μm)	125 to 440	125 to 440	15 to 60

Source: [SHUF84]




FIGURE 3.4 Optical Fiber Transmission Modes

a solid-state device that works on the laser principle in which quantum electronic effects are stimulated to produce a superradiant beam of narrow bandwidth. The LED is less costly, operates over a greater temperature range, and has a longer operational life. The ILD is more efficient and can sustain greater data rates.

The detector used at the receiving end to convert the light into electrical energy is a *photodiode*. Two solid-state devices have been used: the PIN photodiode and the APD detector. The PIN photodiode has a segment of intrinsic (I) silicon between the P and N layers of a diode. The APD, avalanche photodiode, is similar in appearance but uses a stronger elec-

tric field. Both devices are basically photon counters. The PIN is less expensive and less sensitive than the APD.

The amplitude-shift keying technique is commonly used to transmit digital data over optical fiber; in this context, it is known as *intensity modulation*. For LED transmitters, binary one is represented by a short pulse of light and binary zero by the absence of light. Laser transmitters normally have a fixed bias current that causes the device to emit a low light level. This low level represents binary zero, while a higher-amplitude lightwave represents another signal element.

Data rates as high as a few gigabits per second have been demonstrated in the laboratory. Current practical applications are in the range of a few hundreds of megabits per second over a few kilometers.

There is a relationship among the wavelength employed, the type of transmission, and the achievable data rate. Both single mode and multimode can support several different wavelengths of light and can employ laser or LED light sources. In glass-composition fiber, light propagates best in three distinct wavelength "windows," centered on 850, 1300, and 1500 nanometers (nm). The loss is lower at higher wavelengths, allowing greater data rates over longer distances (Table 3.2). Most local applications today use 850-nm LED light sources. Although this is relatively inexpensive, this combination is generally limited to data rates under 100 Mbps and distances of a few kilometers. To achieve higher data rates and longer distances, a 1300-nm LED or laser source is needed. Thus, although the 850-nm source is attractive for LANs, the 1300-nm source is more appropriate for HSLNs. The highest transmission capacities and longest distances achievable today require 1500-nm light sources. These require lasers and are used in some long-distance applications, but are currently too expensive for local networks.

TABLE 3.2 Transmission Losses of Various Types of Optical Fiber [FREE85]

Mode	Material Core/Cladding	Transmission Loss, dB/km		
		850 nm	1300 nm	1500 nm
Single mode	Silica glass/silica glass	2	0.5	0.2
Step-index multimode	Silica glass/silica glass	2	0.5	0.2
	Silica glass/plastic	2.5	High	High
	Multicomponent glass/multicomponent glass	3.4	High	High
Graded-index multimode	Silica glass/silica glass	2	0.5	0.2
	Multicomponent glass/multicomponent glass	3.5	High	High

Currently, a single carrier frequency is used for optical fiber transmission. Future advances will permit practical FDM systems, also referred to as wavelength division multiplexing or color division multiplexing.

Connectivity. The most common use of optical fiber is for point-to-point links. Experimental multipoint systems using a bus topology have been built, but are too expensive to be practical today. In principle, however, a single segment of optical fiber could support many more drops than either twisted pair or coaxial cable, due to lower power loss, lower attenuation characteristics, and greater bandwidth potential.

Geographic Scope. Present technology supports transmission over distances of 6 to 8 km without repeaters. Hence optical fiber is suitable for linking local networks in several buildings via point-to-point links.

Noise Immunity. Optical fiber is not affected by electromagnetic interference or noise. This characteristic permits high data rates over long distance and provides excellent security.

Cost. Fiber optic systems are more expensive than twisted pair and coaxial cable in terms of cost per foot and required components (transmitters, receivers, connectors). While costs of twisted pair and coaxial cable are unlikely to drop, engineering advances should reduce the cost of fiber optics to be competitive with these other media.

Line-of-Sight Media

In this section we look at three techniques for transmitting electromagnetic waves through the atmosphere: microwave, infrared, and laser. All three require a *line-of-sight path* between transmitter and receiver.

Because of the high-frequency ranges at which these devices operate (microwave, 10^9 to 10^{10} Hz; infrared, 10^{11} to 10^{14} Hz; laser, 10^{14} to 10^{15} Hz), there is the potential for very high data rates. Practical systems for short links have been built with data rates of several megabits per second.

These transmission techniques are primarily useful for connecting local networks that are in separate buildings. It is difficult to string cable between buildings, either underground or overhead on poles, especially if the intervening space is public property. The line-of-sight techniques require equipment only at each building.

The *infrared* link consists of a pair of transmitter/receivers (transceivers) that modulate noncoherent infrared light. Transceivers must be within the line of sight, installed on either a rooftop or within a building with data transmitted through adjacent exterior windows. The system

is highly directional; it is extremely difficult to intercept, inject data, or to jam such systems. No licensing is required and the system can be installed in just a few days. Data rates of a few megabits per second over a few kilometers are practical [SEAM82].

A similar system can be installed with *laser* transceivers using coherent light modulation. The major difference is that the Food and Drug Administration (FDA) requires that laser hardware, which emits low-level radiation, be properly shielded. The licensing process takes from 2 to 6 months [CELA82].

Both infrared and laser are susceptible to environmental interference, such as rain and fog. A system with less sensitivity is *microwave*. As with laser and infrared, installation is relatively easy; the major difference is that microwave transceivers can be mounted only externally to a building. Microwave is less directional than either laser or infrared; hence there is a security problem of data eavesdropping, insertion, or jamming. As with all radio-frequency systems, microwave requires Federal Communications Commission (FCC) licensing, which takes about 2 to 3 months. Comparable data rates and distances to laser and infrared can be achieved [RUSH82].

Table 3.3 summarizes the key characteristics of these techniques and includes, for comparison, the use of cable for building-to-building links.

Choice of Transmission Medium

The choice of transmission medium is determined by a number of factors. It is, we shall see, constrained by the topology of the network. Other factors come into play, such as:

- *Capacity*: to support the expected local network traffic
- *Reliability*: to meet availability requirements
- *Types of data supported*: tailored to the application
- *Environmental scope*: to provide service over the range of environments required

And so on. The choice is part of the overall task of designing a network, which is addressed in later chapters. Here we can make a few general observations.

Twisted pair is an inexpensive, well-understood medium. Typically, office buildings are wired to meet the anticipated telephone system demand plus a healthy margin. Compared to coax, the bandwidth is limited. Twisted pair is likely to be the most cost effective for a single-building, low-traffic, local network installation. An office automation system, with a preponderance of dumb terminals and/or intelligent workstations plus a few minis, is a good example.

Coaxial cable is more expensive than twisted pair, but has greater capacity. For a broad range of LAN/MAN requirements, and with the

TABLE 3.3 Transmission Media for Local Networks: Point-to-Point Across Public Property

Medium	Ease of Installation	Regulatory Licensing (months)	Data Rate (Mbps)	Ease of Maintenance	Cost
Infrared	1-2 days, easy	None	1-3	Excellent	Low
Laser	1-2 days, easy	2-6	1-3	Excellent	Low
Microwave	1 week, easy	2-3	1-3	Excellent	Low
Underground coax/optical fiber	1-18 months, moderate to hard	6-18	10+	Fair to good	Moderate to high
Aerial coax/optical fiber	1-6 months, moderate	6-18	10+	Good	Moderate to high

Source: [CELA82].

exception of terminal-intensive systems, it is the medium of choice. For most requirements, a coaxial-based local network can be designed to meet current demand with plenty of room for expansion, at reasonable cost. Coaxial systems excel when there are a lot of devices and a considerable amount of traffic. Examples include large data processing installations and sophisticated office automation systems, which may include facsimile machines, intelligent copiers, and color graphics devices.

At the current state of the art, fiber optic links are suited for point-to-point communications. Hence they do not compete with coaxial cable. The exception is for ring topology networks. However, when the cost of multidrop fiber cable becomes competitive with that of coaxial cable, its advantages—low noise susceptibility, low loss, small size, light weight—will make it a serious contender for many network applications.

The line-of-sight media are not well suited to local network requirements. They are, however, good choices for point-to-point links between buildings, each of which has a twisted-pair or coaxial-based LAN.

3.3

RELATIONSHIP BETWEEN MEDIUM AND TOPOLOGY

Combinations

The choices of transmission medium and topology are not independent. Table 3.4 shows the preferred combinations. The ring topology requires point-to-point links between repeaters. Twisted-pair wire, baseband coaxial cable, and optical fiber can all be used to provide the links. However, broadband coaxial cable would not work well in this topology. Each repeater would have to be capable of receiving and transmitting data simultaneously on multiple channels. It is doubtful that the expense of such devices could be justified. Table 3.5 summarizes representative parameters for transmission media for commercially available ring LANs and MANs. Remember, however, that tables such as this one and Table 3.6 will always be overtaken by new developments in technology.

TABLE 3.4 Relationship Between Medium and Topology

Medium	Topology			
	Bus	Tree	Ring	Star
Twisted pair	X		X	X
Baseband coaxial cable	X		X	
Broadband coaxial cable	X	X		
Optical fiber	X		X	X

TABLE 3.5 Characteristics of Transmission Media for LAN/MAN Ring

Transmission Medium	Data Rate (Mbps)	Repeater Spacing (km)	Number of Repeaters
Twisted Pair	16	0.3	250
Baseband Coaxial Cable	16	1.0	250
Optical Fiber	100	2.0	500

For example, the possibility of standardizing a twisted-pair ring LAN at 100 Mbps is now under study, and a number of manufacturers have developed designs and expressed an interest in introducing such a product [VERE91].

For the bus topology, twisted-pair and both baseband and broadband coaxial cable are appropriate, and numerous products exist for each of these media. Until recently, optical fiber cable has not been considered feasible; the multipoint configuration was not considered cost effective, due to the difficulty in constructing low-loss optical taps. However, recent advances have made the optical fiber bus practical, even at quite high data rates.

The tree topology can be employed with broadband coaxial cable. The unidirectional nature of broadband signaling allows the construction of a tree architecture. On the other hand, the bidirectional nature of baseband signaling on either twisted pair or coaxial cable is not suited to the tree topology. Table 3.6 summarizes representative parameters for transmission media for commercially available bus and tree LANs and MANs.

The reader will note that the performance for a given medium is considerably better for the ring topology compared with the bus/tree topology. In the bus/tree topology, each station is attached to the medium by a tap, and each tap introduces some attenuation and distortion to the

TABLE 3.6 Characteristics of Transmission Media for LAN/MAN: Bus/Tree

Transmission Medium	Data Rate (Mbps)	Range (km)	Number of Taps
Unshielded Twisted Pair	1-2	<2	10's
Baseband Coaxial Cable	10/70	<3/<1	100's/10's
Broadband Coaxial Cable	20 per channel	<30	100's-1,000's
Optical Fiber	45	<150	500

signal as it passes by. In the ring, each station is attached to the medium by a repeater, and each repeater generates a new signal to compensate for effects of attenuation and distortion.

The star topology requires a point-to-point link between each device and the central node. For circuit switching, where the central node performs the circuit-switching task, twisted pair has traditionally been used. The higher data rates of coaxial or fiber would overwhelm the typical circuit-switching node. For packet switching, the performance of the star topology will depend on whether an active or pass star configuration is used.

Layout

One very practical issue is related to the selection of both medium and topology, and that is the actual layout of the transmission medium in the building. To address this issue, we need to make a distinction between topology and geometry. The net illustrations in Figure 3.2 depict the various topologies of local networks; this defines the way in which the devices are interconnected. But, as a practical matter, the actual path that the cable follows is constrained by physical characteristics of the building. The cable must follow routes that accommodate the walls and floors of the building. Typically, predefined cable paths are used, sometimes defined by the existence of conduits. Thus the geometry, or actual layout, of the cable will be distorted to some extent relative to the intended topology.

Let us consider some of the requirements that dictate the layout of the installed cable. Of prime importance is the need to minimize cost while providing the required capacity. One determinant of cost, of course, is the medium itself. As was mentioned, twisted pair is cheaper than coaxial cable, which is in turn cheaper than optical fiber. It is often the case, however, that the installation costs, which are primarily labor costs, far exceed the cost of the materials. This is particularly true in existing buildings, which may present difficulties in finding pathways for new cable. In new buildings, the problems and costs can be minimized if the cable layout for a local network can be designed ahead of time. Then the cable can be installed during construction.

A second important requirement is that the layout be suitable for accommodating equipment relocation and network growth. It is not unusual for 50% of the installed data terminals in an office building to be moved each year [IBM84]. And, with the continued proliferation of personal and other microcomputers, virtually any local network can be expected to grow. The safest way to plan for both relocation and growth is to install a network that reaches every office, or at least to install a smaller network that can easily be expanded to include additional offices with little or no disruption of the existing network. Finally, the layout

TABLE 3.7 The Use of Alternative Wiring Strategies

Medium	Topology		
	Star	Ring	Bus
Twisted pair	S	L, S	L, S
Coaxial cable		L	L
Optical fiber		L	L

L = linear wiring strategy
S = star wiring strategy


should be such as to facilitate servicing and maintenance. When a fault occurs somewhere in the network, we would like to be able to locate the fault, isolate it from the rest of the network, and fix it as soon as possible.

With the above considerations in mind, we can identify two general strategies for laying out the local network transmission medium: linear and star. Table 3.7 summarizes the relationship of the layout strategy to the transmission media and topologies of local networks.


The linear strategy attempts to provide the desired topology with the minimum cable, subject to the physical constraints of the building. The medium is propagated to the subscriber locations, which may be some or all of the offices in the building. Any of the guided media that have been described can be used, and either a bus or ring topology can be provided.

The star layout strategy uses an individual cable from a concentration point to each subscriber location. This is clearly the proper approach for the star topology local network. It can also be used for bus and ring topologies, as depicted in Figure 3.5. In the case of the bus topology, the bus is very short and resides at the concentration point; the drop cables to the attached devices are relatively long. In the case of the ring topology, the ring is distorted so that each link of the ring loops through the concentration point. Typically, this layout is used separately on each floor of a building. The concentration point is referred to as a wiring closet; some or all of the offices on the floor are connected to the closet. Connections between floors are provided by linking the closets. This type of layout is invariably used to support telephones in an office building and is becoming increasingly popular for local networks.

Although the star strategy is logical for the star topology, it may seem inappropriate for the ring and bus topologies. Its main disadvantage is that, for the ring and bus, the star strategy will require more cable than the linear strategy, increasing cost and cable congestion. For this reason, the star strategy is rarely used for coaxial cable or optical fiber local networks. However, the star approach is well suited for twisted-pair local



a. Bus Using Star Wiring



b. Ring Using Star Wiring

FIGURE 3.5 Bus and Ring Topologies Using Star Wiring

networks, where the cost penalty is lower. Some of the advantages of the star strategy are:

1. It lends itself to prewiring of the building. The layout is a regular one and conforms to normal installation practice in office buildings. Furthermore, most existing buildings are prewired with excess unshielded twisted pair. Thus, for local networks that employ unshielded twisted pair, it may be possible to use existing wiring. Even in the case of shielded twisted pair, installation will be easier since the paths for the new cable are well defined.
2. The system can be easily expanded, simply by patching additional cables into the network at the wiring closet.
3. Servicing and maintenance are easier. Diagnosis of problems can be performed from centralized points. Faults can easily be isolated by patching cables out of the network.

Further discussion of the star strategy will be provided as we look at some specific uses in the next two chapters.

3.4

CLASSES OF NETWORKS

There are a number of ways of classifying communications networks. We touched on this topic briefly in Chapter 1 and are now in a position to examine it in more depth.

One way to classify networks is in terms of the technology used: specifically, in terms of topology and transmission medium. This chapter has introduced these basic elements. As we shall see, the same topologies and transmission media are repeated in a wide variety of networks.

Perhaps the most commonly used means of classification is on the basis of geographical scope. Traditionally, networks have been classified as either local area networks (LANs) or wide-area networks (WANs). A category that has recently begun to receive much attention is the metropolitan area network (MAN).

Figure 3.6 illustrates these categories, plus some special cases. Table 3.8 summarizes key performance parameters. We examine each of these in turn.




FIGURE 3.6 Comparison of Multiprocessor Systems, LANs, MANs, and WANs

TABLE 3.8 Characteristics of LANs, MANs, and WANs

Network	Data Rate	Distance Covered
Local Area Network (IEEE 802)	1–20 Mbps	< 25 km
Fiber Distributed Data Interface	100 Mbps	< 200 km
Metropolitan Area Network (IEEE 802.6)	30 Mbps–1+Gps	<160 km
Traditional Wide-Area Network	10 kbps–1.5 Mbps	unlimited
High-Speed Wide-Area Network	50 Mbps–1+Gbps	unlimited

Wide-Area Network

Wide-area networks have traditionally been considered to be those that cover a large geographical area, require the crossing of public right-of-ways, and rely at least in part on circuits provided by a common carrier. Such WANs are typically switched communications networks, consisting of an interconnected set of switching nodes (see Figure 2.9); each station attaches to one of the nodes. As was discussed in Chapter 2, a WAN may be either circuit-switched or packet-switched.

Until recently, WANs have provided only relatively modest capacity to subscribers. For data attachment, either to a packet-switching network or to a circuit-switching network by means of a modem, data rates of 9600 bps or even less have been common. Business subscribers have been able to obtain higher rates, with a service known as T-1, which operates at 1.544 Mbps. The most important recent development in WANs in this range of performance has been the development of the integrated services digital network (ISDN), which provides circuit-switching and packet-switching services at rates up to 1.544 Mbps (2.048 Mbps in Europe). The basic user interface to ISDN is multirate circuit switching using twisted pair.

The continuing development of practical optical fiber facilities has led to the standardization of much higher data rates for WANs, and we can expect these services to become widely available over the next few years, certainly by the end of the 1990s. These high-speed WANs provide user connections in the 10's and 100's of Mbps. The most important effort in this regard is the standardization of a broadband integrated services digital network (B-ISDN) that uses cell relay rather than circuit switching.

WANs are beyond the scope (no pun intended) of this book.

Local Area Network

As with wide-area networks, a local area network is a communications network that interconnects a variety of devices and provides a means for information exchange among those devices. There are several key distinctions between LANs and WANs:

1. The scope of the LAN is small, typically a single building or a cluster of buildings. This difference in geographic scope leads to different technical solutions. In particular, LANs almost invariably employ a shared medium topology (bus, tree, ring, passive, or active star) as opposed to a switching network architecture, and they use a packet broadcasting technique.
2. It is usually the case that the LAN is owned by the same organization that owns the attached devices. For WANs, this is less often the case, or at least a significant fraction of the network assets are not owned. This has two implications. First, care must be taken in the choice of LAN, since there may be a substantial capital investment (compared to dial-up or leased charges for wide-area networks) for both purchase and maintenance. Second, the network management responsibility for a local network falls solely on the user.
3. The internal data rates of LANs are much greater than those of wide-area networks.

LANs have been the focus of a standardization effort by the IEEE 802 committee, and it is perhaps useful to quote their definition of a LAN [IEEE90]:

The LANs described herein are distinguished from other types of data networks in that they are optimized for a moderate size geographic area such as a single office building, a warehouse, or a campus. The IEEE 802 LAN is a shared medium peer-to-peer communications network that broadcasts information for all stations to receive. As a consequence, it does not inherently provide privacy. The LAN enables stations to communicate directly using a common physical medium on a point-to-point basis without any intermediate switching node being required. There is always need for an access sublayer in order to arbitrate the access to the shared medium. The network is generally owned, used, and operated by a single organization. This is in contrast to Wide Area Networks (WANs) that interconnect communication facilities in different parts of a country or are used as a public utility. These LANs are also different from networks, such as backplane buses, that are optimized for the interconnection of devices on a desk top or components within a single piece of equipment.

The committee was given a charter to develop standards for networks in the range of 1 to 20 Mbps.

Metropolitan Area Networks

As the name suggests, a MAN occupies a middle ground between LANs and WANs. Interest in MANs has come about as a result of a recognition that the traditional point-to-point and switched network techniques used in WANs may be inadequate for the growing needs of organizations. While broadband ISDN, with cell relay, holds out promise for meeting a wide range of high-speed needs, there is a requirement now for both private and public networks that provide high capacity at low costs over a large area. The high-speed shared-medium approach of the LAN standards provides a number of benefits that can be realized on a metropolitan scale.

Over many years of research on MANs, a number of alternatives have been explored and rejected. One approach has emerged that has received widespread support from providers and users, and has been standardized by the IEEE 802 committee as IEEE 802.6. Again, it is useful to look at their definition:

A MAN is optimized for a larger geographical area than a LAN, ranging from several blocks of buildings to entire cities. As with local networks, MANs can also depend on communications channels of moderate-to-high data rates. Error rates and delay may be slightly higher than might be obtained on a LAN. A MAN might be owned and operated by a single organization, but usually will be used by many individuals and organizations. MANs might also be owned and operated as public utilities. They will often provide means for internetworking of local networks. Although not a requirement for all LANs, the capability to perform local networking of integrated voice and data (IVD) devices is considered an optional function for a LAN. Likewise, such capabilities in a network covering a metropolitan area are optional functions of a MAN.

Whereas the LANs defined by IEEE 802 are typically used only to support data traffic, the 802.6 MAN is intended for the support of both data and voice traffic. As Figure 3.6 indicates, MANs cover greater distances at higher data rates than LANs, although there is some overlap in geographical coverage.

The primary market for MANs is the customer that has high-capacity needs in a metropolitan area. A MAN is intended to provide the required capacity at lower cost and greater efficiency than obtaining an equivalent service from the local telephone company.

The MAN, as defined in IEEE 802.6, shares many characteristics with LANs. As with LANs, the MAN uses packet broadcasting over a shared transmission medium. As we will see, the MAN defined by 802.6 uses a bus topology.

Fiber Distributed Data Interface

The fiber distributed data interface (FDDI) is a network standard developed by the American National Standards Institute (ANSI) that specifies a 100-Mbps optical fiber ring network. In the literature, FDDI is generally considered to be a LAN and, indeed, most of the existing installations are within a single building or small cluster of buildings. FDDI is designed to provide high end-to-end throughput between expensive, high-speed devices such as mainframes and mass storage devices. It is also used as a backbone network to connect a number of lower-speed LANs.

As with the typical LAN, FDDI was originally defined to use packet broadcasting and to support data traffic. A recent enhancement to the standard, known as FDDI-II, provides support for voice traffic and other applications that normally use circuit switching.

As Table 3.8 and Figure 3.6 indicate, FDDI can be considered either to be a high-speed LAN or a MAN. The classification is somewhat arbitrary. However, for purposes of this book, it is more convenient to present the technical details of FDDI in the chapter on MANs, and that is the course that is followed.

Circuit-Switched Local Networks

In contrast to LANs and MANs, which use packet broadcasting, there is another approach to local networking using circuit switching. Typically, circuit-switching local networks have a star or hierarchical star topology using twisted-pair wire to connect end points to the switch. In the hierarchical star, high-speed trunks of coaxial cable or optical fiber may be used to connect satellite switching units to the central switching unit. Data rates to individual stations are typically low (≤ 64 kbps), but bandwidth is guaranteed and there is essentially no network delay once a connection is made.

One form of circuit-switched local network is the digital private branch exchange (PBX). This is an on-premise switch designed to handle both voice and data connections. Although the strength of these systems is their support from telephones, they are also well suited to terminal-to-host data traffic. Another form of circuit-switched local network is the digital data switch. Devices in this category are designed to handle data only, not voice, and are typically lower in cost than a digital PBX of comparable size.

Traditionally, the term *LAN* has been reserved for the packet-broadcasting, shared-medium networks of relatively high data rates indicated in Figure 3.6. Although a circuit-switched local network certainly provides local area coverage, the differences in architecture and data

rate compared to LANs is such that these networks need to be treated separately.

3.5

RECOMMENDED READING

Detailed descriptions of the transmission characteristics of the transmission media discussed in this chapter can be found in [FREE91] and [BELL90]. A number of books provide detailed coverage of optical fiber transmission. Two that can be recommended are [DIAM90] and [ZANG91]; the former concentrates on the principles of optical transmission, while the latter is concerned with optical communications systems. [STAL93a] contains reprints of many key articles on LAN/MAN technology.

3.6

PROBLEMS

- 3.1 What functions should be performed by the network layer (layer 3) in a bus topology local network? Ring topology? Star topology?
- 3.2 Could HDLC be used as a link layer for a bus topology local network? If not, what is missing? Answer for ring and star.
- 3.3 An asynchronous device, such as a teletype, transmits characters one at a time with unpredictable delays between characters. What problems, if any, do you foresee if such a device is connected to a local network and allowed to transmit at will (subject to gaining access to the medium)? How might such problems be resolved? Answer for ring, bus, and star.
- 3.4 Which combination or combinations of medium and topology would be appropriate for the following applications, and why?
 - a. Terminal intensive: many terminals throughout an office: one or a few shared central computers
 - b. Small network: fewer than 50 devices, all low speed (<56 kbps)
 - c. Office automation: a few hundred devices, mostly terminals and minicomputers
- 3.5 Consider the transfer of a file containing one million characters from one section to another. What is the total elapsed time and effective throughput for the following cases:
 - a. A circuit-switched, star topology local network. Call setup time is negligible, and the data rate on the medium is 64 kbps.
 - b. A bus topology local network with two stations a distance D




FIGURE 3.7 Building Layout for a Local Network

apart, a data rate of B bps, and a packet size P with 80 bits of overhead. Each packet is acknowledged with an 88-bit packet before the text is sent. The propagation speed on the bus is 200 m/μsec. Solve for:

- (1): $D = 1$ km, $B = 1$ Mbps, $P = 256$ bits
 - (2): $D = 1$ km, $B = 10$ Mbps, $P = 256$ bits
 - (3): $D = 10$ km, $B = 1$ Mbps, $P = 256$ bits
 - (4): $D = 1$ km, $B = 50$ Mbps, $P = 10,000$ bits
- c. A ring topology with a total circular length of $2D$, with the two stations a distance D apart. Acknowledgment is achieved by allowing a packet to circulate past the destination station, back to the source station. There are N repeaters on the ring, each of which introduces a delay of 1 bit time. Repeat the calculation for each of b1 through b4 for $N = 10; 100; 1000$.
- 3.6 A 10-story office building has the floor plan of Figure 3.7 for each floor. A local network is to be installed that will allow attachment of a device from each office on each floor. Attachment is to take place along the outside wall at the baseboard. Cable or wire can be run vertically through the indicated closet and horizontally along the baseboards. The height of each story is 10 ft. What is the minimum total length of cable or wire required for bus, tree, ring, and star topologies?
- 3.7 A tree-topology local network is to be provided that spans two buildings. If permission can be obtained to string cable between the two buildings, then one continuous tree layout will be used. Otherwise, each building will have an independent tree topology network and a point-to-point link will connect a special communications station on one network with a communications station on the other network. What functions must the communications stations perform? Repeat for ring and star.

CHAPTER 4

Topologies and Transmission Media for LANs and MANs

Recall that the three principal characteristics of a LAN or MAN are transmission medium, topology, and medium access control. In this chapter, we examine the first two of these in some detail. The complex subject of medium access control is dealt with for LANs and MANs separately in the following two chapters. The first two sections examine metallic media (twisted pair, coaxial cable) and the various topologies that are employed. The next two sections concentrate on optical fiber. Throughout, reference is made to the IEEE 802 standards, developed by a committee of the Institute of Electrical and Electronic Engineers, and the FDDI standard, developed by a committee of the American National Standards Institute. These committees and their work are discussed in the appendices to Chapters 5 and 6.

4.1

METALLIC MEDIA: BUS/TREE TOPOLOGY

Characteristics of Bus/Tree LANs

Of the topologies discussed in the preceding chapter, only the bus/tree topology is a multipoint medium. That is, there are more than two devices connected to and capable of transmitting on the medium.

The operation of this type of LAN can be summarized briefly. Because multiple devices share a single data path, only one may transmit

at a time. A station usually transmits data in the form of a packet containing the address of the destination. The packet propagates throughout the medium and is received by all other stations. The addressed station copies the packet as it goes by.

Two transmission techniques are in use for bus/tree LANs on metallic media: baseband and broadband. *Baseband*, using digital signaling, can be employed on twisted-pair or coaxial cable. *Broadband*, using analog signaling in the radio-frequency (RF) range, employs coaxial cable. Some of the differences are highlighted in Table 4.1, and this section explores the two methods in some detail. There is also a variant, known as *single-channel broadband*, that has the signaling characteristics of broadband but some of the restrictions of baseband. This is also covered below.

The multipoint nature of the bus/tree topology gives rise to several rather stiff problems. First is the problem of determining which station on the medium may transmit at any point in time. With point-to-point links (only two stations on the medium), this is a fairly simple task. If the line is full-duplex, both stations may transmit at the same time; if the line is half-duplex, a rather simple mechanism is needed to ensure that the two stations take turns. Historically, the most common shared access scheme has been the multidrop line, in which access is determined by polling from a controlling station. The controlling station may send data to any other station, or it may issue a poll to a specific station, asking for an immediate response. This method, however, negates some of the advantages of a distributed system and is awkward for communication between two noncontroller stations. A variety of distributed strategies, referred to as *medium access control protocols*, have now been developed for bus and tree topologies. These are discussed in Chapter 5.

A second problem has to do with signal balancing. When two devices exchange data over a link, the signal strength of the transmitter must be adjusted to be within certain limits. The signal must be strong enough so that after attenuation across the medium it meets the receiver's minimum signal strength requirements. It must also be strong enough to

TABLE 4.1 Bus/Tree Transmission Techniques

Baseband	Broadband
Digital signaling	Analog signaling (requires RF modem)
Entire bandwidth consumed by signal—no FDM	FDM possible—multiple data channels, video, audio
Bidirectional	Unidirectional
Bus topology	Bus or tree topology
Distance: up to a few kilometers	Distance: up to 10's of kilometers

maintain an adequate signal-to-noise ratio. On the other hand, the signal must not be so strong that it overloads the circuitry of the transmitter, which creates harmonics and other spurious signals. Although easily done for a point-to-point link, signal balancing is no easy task for a multipoint line. If any device can transmit to any other device, then the signal balancing must be performed for all permutations of stations taken two at a time. For n stations that works out to $n \times (n - 1)$ permutations. So for a 200-station network (not a particularly large system), 39,800 signal strength constraints must be satisfied simultaneously. With interdevice distances ranging from tens to thousands of meters, this is an impossible task for any but small networks. In systems that use radio-frequency (RF) signals, the problem is compounded because of the possibility of RF signal interference across frequencies. The solution is to divide the medium into segments within which pairwise balancing is possible, using amplifiers or repeaters between segments.

Baseband Systems

The principal characteristics of a baseband system are listed in Table 4.1. As mentioned earlier, a baseband LAN is defined as one that uses digital signaling. (This is a restricted use of the word *baseband*, which has become accepted in local network circles. More generally, *baseband* refers to the transmission of an analog or digital signal in its original form, without modulation.) Digital signals are inserted on the line as voltage pulses, usually using either Manchester or Differential Manchester encoding. The entire frequency spectrum of the medium is used to form the signal; hence frequency-division multiplexing (FDM) cannot be used. Transmission is bidirectional. That is, a signal inserted at any point on the medium propagates in both directions to the ends, where it is absorbed (Figure 4.1a). The digital signaling requires a bus topology. Unlike analog signals, digital signals cannot easily be propagated through the splitters and joiners required for a tree topology. Baseband systems can extend only a limited distance, about 1 km at most. This is because the attenuation of the signal, which is most pronounced at higher frequencies, causes a blurring of the pulses and a weakening of the signal to the extent that communication over larger distances is impractical.

Baseband Coax. The most well-known form of baseband bus LAN uses coaxial cable. We concentrate on those systems in this section. Unless otherwise indicated, the discussion is based on the Ethernet system [METC76, SHOC82, DIGI80] and the almost-identical IEEE standard [IEEE90b].

Most baseband coaxial cable systems use a special 50-ohm cable rather than the standard CATV 75-ohm cable. These values refer to the

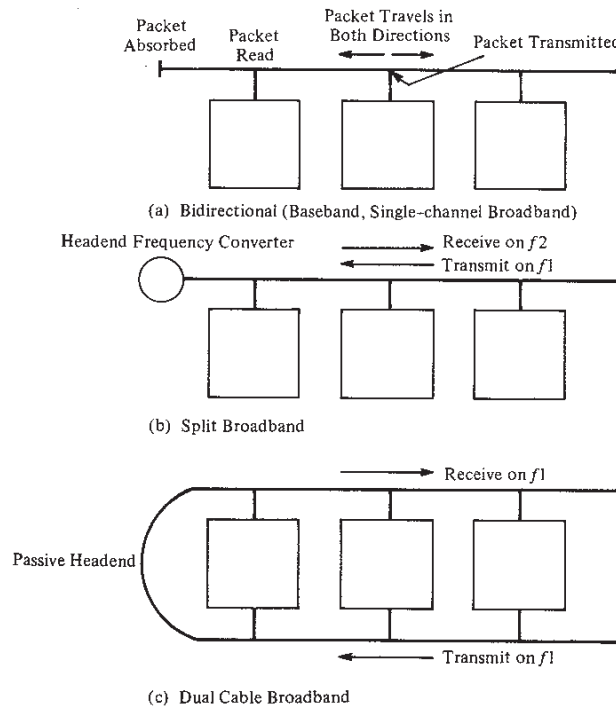


FIGURE 4.1 Baseband and Broadband Transmission Techniques

impedance of the cable. Roughly speaking, impedance is a measure of how much voltage must be applied to the cable to achieve a given signal strength (see Appendix 4A). For digital signals, the 50-ohm cable suffers less intense reflections from the insertion capacitance of the taps and provides better immunity against low-frequency electromagnetic noise. The simplest baseband coaxial bus LAN consists of an unbranched length of coaxial cable with a terminating resistance at each end. The value of the resistance is set equal to the impedance of the cable; this prevents reflection by absorbing any signal on the cable.

As with any transmission system, there are engineering trade-offs involving data rate, cable length, number of taps, and the electrical characteristics of the transmit and receive components for a baseband coaxial system. For example, the lower the data rate, the longer the cable can be. That latter statement is true for the following reason: when a signal is propagated along a transmission medium, the integrity of the signal suffers due to attenuation, noise, and other impairments. The longer the length of propagation, the greater the effect, increasing the probability of error. However, at a lower data rate, the individual pulses of a digital

signal last longer and can be recovered in the presence of impairments more easily than higher-rate, shorter pulses.

With the above in mind, we give one example that illustrates some of the trade-offs. The Ethernet specification and the original IEEE standard specified the use of 50-ohm cable with a 0.4-inch diameter and a data rate of 10 Mbps. With these parameters, the maximum length of the cable is set at 500 meters. Stations attach to the cable by means of a tap, with the distance between any two taps being a multiple of 2.5 m; this is to ensure that reflections from adjacent taps do not add in phase [YEN83]. A maximum of 100 taps is allowed. In IEEE jargon, this system is referred to as "10base5." The first two digits give the data rate in megabits per second; the four letters are an abbreviation for the medium (baseband); and the final digit is the maximum cable length in hundreds of meters.

To provide a lower-cost system for personal computer local networks, IEEE later added a 10base2 specification [METC83, FLAT84, JONE85]. Table 4.2 compares this system, dubbed Cheapernet, with the 10base5 specification. The key difference is the thinner (0.25 in) cable used in products like public address systems. The thinner cable is more flexible; thus it is easier to bend around corners and bring to a workstation cabinet rather than installing cable in the wall and having to provide a drop cable to the station. The cable is easier to install and requires cheaper electronics than the thicker cable. On the other hand, the thinner cable suffers greater attenuation and lower noise resistance than the thicker cable. Thus it supports fewer taps over a shorter distance.

Figure 4.2, from the Ethernet specification, illustrates typical components and their functions. The main components are:

- Transceiver
- Transceiver cable
- Controller
- 50-ohm coaxial cable
- 50-ohm terminators

TABLE 4.2 IEEE Specifications for 10-Mbps Baseband Coaxial Bus Local Networks

Parameter	10base5	10base2
Data Rate	10 Mbps	10 Mbps
Maximum Segment Length	500 m	200 m
Network Span	2500 m	1000 m
Nodes per Segment	100	30
Node Spacing	2.5 m	0.5 m
Cable Diameter	0.4 in	0.25 in

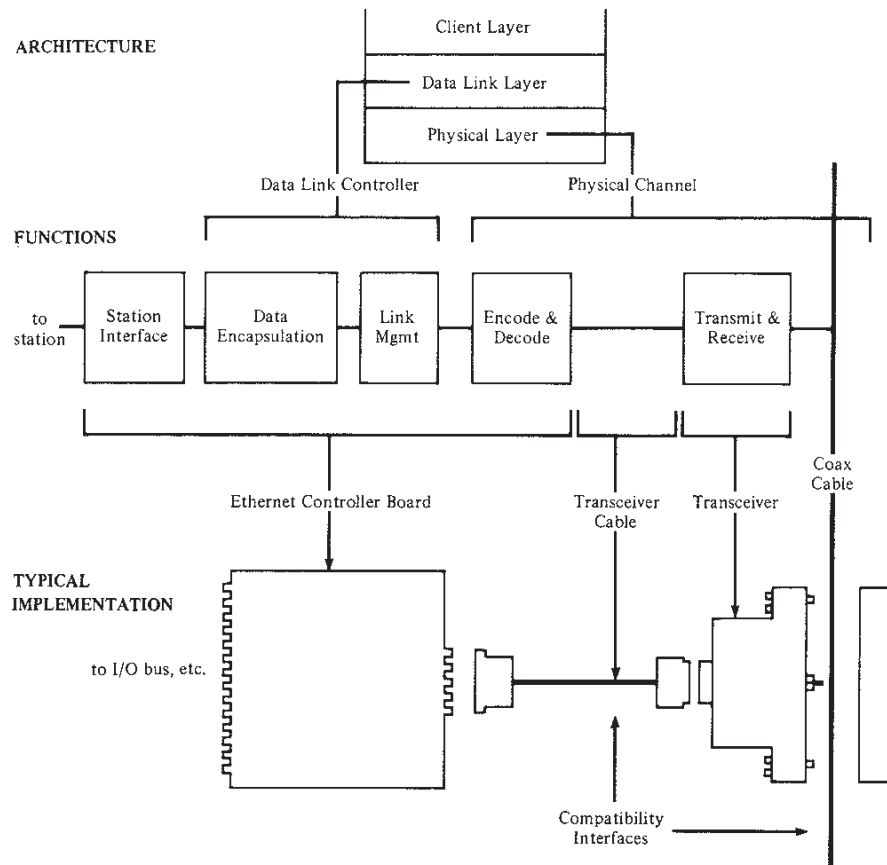


FIGURE 4.2 Ethernet Architecture and Typical Implementation (from [DIG180])

The transceiver taps into the coaxial cable. It transmits signals from the station to the cable, and vice versa. It also contains the electronics necessary to recognize the presence of a signal on the coaxial cable and to recognize a collision of two signals. This last function is needed for Ethernet and 802 because of the CSMA/CD protocol that they use (discussed in Chapter 5). A baseband bus LAN using some other protocol would not require this complexity. The transceiver also provides ground isolation between the signals from the station and the signals on the cable. Since two local grounds may differ by several volts, connection of local grounds to the cable could cause a large current to flow through the cable's shield, introducing noise and creating a safety hazard.

The transceiver cable comprises two twisted pair (referred to as twin pair) and connects the transceiver to the controller, which contains the

bulk of the intelligence required to communicate over the LAN. This split is arbitrary: all of the electronics could be included at the transceiver end. The split is motivated by the assumption that the station will be located some distance from the cable and that the cable tap may be in a relatively inaccessible location. Hence the electronics at the tap should be as simple as possible to reduce maintenance costs. The cable supplies power to the transceiver and passes data signals between the transceiver and the controller as well as control signals. The latter includes a collision presence signal from transceiver to controller. Other signals are possible. For example, the 802 standard has isolate and cease-to-isolate signals that allow the controller to enable and disable the transceiver.

The controller is an implementation of all the functions (other than those performed by the transceiver) needed to manage access to the coax cable for the purpose of exchanging packets between the coax cable and the attached station. More will be said about the particular functions in Chapter 5.

Finally, the transmission system consists of 50-ohm coaxial cable and terminators. The terminators absorb signals, preventing reflection from the ends of the bus.

These five types of components are sufficient for building a baseband bus LAN of up to about 1 km with up to about 100 stations. In many cases, this will be enough, but for greater requirements, an additional component is needed: the repeater.

The repeater is used to extend the length of the network. It consists, in essence, of two transceivers joined together and connected to two different segments of coaxial cable. The repeater passes digital signals in both directions between the two segments, amplifying and regenerating the signals as they pass through. A repeater is transparent to the rest of the system; since it does no buffering, it in no sense isolates one segment from another. So, for example, if two stations on different segments attempt to transmit at the same time, their packets will interfere with each other (collide). To avoid multipath interference, only one path of segments and repeaters is allowed between any two stations. The 802 standard allows a maximum of four repeaters in the path between any two stations, extending the effective cable length of 2.5 km. Figure 4.3 is an example of a baseband system with three segments and two repeaters.

Broadband Systems

Like the term *baseband*, the term *broadband* is co-opted into the local network vocabulary from the telecommunications world, with a change in meaning. In general, broadband refers to any channel having a bandwidth greater than a voice-grade channel (4 kHz). In the local network




FIGURE 4.3 Baseband Configuration

context, the term refers to coaxial cable on which analog signaling is used. A further restriction to transmission techniques that allow frequency-division multiplexing (FDM) on the cable is usually applied. We will mean systems capable of FDM when using the term *broadband*. Systems intended to carry only a single analog signal will be referred to as *carrierband*.

Table 4.1 summarizes the key characteristics of broadband systems. As mentioned, broadband implies the use of analog signaling. FDM is possible: the frequency spectrum of the cable can be divided into channels or sections of bandwidth. Separate channels can support data traffic, television, and radio signals. Broadband components allow splitting and joining operations; hence both bus and tree topologies are possible. Much greater distance—tens of kilometers—are possible with broadband compared to baseband. This is because the analog signals that carry the digital data can propagate greater distances before the noise and attenuation damage the data.

Dual and Split Configurations. As with baseband, stations on a broadband LAN attach to the cable by means of a tap. Unlike baseband, however, broadband is inherently a unidirectional medium; the taps that are used allow signals inserted onto the medium to propagate in only one direction. The primary reason for this is that it is infeasible to build amplifiers that will pass signals of one frequency in both directions. This unidirectional property means that only those stations “downstream”

from a transmitting station can receive its signals. How, then, can full connectivity be achieved?

Clearly, two data paths are needed. These paths are joined at a point on the network known as the **headend**. For a bus topology, the headend is simply one end of the bus. For a tree topology, the headend is the root of the branching tree. All stations transmit on one path toward the headend (inbound). Signals arriving at the headend are then propagated along a second data path away from the headend (outbound). All stations receive on the outbound path.

Physically, two different configurations are used to implement the inbound and outbound paths. (Figure 4.1 b and c). On a **dual-cable** configuration, the inbound and outbound paths are separate cables, with the headend simply a passive connector between the two. Stations send and receive on the same frequency.

By contrast, on a **split** configuration, the inbound and outbound paths are different frequency bands on the same cable. Bidirectional amplifiers¹ pass lower frequencies inbound and pass higher frequencies outbound. Between the inbound and outbound frequency bands is a guardband that carries no signals and serves merely as a separator. The headend contains a device for converting inbound frequencies to outbound frequencies.

The frequency-conversion device at the headend can be either an analog or a digital device. An analog device, known as a **frequency translator**, converts a block of frequencies from one range to another. A digital device, known as a **remodulator**, recovers the digital data from the inbound analog signal and then retransmits the data on the outbound frequency. Thus, a remodulator provides better signal quality by removing all of the accumulated noise and attenuation and transmitting a cleaned-up signal.

Split systems are categorized by the frequency allocation of the two paths, as shown in Table 4.3. Subsplit, commonly used by the cable television industry, was designed for metropolitan area television distribution, with limited subscriber-to-central office communication. It provides the easiest way to upgrade existing one-way cable systems to two-way operation. Subsplit has limited usefulness for local area networking because a bandwidth of only 25 MHz is available for two-way communication. Midsplit is more suitable for LANs, since it provides a more equitable distribution of bandwidth. However, midsplit was developed at a time when the practical spectrum of a cable-TV cable was 300 MHz, whereas a spectrum of 400 to 450 MHz is now available. Ac-

¹Unfortunately, this terminology is confusing, since we have said that broadband is inherently a unidirectional medium. At a given frequency, broadband is unidirectional. However, there is no difficulty in having signals in nonoverlapping frequency bands traveling in opposite directions on the cable, and in amplifying those signals.

TABLE 4.3 Common Cable Frequency Splits

Format	Inbound Frequency Band	Outbound Frequency Band	Maximum Two-way Bandwidth
Subsplit	5 to 30 MHz	54 to 400 MHz	25 MHz
Midsplit	5 to 116 MHz	168 to 400 MHz	111 MHz
High-split	5 to 174 MHz	232 to 400 MHz	168 MHz
Dual Cable	40 to 400 MHz	40 to 400 MHz	360 MHz

cordingly, a high-split specification has been developed to provide greater two-way bandwidth for a split-cable system.

The differences between split and dual configurations are minor. The split system is useful when a single-cable plant is already installed in a building. If a large amount of bandwidth is needed, or the need is anticipated, then a dual-cable system is indicated. Beyond these considerations, it is a matter of a trade-off between cost and size. The single-cable system has the fixed cost of the headend remodulator or frequency translator. The dual-cable system makes use of more cable, taps, splitters, and amplifiers. Thus, dual cable is cheaper for smaller systems, where the fixed cost of the headend is noticeable, and single cable is cheaper for larger systems, where incremental costs dominate.

Broadband Components. Broadband systems use standard, off-the-shelf cable television components, including 75-ohm coaxial cable. All endpoints are terminated with a 75-ohm terminator to absorb signals (see Appendix 4A). Broadband is suitable for tens of kilometers radius from the headend and hundreds or even thousands of devices. The main components of the system are:

- Cable
- Terminators
- Amplifiers
- Directional couplers
- Modems
- Controllers

Cables used in broadband networks are of three types. **Trunk cable** forms the spine of a large LAN system. Trunk cables use a semirigid construction. As the name implies, semirigid cable is not flexible. The outer portion of the cable is made of solid aluminum. The cable can be bent, but not too many times and not very easily. Trunk lines come in six sizes, ranging from 0.412 to 1 inch in diameter. The greater the diameter of the cable, the lower the attenuation. Semirigid cable has excellent noise rejection characteristics and can be used indoors and

outdoors. Typically, a trunk cable will extend from a few kilometers to tens of kilometers.

Distribution cables, or **feeder cables**, are used for shorter distances and for branch cables. They may be semirigid or flexible, and are typically 0.4 to 0.5 inch in diameter. Whereas trunk cables may be used indoors or outdoors, feeder cables are generally limited to indoor use. The choice of cable depends on a number of criteria [COOP84]:

- The physical constraints of the route: smaller-diameter cables are easier to install.
- The required signal level for the distribution network: larger-diameter cables have less signal loss.
- Local and national building codes.

The flexible cable most commonly used for feeder cable has the designation RG-11. With a diameter of 0.405 inches, and with poorer noise resistance than semirigid cable, distance is limited to about 800 meters.

Drop cables are used to connect outlets and stations to distribution cables. These are short (10 to 50 feet) and therefore need not be very large in diameter; although attenuation per unit length is greater for narrower cable, the short distance means that the total attenuation will be small even with a narrow cable. The cables used are flexible and include RG-59 (0.242 in diameter), RG-6 (0.332 inch), and RG-11 (0.405 inch) cables.

Amplifiers may be used on trunk and distribution cables to compensate for cable attenuation. As Figure 4.4 indicates, attenuation on a cable is an increasing function of frequency. Therefore, amplifiers must have a slope to account for the variability of attenuation. For split systems, amplifiers must be bidirectional, passing and amplifying lower frequencies in one direction and higher frequencies in the other.

Directional couplers provide a means for dividing one input into two outputs and combining two inputs into one output. **Splitters**, used to branch the cable, provide roughly equal attenuation along the split branches. **Taps**, used to connect drop cables and hence stations to the LAN, provide more attenuation to the drop cable. Figure 4-5 illustrates these concepts.

Modems are needed to convert between the digital data on the attached stations and the analog signal on the medium. A variety of modulation techniques are in use. The two most common, which are endorsed for use on IEEE-802-standard LANs (see Appendix 5A), are differential phase-shift keying (DPSK), used with IEEE 802.3 and duobinary AM/PSK, used with IEEE 802.4.

In ordinary PSK, a binary zero is represented by a carrier with a particular phase, and a binary one is represented by a carrier with the opposite phase (180-degree difference). DPSK makes use of differential




FIGURE 4.4 Cable Attenuation versus Frequency for Various Sizes of Coaxial Cable

encoding, in which a change of phase occurs when a zero occurs, and there is no change of phase when a one occurs. The advantage of differential encoding is that it is easier for the receiver to detect the presence or absence of a change of phase than it is to determine the phase itself.

In duobinary AM/PSK, a special narrow-bandwidth pulse is created that is used to amplitude-modulate an RF carrier. Such pulse is illustrated in Figure 4.6 for a 10-Mbps data rate; the pulse of opposite polarity is also used. Note that the pulse spreads over a number of bit time. Thus, pulses that are generated in nearby bit slots will overlap. However, the overlap is highly predictable: at each sample point, a pulse has a value of 0 or 1. Thus, at any sample point, a 0, 1, or 2 can be detected. To encode digital data, two pulses, one bit time apart, are used. A binary one is represented by two consecutive pulses of the same polarity, which will produce a sample of +2 or -2, and a binary zero is represented by two consecutive pulses of opposite polarity, which produces a sample of 0. Each pulse participates in 2 bits; that is, each pulse is both the second pulse of one bit and the first pulse of the next bit.

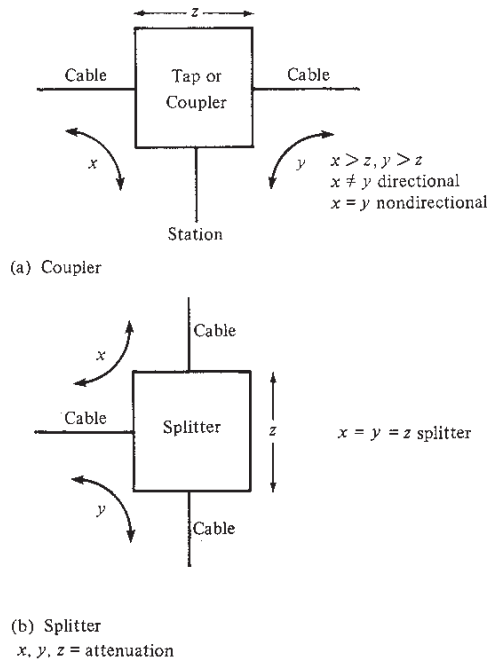


FIGURE 4.5 Directional Couplers and Splitters

A characteristic common to virtually all broadband LAN modems is the use of scrambling. This gives the data a pseudorandom nature that helps the receiver extract bit-timing information. It also improves the spectral characteristics of the signal, giving it a more uniform power distribution, as opposed to the potentially strong discrete spectral lines




FIGURE 4.6 Input Pulse for Duobinary AM/PSK at 10 Mbps

in nonscrambled data. This gives the signal better noise resistance. The scrambling process is explained in Appendix 4C.

Finally, **controllers** are needed, as in baseband, to provide the basic LAN service.

Data Transmission Services. As mentioned earlier, the broadband LAN can be used to carry multiple channels, some used for analog signals, such as video and voice, and some for digital. Digital channels can generally carry a data rate of somewhere between 0.5 and 2 bps/Hz. Figure 4.7 shows a possible allocation of a 350-MHz cable.

Three kinds of digital data transfer service are possible on a broadband cable: dedicated, switched, and multiple access (Figure 4.8). For dedicated service, a small portion of the cable's bandwidth is reserved for exclusive use by two devices. No special protocol is needed. Each of the two devices attaches to the cable through a modem; both modems are tuned to the same frequency. This technique is analogous to securing a dedicated leased line from the telephone company. The dedicated service could be used to connect two devices when a heavy traffic pattern is expected; for example, one computer may act as a standby for another and may need to get frequent updates of state information and file and database changes. Transfer rates of up to 20 Mbps are achievable.

The switched technique requires the use of a number of frequency bands. Devices are attached through *frequency-agile modems*, capable of changing their frequency by electronic command. Initially, all attached devices, together with a controller, are tuned to the same frequency. A station wishing to establish a connection sends a request to the controller, which assigns an available frequency to the two devices and signals their modems to tune to that frequency. This technique is analogous to a dial-up line. Because the cost of frequency-agile modems rises dramatically with data rate, rates of 56 kbps or less are typical. The switched technique is used in Wang's local network for terminal-to-host connections [STAH82] and could also be used for voice service.

Finally, the multiple-access service allows a number of attached devices to be supported at the same frequency. This provides for distributed peer communications among many devices, which is the primary motivation for a local network. As with baseband, some form of medium access control protocol is needed to control transmission. These protocols are discussed in Chapter 5.

Baseband versus Broadband

Table 4.4 summarizes the pros and cons of the two technologies. Baseband has the advantage of simplicity, and, in principle, lower cost. The layout of a baseband cable plant is simple; there are just five rules for




FIGURE 4.7 Dual-Cable Broadband Spectrum Allocation

trunk layout in the Ethernet specification. An office-building electrician should be able to do the job.

The potential disadvantages of baseband include the limitations in capacity and distance—disadvantages only if your requirements exceed those limitations. Another concern has to do with grounding. Because dc components are on the cable, it can be grounded in only one place.

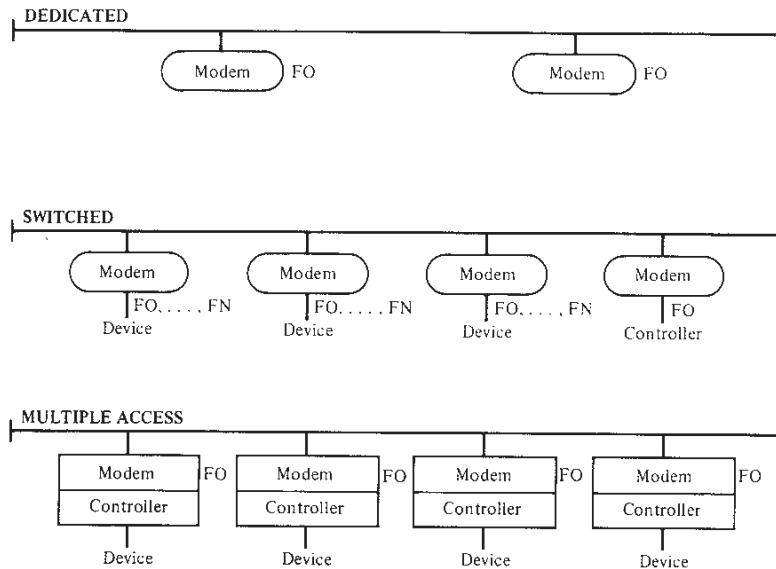


FIGURE 4.8 Broadband Data Transfer Services

Care must be taken to avoid potential shock hazards and antenna effects.

Broadband's strength is its tremendous capacity; it can carry a wide variety of traffic on a number of channels. With the use of amplifiers, broadband can achieve very wide area coverage. Also, the system is based on a mature CATV technology. Components are reliable and readily available.

TABLE 4.4 Baseband versus Broadband

Advantages	Disadvantages
	<i>Baseband</i>
Cheaper—no modem	Single channel
Simpler technology	Limited capacity
Easy to install	Limited distance
	Grounding concerns
	<i>Broadband</i>
High capacity	Modem cost
Multiple traffic types	Installation and maintenance complexity
More flexible configurations	Doubled propagation delay
Large area coverage	
Mature CATV technology	




FIGURE 4.9 Signal Levels from Trunk to Outlet [COOP84]

Broadband systems are more complex than baseband to install and maintain. The layout design must include cable type selection, and placement and setting of all amplifiers and taps. To get some feeling for the complexity of broadband cable layout design, consider Figure 4.9, which shows a small portion of a cable plant.² In order to assure that the signal level at each station or outlet is within prescribed tolerances, the engineer must consider the attenuation loss along each cable segment, the loss at each splitter and tap, and the gain at each amplifier. These losses and gains must be balanced to provide proper signal levels throughout the LAN. Maintenance involves periodic testing and alignment of all network parameters. These are jobs for experienced radio-frequency engineers.

Finally, the average propagation delay between stations on broadband is twice that for a comparable baseband system. This reduces the efficiency and performance of the broadband system, as discussed in Chapter 9.

As with all other network design choices, the selection of baseband or broadband must be based on relative costs and benefits. It is likely that some installations will have both types. Neither is likely to win the LAN war.

Carrierband Systems

There is another application of analog signaling on a LAN, known as carrierband, or single-channel broadband. In this case, the entire spec-

²The figure uses dB and dBmV units: these are explained in Appendix 4B.

trum of the cable is devoted to a single transmission path for the analog signals; no frequency-division multiplexing is possible.

Typically, a carrierband LAN has the following characteristics. Bidirectional transmission, using a bus topology, is employed. Hence there can be no amplifiers, and there is no need for a headend. Although the entire spectrum is used, most of the signal energy is concentrated at relatively low frequencies. This is an advantage, because attenuation is less at lower frequencies.

Because the cable is dedicated to a single task, it is not necessary to take care that the modem output be confined to a narrow bandwidth. Energy can spread over the entire spectrum. As a result, the electronics are simple and relatively inexpensive. Typically, some form of frequency-shift keying (FSK) is used.

Carrierband would appear to give comparable performance, at a comparable price, to baseband.

4.2

METALLIC MEDIA: STAR TOPOLOGY

In recent years, there has been increasing interest in the use of twisted pair as a transmission medium for LANs. From the earliest days of commercial LAN availability, twisted-pair bus LANs have been popular. However, such LANs suffer in comparison with a coaxial cable LAN. First of all, the apparent cost advantage of twisted pair is not as great as it might seem when a linear bus layout is used. True, twisted-pair cable is less expensive than coaxial cable. On the other hand, much of the cost of LAN wiring is the labor cost of installing the cable, which is no greater for coaxial cable than for twisted pair. Second, coaxial cable provides superior signal quality, and therefore it can support more devices over longer distances at higher data rates than twisted pair.

The renewed interest in twisted pair, at least in the context of bus/tree-type LANs, is in the use of unshielded twisted pair in a star wiring arrangement (see discussion in Section 3.3). The reason for the interest is that unshielded twisted pair is simply telephone wire, and virtually all office buildings are equipped with spare twisted pairs running from wiring closets to each office. This yields two benefits when deploying a LAN:

1. There is essentially no installation cost with unshielded twisted pair, since the wire is already there. Coaxial cable has to be pulled. In older buildings, this may be difficult since existing conduits may be crowded.
2. In most office buildings, it is impossible to anticipate all the locations where network access will be needed. Since it is extrava-

gantly expensive to run coaxial cable to every office, a coaxial cable-based LAN will typically cover only a portion of a building. If equipment subsequently has to be moved to an office not covered by the LAN, a significant expense is involved in extending the LAN coverage. With telephone wire, this problem does not arise, since all offices are covered.

The most popular approach to the use of unshielded twisted pair for a LAN is therefore a star-wiring approach. In Figure 3.5a we indicated how a star-wiring approach was compatible with a bus topology. In general, however, the products on the market use a scheme suggested by Figure 4.10, in which the central element of the star is an active element, referred to as the **hub**. Each station is connected to the hub by two twisted pairs (transmit and receive). The hub acts as a repeater: when a single station transmits, the hub repeats the signal on the outgoing line to each station.

Note that although this scheme is physically a star, it is logically a bus: a transmission from any one station is received by all other stations, and if two stations transmit at the same time, there will be a collision.

Multiple levels of hubs can be cascaded in a hierarchical configuration. Figure 4.11 illustrates a two-level configuration. There is one **header hub** (HHUB) and one or more **intermediate hubs** (IHUB). Each hub may have a mixture of stations and other hubs attached to it from below. This layout fits well with building wiring practices. Typically, there is a wiring closet on each floor of an office building and a hub can be placed in each one. Each hub could service the stations on its floor.

Figure 4.12 shows an abstract representation of the intermediate and header hubs. The header hub performs all the functions described previously for a single-hub configuration. In the case of an intermediate





FIGURE 4.10 Twisted-Pair, Star-Wiring, Logical-Bus Arrangement



HHUB = Header Hub
 IHUB = Intermediate Hub

FIGURE 4.11 Two-Level Hierarchy

hub, any incoming signal from below is repeated upward to the next higher level. Any signal from above is repeated on all lower-level outgoing lines. Thus, the logical bus characteristic is retained: a transmission from any one station is received by all other stations, and if two stations transmit at the same time, there will be a collision.

The initial version of the above scheme employed a data rate of 1 Mbps and was dubbed StarLAN [PARL85]. More recently, products op-




FIGURE 4.12 Intermediate and Header Hubs

erating at 10 Mbps have begun to appear [SCHM88, MULQ88b]. These are intended to be compatible with the 10-Mbps baseband coaxial cable bus systems, requiring only a change of transceiver. Although there is now a fair amount of practical experience with these higher-speed systems, there remains a controversy about their practicality [CLAI88, ORLO88]. Two reasons for this controversy can be stated:

1. Existing telephone wire in buildings can be inadequate for data transmission. Problems include twisted pair that is not twisted, splicing and other connections, and other faults that are not noticeable for voice transmission but that would produce very high error rates at 10 Mbps.
2. Twisted-pair cables are rather tightly packed together in conduits. The mutual capacitance from adjacent pairs adversely affects attenuation, cross-talk, and velocity of propagation. The effects on data transmission may not be noticeable at 1 Mbps, but become a problem at 10 Mbps.

These problems can to some extent be overcome by the use of signal processing techniques and by careful design of the transceiver. However, just as we saw with the 10-Mbps coaxial cable bus, there are trade-offs to be made. In this case, IEEE recommends a maximum distance between station and hub of 250 meters at 1 Mbps and 100 meters at 10 Mbps.

4.3

METALLIC MEDIA: RING TOPOLOGY

Description

The ring consists of a number of repeaters, each connected to two others by unidirectional transmission links to form a single closed path (Figure 4.13). Data are transferred sequentially, bit by bit, around the ring from one repeater to the next. Each repeater regenerates and retransmits each bit.

For a ring to operate as a communications network, three functions are required: data insertion, data reception, and data removal. These functions are provided by the repeaters. Each repeater, in addition to serving as an active element on the ring, serves as a device attachment point for data insertion. Data are transmitted in packets, each of which contains a destination address field. As a packet circulates past a repeater, the address field is copied to the attached station. If the station recognizes the address, then the remainder of the packet is copied.

A variety of strategies can be used for determining how and when packets are added to and removed from the ring. The strategy can be

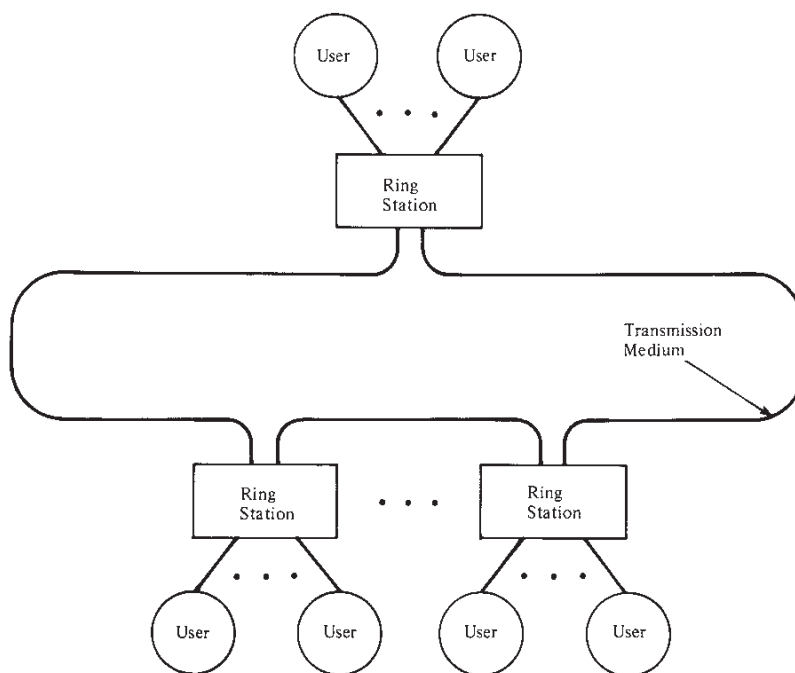


FIGURE 4.13 Ring System

viewed, at least conceptually, as residing in a medium access control layer, discussed in Chapter 5.

Repeaters perform the data insertion and reception functions in a manner not unlike that of taps, which serve as device attachment points on a bus or tree. Data removal, however, is more difficult on a ring. For a bus or tree, signals inserted onto the line propagate to the end points and are absorbed by terminators. Hence, shortly after transmission ceases, the bus or tree is clear of data. However, because the ring is a closed loop, data will circulate indefinitely unless removed. A packet may be removed by the addressed repeater. Alternatively, each packet could be removed by the transmitting repeater after it has made one trip around the loop. The latter approach is more desirable because (1) it permits automatic acknowledgement, and (2) it permits multicast addressing: one packet sent simultaneously to multiple stations.

The repeater, then, can be seen to have two main purposes: (1) to contribute to the proper functioning of the ring by passing on all the data that come its way, and (2) to provide an access point for attached stations to send and receive data. Corresponding to these two purposes are two states (Figure 4.14): the listen state and the transmit state.

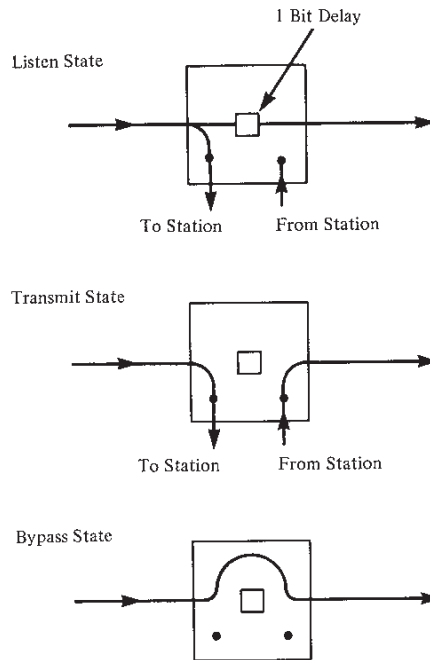


FIGURE 4.14 Ring Repeater States

In the *listen state*, each bit that is received is retransmitted with a small delay, required to allow the repeater to perform necessary functions. Ideally, the delay should be on the order of 1 bit time (the time it takes for a repeater to transmit 1 complete bit onto the outgoing line). These functions are:

- Scan passing bit stream for pertinent patterns. Chief among these is the address or addresses of attached devices. Another pattern, used in the token control strategy explained later, indicates permission to transmit. Note that to perform the scanning function, the repeater must have some knowledge of packet format.
- Copy each incoming bit and send it to the attached station, while continuing to retransmit each bit. This will be done for each bit of each packet addressed to this station.
- Modify a bit as it passes by. In certain control strategies, bits may be modified to, for example, indicate that the packet has been copied. This would serve as an acknowledgment.

When a repeater's station has data to send and when the repeater, based on the control strategy, has permission to send, the repeater enters the *transmit state*. In this state, the repeater receives bits from the

station and retransmits them on its outgoing link. During the period of transmission, bits may appear on the incoming ring link. There are two possibilities, and they are treated differently:

1. The bits could be from the same packet that the repeater is still sending. This will occur if the bit length of the ring is shorter than the packet. In this case, the repeater passes the bits back to the station, which can check them as a form of acknowledgment.
2. For some control strategies, more than one packet could be on the ring at the same time. If the repeater, while transmitting, receives bits from a packet it did not originate, it must buffer them to be transmitted later.

These two states, listen and transmit, are sufficient for proper ring operation. A third state, the *bypass state*, is also useful. In this state, a bypass relay is activated, so that signals propagate past the repeater with no delay other than medium propagation. The bypass relay affords two benefits: (1) it provides a partial solution to the reliability problem, discussed later, and (2) it improves performance by eliminating repeater delay for those stations that are not active on the network.

Ring Benefits

A good deal of research into overcoming some of the weaknesses of the ring has been done at Massachusetts Institute of Technology [SALT79, SALT83] and at IBM [BUX83, STRO83, DIXO83]. The result has been a proliferation of ring-based LAN products, most notably the appearance of the IBM product in 1985, followed by a number of compatible products from other vendors [DERF86, STRO86].

Like the bus and tree, the ring is a shared-access or multiaccess network (although the medium itself is a collection of point-to-point links). Hence the ring shares the same benefits as the bus/tree, including ability to broadcast and incremental cost growth. There are other benefits provided by the ring that are not shared by the bus/tree topology.

The most important benefit or strength of the ring is that it uses point-to-point communication links. There are a number of implications of this fact. First, because the transmitted signal is regenerated at each node, greater distances can be covered than with baseband bus. Broadband bus/tree can cover a similar range, but cascaded amplifiers can result in loss of data integrity at high data rates. Second, the ring can accommodate optical fiber links that provide very high data rates and excellent electromagnetic interference (EMI) characteristics. Finally, the electronics and maintenance of point-to-point lines are simpler than for multipoint lines.

Another benefit of the ring is that fault isolation and recovery are simpler than for bus/tree. This is discussed in more detail later in this section.

With the ring, the duplicate address problem is easily solved. If, on a bus or tree, two stations are by accident assigned the same address, there is no easy way to sort this out. A relatively complex algorithm must be incorporated into the LAN protocol. On a ring, the first station with an address match that is encountered by a packet can modify a bit in the packet to acknowledge reception. Subsequent stations with the same address will easily recognize the problem.

Finally, there is the potential throughput of the ring. Under certain conditions, the ring has greater throughput than a comparable bus or tree LAN. This topic is explored in Chapter 9.

Potential Ring Problems

The potential problems of a ring are, at first blush, more obvious than the benefits:

1. *Cable vulnerability:* A break on any of the links between repeaters disables the entire network until the problem can be isolated and a new cable installed. The ring may range widely throughout a building and is vulnerable at every point to accidents.
2. *Repeater failure:* As with the links, a failure of a single repeater disables the entire network. In many networks, it will be common for many of the stations not to be in operation at any time; yet all repeaters must always operate properly.
3. *Perambulation:* When either a repeater or a link fails, locating the failure requires perambulation of the ring, and thus access to all rooms containing repeaters and cable. This is known as the "pocket full of keys" problem.
4. *Installation headaches:* Installation of a new repeater to support new devices requires the identification of two nearby, topologically adjacent repeaters. It must be verified that they are in fact adjacent (documentation could be faulty or out of date), and cable must be run from the new repeater to both of the old repeaters. There are several unfortunate consequences. The length of cable driven by the source repeater may change, possibly requiring retuning. Old cable, if not removed, accumulates. In addition, the geometry of the ring may become highly irregular, exacerbating the perambulation problem.
5. *Size limitations:* There is a practical limit to the number of repeaters on a ring. This limit is suggested by the reliability and maintenance problems cited earlier, the timing jitter discussed below, and

the accumulating delay of large numbers of repeaters. A limit of a few hundred repeaters seems reasonable.

6. *Initialization and recovery:* To avoid designating one ring node as a controller (negating the benefit of distributed control), a strategy is required to assure that all stations can cooperate smoothly when initialization and recovery are required. This need arises, for example, when a packet is garbled by a transient line error; in that case, no repeater may wish to assume the responsibility of removing the circulating packet.
7. *Timing jitter:* This is a subtle problem having to do with the clocking or timing of a signal in a distributed network. It is discussed below.

Problems 1 and 2 are reliability problems. However, these two problems, together with problems 3, 4, and 5 can be ameliorated by a refinement in the ring architecture, explained in the next section. Problem 6 is a software problem, to be dealt with by the various LAN protocols discussed in Chapter 5. Problem 7 is discussed next.

Timing Jitter. On a twisted-pair or coaxial-cable ring LAN, digital signaling is generally used with biphase encoding, typically Differential Manchester. As data circulate around the ring, each receiver must recover the binary data from the received signal. To do this, the receiver must know the starting and ending times of each bit, so that it can sample the received signal properly. This requires that all the repeaters on the ring be synchronized, or clocked, together. Recall from Chapter 2 that biphase codes are self-clocking; the signal includes a transition in the middle of each bit time. Thus each repeater recovers clocking as well as data from the received signal. This clock recovery will deviate in a random fashion from the mid-bit transitions of the received signal for several reasons, including noise during transmission and imperfections in the receiver circuitry. The predominant reason, however, is delay distortion. Delay distortion is caused by the fact that the velocity of propagation of a signal through a guided medium varies with frequency. The effect is that some of the signal components of one pulse will spill over into other pulse positions; this is known as *intersymbol interference*. The deviation of clock recovery is known as *timing jitter*.

As each repeater receives data, it recovers the clocking for two purposes: first to know when to sample the incoming signal to recover the data, and second, to use the clocking for transmitting the Differential Manchester signal to the next repeater. The repeater issues a clean signal with no distortion. However, since the clocking is recovered from the incoming signal, the timing error is not eliminated. Thus the digital pulse width will expand and contract in a random fashion as the signal travels around the ring and the timing jitter accumulates. The cumulative effect of the jitter is to cause the bit latency, or bit length, of the ring

to vary. However, unless the latency of the ring remains constant, bits will be dropped (not retransmitted) as the latency of the ring decreases or added as the latency increases.

Thus timing jitter places a limitation on the number of repeaters in a ring. Although this limitation cannot be entirely overcome, several measures can be taken to improve matters [KELL83, HONG86]; these are illustrated in Figure 4.15. First, each repeater can include a phase-locked loop (PLL). This is a device that uses feedback to minimize the deviation from one bit time to the next. Although the use of phase-locked loops reduces the jitter, there is still an accumulation around the ring. A supplementary measure is to include a buffer in one of the repeaters, usually designated as the monitor repeater or station. Bits are written in using the recovered clock and are read out using a crystal master clock. The buffer is initialized to hold a certain number of bits and expands and contracts as needed. For example, the IEEE standard specifies a 6-bit buffer, which is initialized to hold 3 bits. That is, as bits come in, they are placed in the buffer for 3 bit times before being retransmitted. If the received signal at the monitor station is slightly faster than the master clock, the buffer will expand, as required, to 4, 5, or 6 bits to avoid dropping bits. If the received signal is slow, the buffer will contract to 2, 1, or 0 bits to avoid adding bits to the repeated bit stream. Thus the cleaned-up signals that are retransmitted are purged of the timing jitter. This combination of PLLs and a buffer significantly increases maximum feasible ring size. The actual limit will depend on the characteristics of the transmission medium, which determine the amount of delay distortion and therefore the amount of accumulated jitter. For example, the IBM ring product specifies a maximum of 72 repeaters in a ring using

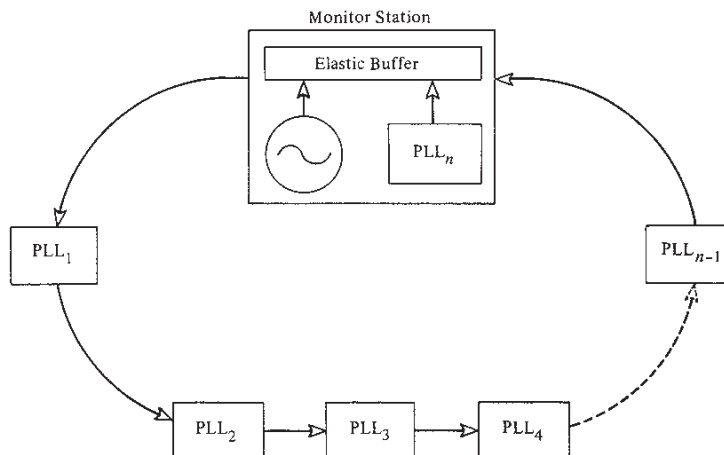


FIGURE 4.15 Ring Synchronization

unshielded twisted pair, and a maximum of 260 repeaters in a ring using shielded twisted pair.

The Star-Ring Architecture

Two observations can be made about the basic ring architecture described above. First, there is a practical limit to the number of repeaters on a ring. As was mentioned above, a number of factors combine to limit the practical size of a ring LAN to a few hundred repeaters. Second, the cited benefits of the ring do not depend on the actual routing of the cables that link the repeaters.

These observations have led to the development of a refined ring architecture, the star ring, which overcomes some of the problems of the ring and allows the construction of large local networks [SALW83]. This architecture uses the star wiring strategy discussed in the previous chapter. It is the basis of IBM's ring product and similar products.

As a first step, consider the rearrangement of a ring into a star. This is achieved by having the interrepeater link all threads through a single site (Figure 4.16). This ring wiring concentrator has a number of advantages. Because there is access to the signal on every link, it is a simple matter to isolate a fault. A message can be launched into the ring and tracked to see how far it gets without mishap. A faulty segment can be disconnected—no pocket full of keys needed—and repaired at a later time. New repeaters can easily be added to the ring: simply run two cables from the new repeater to the site of ring wiring concentration and splice into the ring.

The bypass relay associated with each repeater can be moved into the ring wiring concentrator. The relay can automatically bypass its repeater and two links for any malfunction. A nice effect of this feature is that the transmission path from one working repeater to the next is approximately constant; thus the range of signal levels to which the transmission system must automatically adapt is much smaller.

The ring wiring concentrator greatly alleviates the perambulation and installation problems mentioned earlier. It also permits rapid recovery from a cable or repeater failure. Nevertheless, a single failure could, at least temporarily, disable the entire network. Furthermore, throughput and jitter considerations still place a practical upper limit on the number of repeaters in a ring. Finally, in a spread-out network, a single wire concentration site dictates a lot of cable.

To attack these remaining problems, consider a local network consisting of multiple rings. Each ring consists of a connected sequence of wiring concentrators, and the set of rings is connected by a bridge (Figure 4.17). The bridge routes data packets from one ring subnetwork to another, based on addressing information in the packet so routed. From a physical point of view, each ring operates independently of the other

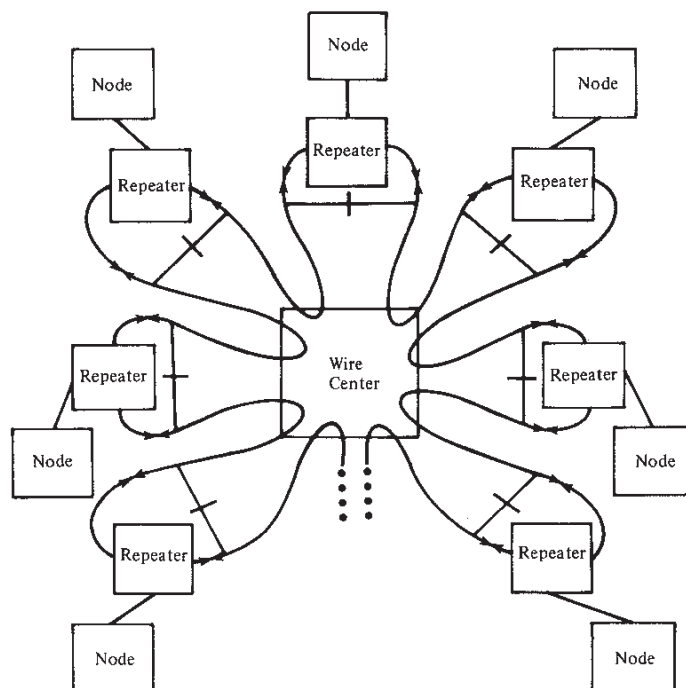


FIGURE 4.16 Ring Wiring Concentrator

rings attached to the bridge. From a logical point of view, the bridge provides transparent routing among the rings.

The bridge must perform five functions:

1. *Input filtering*: For each ring, the bridge monitors the traffic on the ring and copies all packets addressed to other rings on the bridge. This function can be performed by a repeater programmed to recognize a family of addresses rather than a single address.
2. *Input buffering*: Received packets may need to be buffered, either because the inter-ring traffic is peaking, or because the target output buffer is temporarily full.
3. *Switching*: Each packet must be routed through the bridge to its appropriate destination ring.
4. *Output buffering*: A packet may need to be buffered at the threshold of the destination ring, waiting for an opportunity to be inserted.
5. *Output transmission*: This function can be performed by an ordinary repeater.

For a small number of rings, a bridge can be a reasonably simple device. As the number of rings on a bridge grows, the switching com-

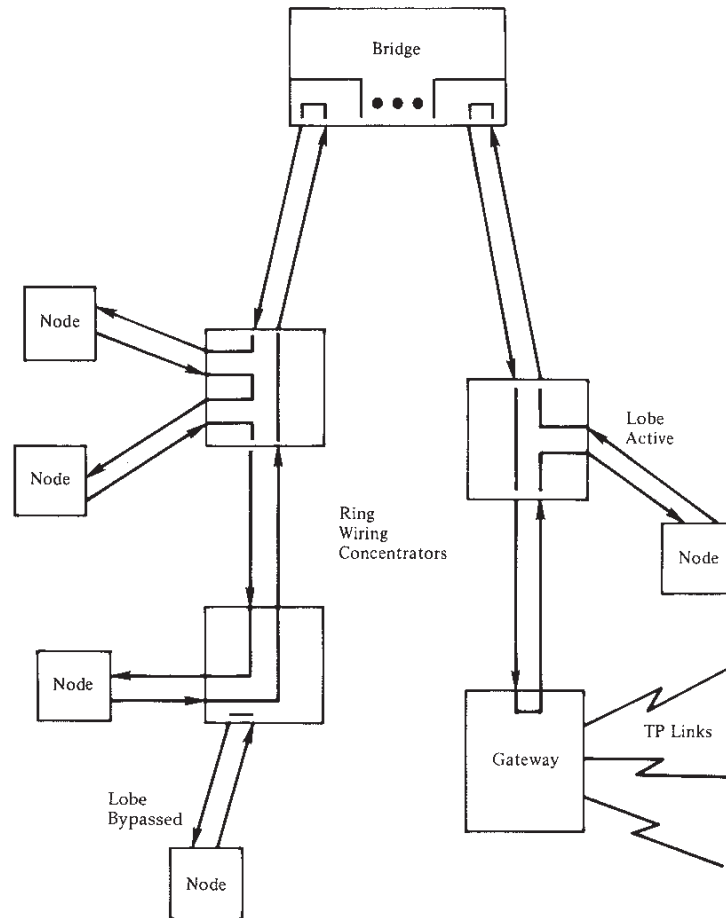


FIGURE 4.17 Ring Bridge

plexity and load on the bridge also grow. For very large installations, multiple bridges, interconnected by high-speed trunks, may be needed (Figure 4.18).

Three principal advantages accrue from the use of a bridge. First, the timing jitter problem, which becomes more difficult as the number of repeaters on a ring grows, is bounded by restricting the size of the ring. Second, the failure of a ring, for whatever reason, will disable only a portion of the network; failure of the bridge does not prevent intraring traffic. Finally, multiple rings may be employed to obtain a satisfactory level of performance when the throughput capability of a single ring is exceeded.

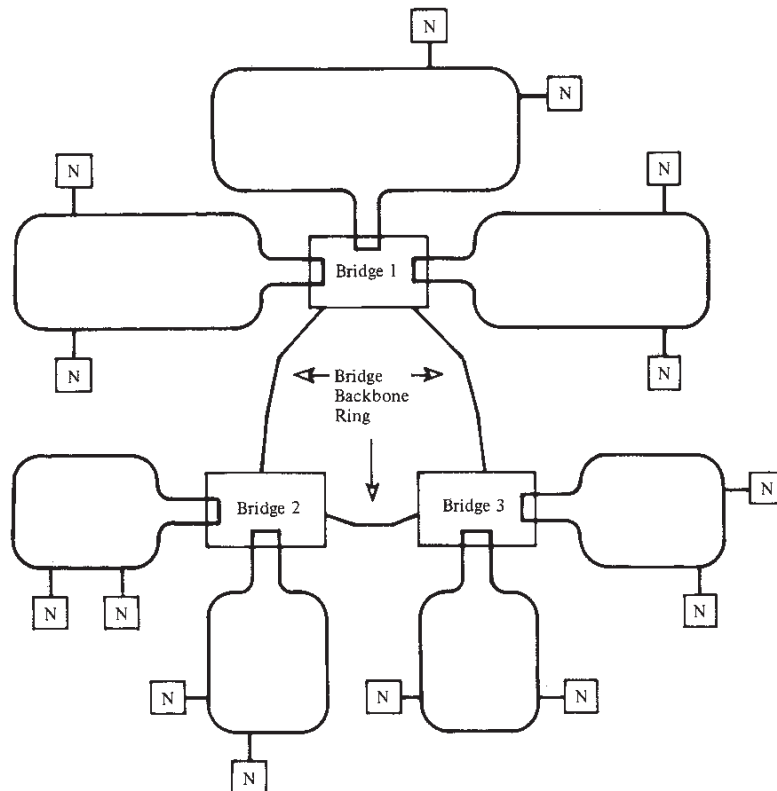


FIGURE 4.18 Multiple Bridges

There are several pitfalls to be noted. First, the automatic acknowledgment feature of the ring is lost; higher-level protocols must provide acknowledgment. Second, performance may not significantly improve if there is a high percentage of inter-ring traffic. If it is possible to do so, network devices should be judiciously allocated to rings to minimize inter-ring traffic.

4.4

OPTICAL FIBER STAR

The earliest work on optical fiber LANs focused on the star topology. Two general approaches have been investigated: the passive star and the active star. We examine each of these in turn.

Passive Star

One of the first commercially available approaches for fiber LANs was the passive star coupler [RAWS78, SCHO88]. The passive star coupler is fabricated by fusing together a number of optical fibers. Any light input to one of the fibers on one side of the coupler will be equally divided among and output through all the fibers on the other side. To form a network, each device is connected to the coupler with two fibers, one for transmit and one for receive (Figure 4.19). All of the transmit fibers enter the coupler on one side, and all of the receive fibers exit on the other side. Thus, although the arrangement is physically a star, it acts like a bus: a transmission from any one device is received by all other devices, and if two devices transmit at the same time, there will be a collision.

Two methods of fabrication of the star coupler have been pursued: the biconic fused coupler, and the mixing rod coupler. In the biconic fused coupler [STRA87], the fibers are bundled together. The bundled fibers are heated with an oxyhydrogen flame and pulled into a biconical tapered shape. That is, the rods come together into a fused mass that tapers into a conical shape and then expands back out again. The mixing rod approach [OHS86] begins in the same fashion. Then, the biconical taper is cut at the waist and a cylindrical rod is inserted between the tapers and fused to the two cut ends. This latter technique allows the use of a less narrow waist and is easier to fabricate.




FIGURE 4.19 Optical Fiber Passive Star Configuration

Commercially available passive star couplers can support a few tens of stations at a radial distance of up to a kilometer or more. Figure 4.20 shows the operating range of the two types of couplers. The limitations on number of stations and distances are imposed by the losses in the network. With today's equipment, the optical power loss between transmitter and receiver that can be tolerated is on the order of 25 to 30 dB. In the figure, the outer edge of each region is defined by a maximum end-to-end attenuation of 30 dB. The attenuation that will occur in the network consists of the following components:

- *Optical connector losses:* Connectors are used to splice together cable segments for increased length. Typical connector losses are 1.0 to 1.5 dB per connector. A typical passive star network will have from 0 to 4 connectors in a path from transmitter to receiver, for a total maximum attenuation of 4 to 6 dB.
- *Optical cable attenuation:* Typical cable attenuation for the cable that has been used in these systems ranges from 3 to 6 dB per kilometer.
- *Optical power division in the coupler:* The coupler divides the optical power from one transmission path equally among all reception paths. Expressed in decibels, the loss seen by any node is $10 \log N$, where N is the number of nodes. For example, the effective loss in a 16-port coupler is about 12 dB.

As Figure 4.20 indicates, the passive star coupler is quite limited. One promising approach to improving performance is to use an optical amplifier. In 1989, it was demonstrated that an optical signal can be directly




FIGURE 4.20 Operating Range for Optical Fiber Passive Star LAN [SCHO88]

amplified with low noise by an amplifier known as the erbium-doped fiber amplifier (EDFA). The EDFA is applicable as a high-power post-amplifier at the sending port, as a low-noise preamplifier at the receiving port, and as an intermediate repeating in-line optical amplifier [NAKA90, PARK92]. In the context of the passive star topology, the EDFA can be used to amplify signals as they pass through the star coupler.

The basic technique of the EDFA is as follows (Figure 4.21). A short segment of optical fiber is doped with erbium atoms. A constant laser input at a given wavelength, known as a laser pump, is applied to that portion of the fiber. When a laser signal at a different wavelength encounters erbium atoms that are excited to higher energy levels by a pumping light, the power of the signal light gradually increases along the optical fiber.

The use of the EDFA at the star coupler allows the implementation of networks with a greater number of stations operating over longer distances at higher data rates than can be achieved with an ordinary passive star coupler [IRSH92]. As yet, these devices are not commercially practical, but we can expect to see products in a few years.

Active Star

For a number of years, work has been underway at the Xerox Palo Alto Research Center to develop an improved version of the star topology fiber LAN. The result is Fibernet II [SCHM83], which differs from the passive star only in that the central coupler is an active repeater rather than a passive device. However, like the passive star, the active star appears as a bus to the attached devices: a transmission from any one device is received by all other devices, and only one device at a time can successfully transmit.

Figure 4.22 is a schematic diagram of Fibernet II. As before, each device attaches to the central node through two optical fiber cables, one for transmit and one for receive. Figure 4.16 reveals the internal organization of the node. When a station transmits, the receiver module de-


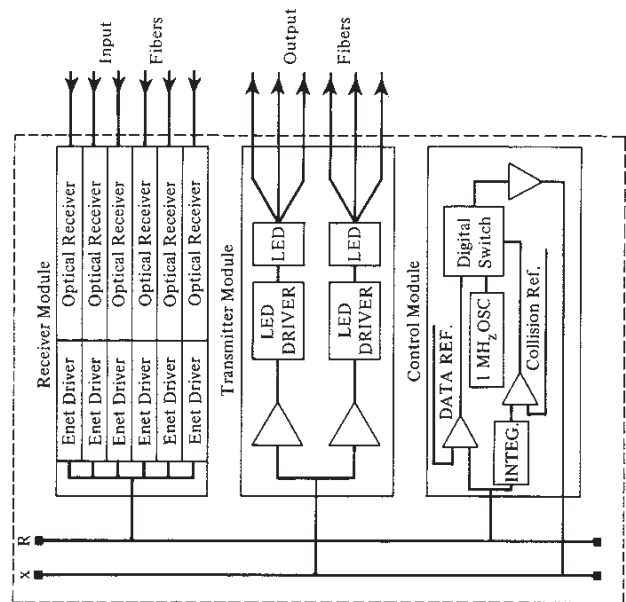
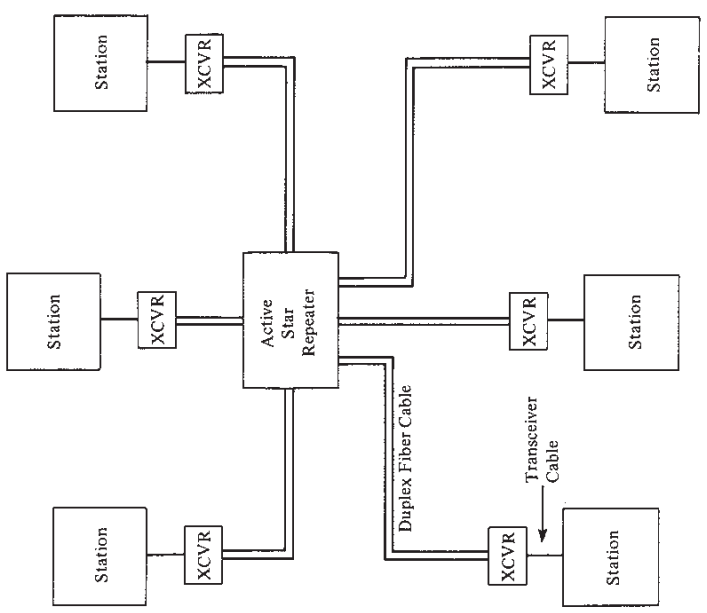


FIGURE 4.21 Erbium-Doped Fiber Amplifier [NAKA90]



b. Active Star Organization



a. Overall Diagram

FIGURE 4.22 Fibernet II Active Star Configuration

pects the optical signal on the inbound optical fiber and retransmits it on a backplane bus designated *R*. The bus is in fact a miniature 50-ohm coaxial cable. The signal on the *R* bus is received by the control module, which retransmits it on another 50-ohm coaxial bus designated *X*. The purpose of this intermediate module is to perform collision detection, a function discussed in the next chapter. Finally, the transmitter module picks up the signal from bus *X* and retransmits it in optical form on all output fibers. The delay for this entire process is on the order of a few bit times. It can be seen that this arrangement is in fact simply a bus topology using star wiring, as depicted in Figure 3.5a.

The active star has several advantages over the passive star. In the passive star, the incoming signal is split equally among all outgoing fibers, so that the greater the number of fibers, the greater the loss on any one path. With the active star, this loss does not occur. Thus the active star can support more devices over a greater distance. Fibernet II is designed to support up to over a hundred devices at a maximum radius of 2.5 km. The disadvantage of the active star is that it is more expensive due to the active components in the central node.

4.5

OPTICAL FIBER RING

Even with the use of EDFAs, the optical fiber star configuration, at least for the near future, will be limited to relatively low speeds (for optical fiber) and modest distances. The optical fiber ring, on the other hand, is well suited to providing high data rates over long distances, better exploiting the potential of optical fiber. The ring consists of a series of point-to-point links, and the technology for point-to-point fiber transmission is well understood and widely available. In addition to the other advantages of fiber cited earlier, it exhibits significantly less delay distortion than coaxial cable or twisted pair and hence suffers less from timing jitter, which means that larger ring networks can be constructed.

Because of the high data rates attainable with optical fiber, the fiber ring is a natural choice for a very high-speed LAN or for a MAN. The fiber-distributed data interface (FDDI) is such a network, and we look at the details of its physical-layer specification in Chapter 6. In this section, we briefly look at some considerations for a fiber ring LAN of lower speed and smaller geographic extent than FDDI. The trade-off, clearly, is one of cost. By limiting the design to a relatively low speed and to a relatively short distance, a relatively inexpensive fiber LAN can be developed.

As an example, we will use the specifications developed by IBM for its fiber ring product [SEE86]; these are representative of what is com-

TABLE 4.5 IBM Optical Fiber Ring Specification (SEE86)

Core Diameter	100 μm
Cladding Diameter	143 μm
Wavelength	850 nm
Attenuation	<6 dB/km
Bandwidth	>150 MHz
Data Rate	up to 20 Mbps
Distance	1.5 to 2.0 km

mercially available and commercially feasible. IBM's fiber ring specification was written to satisfy current transmission requirements using a light wavelength of 840 nm. The fiber specification also supports upward migration to higher-performance networks operating at a wavelength of 1300 nm. Although the latter could support a higher data rate, the transmitters and receivers operating at that wavelength are considerably more expensive than 850-nm devices.

Table 4.5 lists the key parameters of the specification. At an 850-nm wavelength, relatively low-cost LED transmitters and PIN detectors are used. Transmission is in the graded-index mode. A data rate of up to 20 Mbps is achievable with a maximum single-link distance of up to 1.5 to 2 km. The system should be able to support at least as many repeaters on a single ring as shielded twisted pair, on the order of 250.

4.6


OPTICAL FIBER BUS

Several approaches can be taken in the design of a fiber bus topology LAN or MAN [MUKH91]. The differences have to do with the nature of the taps into the bus and the detailed topology.


Optical Fiber Taps

With an optical fiber bus, either an active or passive tap can be used; both are permissible with the 802.6 standard. In the case of an active tap (Figure 4.23a), the following steps occur:

1. Optical signal energy enters the tap from the bus.
2. Clocking information is recovered from the signal and the signal is converted to an electrical signal.



a. Active Tap



b. Passive Tap

FIGURE 4.23 Optical Fiber Bus Taps

3. The converted signal is presented to the node and perhaps modified by the latter.
4. The optical output (a light beam) is modulated according to the electrical signal and launched into the bus.

In effect, the bus consists of a chain of point-to-point links, and each node acts as a repeater. Each tap actually consists of two of these active couplers and requires two fibers. This is because of the inherently unidirectional nature of the device in Figure 4.23a.

In the case of a passive tap (Figure 4.23b), the tap extracts a portion of the optical energy from the bus for reception and it injects optical energy directly into the medium for transmission. Thus, there is a single run of cable rather than a chain of point-to-point links. This passive approach is equivalent to the type of taps typically used for twisted pair and coaxial cable. Each tap must connect to the bus twice, once for transmit and once for receive.

The electronic complexity and interface cost are drawbacks for the implementation of the active tap. Also, each tap will add some increment of delay, just as in the case of a ring. For passive taps, the lossy nature of pure optical taps limits the number of devices and the length

of the medium. However, the performance of such taps has improved sufficiently in recent years to make fiber bus networks practical [ZANG91].

Optical Fiber Bus Configurations

A variety of configurations for the optical fiber bus have been proposed. All of these fall into two categories: those that use a single bus and those that use two buses.

Figure 4.24a shows a typical single-bus configuration, referred to as a loop bus. The operation of this bus is essentially the same as for the dual-bus broadband coaxial system described earlier. Each station transmits on the bus in the direction toward the headend and receives on the bus in the direction away from the headend. In addition to the two connections shown, some medium access control (MAC) protocols require that each station have an additional *sense tap* on the inbound (toward the headend) portion of the bus. The sense tap is able to sense the presence or absence of light on the fiber, but is not able to recover data.

Figure 4.24b shows the two-bus configuration. Each station attaches to both buses and has both transmit and receive taps on both buses. On each bus, a station may transmit only to those stations downstream from it. By using both buses, a station may transmit to and receive from




FIGURE 4.24 Optical Fiber Bus Configurations

all other stations. A given node, however, must know which bus to use to transmit to another node; if not, all data would have to be sent out on both buses. This is the configuration used in the IEEE 802.6 MAN, and it is described in Chapter 6.

4.7

RECOMMENDED READING

[MART89] covers many of the topics in this chapter and discusses commercial implementations. [NAUG91] does the same at a less technical level. [SLON91] contains a number of useful papers on these topics.

For a detailed look at baseband systems, the original Ethernet article [METC76] and a later follow-up article [SHOC82] remain informative. Detailed discussions of broadband LANs can be found in [COOP84] and [KIM88]. [MUKH91] and [HENR89] are good surveys of fiber LAN/MAN technology.

4.8

PROBLEMS

- 4.1 Consider a baseband bus with a number of equally spaced stations. As a fraction of the end-to-end propagation delay, what is the mean delay between stations? What is it for broadband bus? Now, rearrange the broadband bus into a tree with N equal-length branches emanating from the headend; what is the mean delay?
- 4.2 Give examples of appropriate applications of the broadband dedicated service and the switched service.
- 4.3 Consider a baseband bus with a number of equally spaced stations with a data rate of 10 Mbps and a bus length of 1 km. What is the average time to send a packet of 1000 bits to another station, measured from the beginning of transmission to the end of reception? Assume a propagation speed of 200 m/ μ s. If two stations begin to transmit at exactly the same time, their packets will interfere with each other. If each transmitting station monitors the bus during transmission, how long before it notices an interference, in seconds? In bit times?
- 4.4 Repeat Problem 4.3 for a data rate of 1 Mbps.
- 4.5 Repeat Problems 4.3 and 4.4 for broadband bus.
- 4.6 Repeat Problems 4.3 and 4.4 for a broadband tree consisting of 10 cables of length 100 m emanating from a headend.

- 4.7 Reconsider Problem 3.6. Can a baseband bus following the IEEE 802 rules (500-m segments, maximum of four repeaters in a path) span the building? If so, what is the total cable length?
- 4.8 Reconsider Problem 3.6 for a broadband tree. Can the total length be reduced compared to the broadband bus?
- 4.9 Reconsider Problem 3.6, but now assume that there are two rings, with a bridge on floor 5 and a ring wiring concentrator on each floor. The bridge and concentrators are located in closets along the vertical shaft.
- 4.10 At a propagation speed of 200 m/ μ s, what is the effective length added to a ring by a bit delay at each repeater:
 - a. At 1 Mbps?
 - b. At 40 Mbps?
- 4.11 System A consists of a single ring with 300 stations, one per repeater. System B consists of three 100-station rings linked by a bridge. If the probability of a link failure is P_l , a repeater failure is P_r , and a bridge failure is P_b , derive an expression for parts (a) through (d):
 - a. Probability of failure of system A.
 - b. Probability of complete failure of system B.
 - c. Probability that a particular station will find the network unavailable, for systems A and B.
 - d. Probability that any two stations, selected at random, will be unable to communicate, for systems A and B.
 - e. Compare values of 4a through 4d for $P_l = P_b = P_r = 10^{-2}$.
- 4.12 Consider two rings of 100 stations each joined by a bridge. The data rate on each link is 10 Mbps. Each station generates data at a rate of 10 packets of 2000 bits each per second. Let F be the fraction of packets on each ring destined for the other. What is the minimum throughput of the bridge required to keep up?

APPENDIX 4A: CHARACTERISTIC IMPEDANCE

An important parameter associated with any transmission line is its characteristic impedance. To understand its significance, we need to consider the electrical properties of a transmission line. Any transmission line has both inductance and capacitance, which are distributed along the entire length of the line. These quantities can be expressed in terms of inductance and capacitance per unit length.

An infinite transmission line has similar electrical properties to the circuit depicted in Figure 4.24a and b. Of course, the actual inductance

and capacitance are distributed uniformly along the line and not lumped as shown in the figure, but the equivalent circuit is good enough to explain the behavior of an actual line.

Figure 4.25a shows a section of an infinite line connected to a voltage source. Closing the switch (Figure 4.25b) will cause current to flow. Now, in a finite line, at steady state, the inductors will behave as short circuits (zero resistance) and the capacitors as open circuits (infinite resistance). However, at the instance that the switch is closed, current will flow and be resisted by the inductance and capacitance. The process will continue indefinitely because there is an infinite number of capacitors to be charged. There will be a definite relationship between the applied voltage and the amount of current that will flow. The relationship will depend only on the value of inductance and capacitance, which in turn depend on the physical dimensions of the line. In our example, an applied voltage of 100 volts causes a current of 2 amperes to flow into the

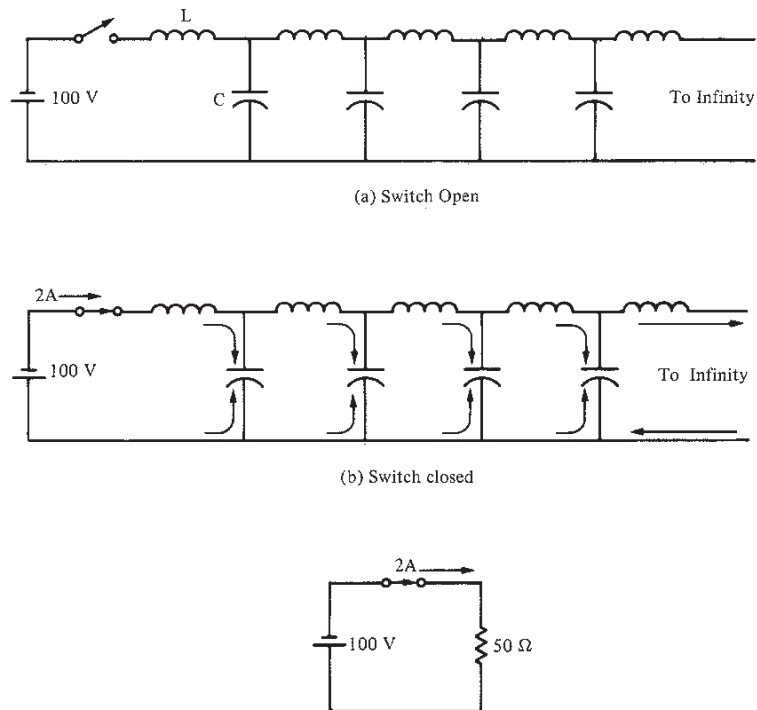


FIGURE 4.25 Characteristic Impedance

line when the switch is closed. As far as the source is concerned, it has no way of knowing whether it is connected to a transmission line that is infinitely long or to a 50-ohm resistor, as shown in Figure 4.25. In both cases, a current of 2 amperes would flow. For this reason, we say that this particular line has a characteristic impedance, or surge impedance, of 50 ohms.

The characteristic impedance is given by the equation:

$$Z_0 = \sqrt{\frac{L}{C}}$$

where

Z = characteristic impedance of the line, in ohms

L = inductance, in henrys per unit length

C = capacitance, in farads per unit length

Since the inductance and capacitance depend on the construction of the line, the characteristic impedance can also be determined from the physical dimensions of the line. In particular, for coaxial cable,

$$Z_0 = \frac{138}{\sqrt{\epsilon}} \log \frac{D}{d}$$

where

log = logarithm to the base 10

D = diameter of outside conductor

d = diameter of inside conductor

ϵ = dielectric constant of the insulating material between the two conductors; for air, the value is 1

For a dielectric of 1 and an impedance of 50 ohms, the ratio D/d is 2.3, and for an impedance of 75 ohms, the ratio is 3.5.

It is important to realize that the characteristic impedance of a transmission line is a function of the construction of the line itself; it does not depend on the signal carried or on what is connected to the line.

The significance of characteristic impedance is this: When a line is terminated in its characteristic impedance, any signal on the line is absorbed when it reaches the terminating resistance. There are no reflections. Obviously, such reflections are to be avoided since they would interfere with the signal being transmitted.

More detail on these matters can be found in any text on transmission line theory, for example, [LIBO85].

APPENDIX 4B: DECIBELS

An important parameter in any transmission system is the strength of the signal being transmitted. As a signal propagates along a transmission medium, there will be a loss, or attenuation, of signal strength. Additional losses occur at taps and splitters. To compensate, amplifiers may be inserted at various points to impart a gain in signal strength. It is customary to express gains, losses, and relative levels in decibels, because:

- Signal strength often falls off logarithmically, so loss is easily expressed in terms of the decibel, which is a logarithmic unit.
- The net gain or loss in cascaded transmission path can be calculated with simple addition and subtraction.

The decibel is a measure of the difference in two signal levels:

$$N_{\text{dB}} = 10 \log \frac{P_1}{P_2}$$

where

N_{dB} = number of decibels

$P_{1,2}$ = voltage values

For example, if a signal with a power level of 10 mw is inserted onto a transmission line and the measured power some distance away is 5 mw, the loss can be expressed as

$$\text{LOSS} = 10 \log(5/10) = 10(-.03) = -3 \text{ dB}$$

Note that the decibel is a measure of relative, not absolute, difference. A loss from 1000 mw to 500 mw is also a -3 dB loss. Thus, a loss of 3 dB halves the voltage level; a gain of 3 dB doubles the magnitude.

The decibel is also used to measure the difference in voltage, taking into account that power is proportional to the square of the voltage:

$$P = \frac{V^2}{R}$$

where

P = power dissipated across resistance R

V = voltage across resistance R

Thus

$$N_{\text{dB}} = 10 \log \frac{P_1}{P_2} = 10 \log \frac{V_1^2/R}{V_2^2/R} = 20 \log \frac{V_1}{V_2}$$

Decibel values refer to relative magnitudes or changes in magnitude, not to an absolute level. It is convenient to be able to refer to an absolute level of voltage in decibels so that gains and losses with reference to an initial signal level may easily be calculated. One unit in common use in cable television and broadband LAN applications is the dBmV (decibel-millivolt). This is an absolute unit with 0 dBmV equivalent to 1 mV. Thus

$$\text{Voltage(dBmV)} = 20 \log \frac{\text{Voltage(mV)}}{1\text{mV}}$$

The voltage levels are assumed to be across a 75-ohm resistance.

The decibel is convenient for determining overall gain or loss in a signal path. For example, Figure 4.9 shows a path from a point on a broadband trunk cable at which the signal level is 20 dBmV to an outlet. The amplifier gain and the losses due to the cables, tap, and splitter are expressed in decibels. By using simple addition and subtraction, the signal level at the outlet is easily calculated to be 7.6 dBmV.

APPENDIX 4C: SCRAMBLING AND DESCRAMBLING

For some digital data encoding techniques, a long string of binary zeros or ones in a transmission can degrade system performance. For example, in the differential phase-shift keying (DPSK) scheme used in some broadband LAN modems, a phase shift occurs only when the input is a zero bit. If there is a long string of ones, it is difficult for the receiver to maintain synchronization with the transmitter. A similar problem arises with the other common broadband LAN modulation scheme, duobinary AM/PSK. Also, other transmission properties are enhanced if the data are more nearly of a random nature rather than constant or repetitive [BELL82a]. A technique commonly used with modems to improve signal quality is scrambling and descrambling. The scrambling process tends to make the data appear more random.

The scrambling process consists of a feedback shift register, and the matching descrambler consists of a feedforward shift register. An example is shown in Figure 4.26. In this example, the scrambled data sequence may be expressed as follows:

$$B_m = A_m \oplus B_{m-3} \oplus B_{m-5}$$

where \oplus indicates the *exclusive or* operation. The descrambled sequence is

$$\begin{aligned} C_m &= B_m + B_{m-3} \oplus B_{m-5} \\ &= (A_m \oplus B_{m-3} \oplus B_{m-5}) \oplus B_{m-3} \oplus B_{m-5} \\ &= A_m \end{aligned}$$






FIGURE 4.26 Scrambler and Descrambler

As can be seen, the descrambled output is the original sequence.

We can represent this process with the use of polynomials. Thus, for this example, the polynomial is $P = 1 + X^{-3} + X^{-5}$. The input is divided by this polynomial to produce the scrambled sequence. At the receiver the received scrambled signal is multiplied by the same polynomial to reproduce the original input. Figure 4.27 is an example using the polynomial P and an input of 101010100000111. The scrambled transmission, produced by dividing by $P(100101)$, is 101110001101001. When this number is multiplied by P , we get the original input. Note that the input sequence contains the periodic sequence 10101010 as well as a long string of zeros. The scrambler effectively removes both patterns.



(a) Scrambling



(b) Descrambling

FIGURE 4.27 Example of Scrambling with $p(x) = 1 + x^{-3} + x^{-5}$

CHAPTER 5

Local Area Network Architecture

The preceding chapter examined some key issues relating to the architecture and physical properties of LANs. Because of its scope and importance, the subject of communications architecture or protocols was deferred and is presented here in its own chapter.

This chapter begins with an overall discussion of LAN protocols and seeks to determine what layers of functionality are required. Then the specific areas of link control and medium access control are explored. Throughout, reference is made to the IEEE 802 standard. This is for two reasons:

1. The standard is well thought out, providing a framework for exposing and clarifying LAN communication architectural issues.
2. The standard has had a major influence on LAN products.

A brief rationale and summary of the IEEE 802 standard is contained in an appendix to this chapter.

5.1

LAN PROTOCOLS

A LAN Reference Model

Chapter 2 summarized an architecture for communications, the OSI reference model, based on seven layers of protocols. We saw in that dis-

cussion (see Figure 2.14) that layers 1, 2, and 3 were required for the functioning of a packet-switching network. To recall, these layers were described as follows:

1. *Physical layer*: concerned with transmission of unstructured bit stream over physical link. Involves such parameters as signal voltage swing and bit duration. Deals with the mechanical, electrical, and procedural characteristics to establish, maintain, and deactivate the physical link.
2. *Data link layer*: provides for the reliable transfer of data across the physical link; sends blocks of data (frames) with the necessary synchronization, error control, and flow control.
3. *Network layer*: provides upper layers with independence from the data transmission and switching technologies used to connect systems; responsible for establishing, maintaining, and terminating connections.

We now turn to the question of what layers are required for the proper operation of the LAN. For the sake of clarity, we examine the question in the context of the OSI reference model. Two characteristics of LANs are important in this context. First, data are transmitted in addressed frames. Second, there is no intermediate switching, hence no routing required (repeaters are used in rings and may be used in baseband bus LANs, but do not involve switching or routing). One exception to the second characteristic is the ring bridge. A discussion of that and other exceptions is deferred until Chapter 10.

These two characteristics essentially determine the answer to the question: What OSI layers are needed? Layer 1, certainly. Physical connection is required. Layer 2 is also needed. Data transmitted across the LAN must be organized into frames and control must be exercised. But what about layer 3? The answer is yes and no. If we look at the functions performed by layer 3, the answer would seem to be no. First, there is routing. With a direct link available between any two points, this is not needed. The other functions—addressing, sequencing, flow control, error control, and so on—are, we learned, also performed by layer 2. The difference is that layer 2 performs these functions across a single link, whereas layer 3 may perform them across the sequence of links required to traverse the network. But since only one link is required to traverse the LAN, these layer 3 functions are redundant and superfluous!

From the point of view of an attached device, the answer would seem to be yes, the LAN must provide layer 3. The device sees itself attached to an access point into a network supporting communication with multiple devices. The layer for assuring that a message sent across that access point is delivered to one of a number of each points would seem to be a layer 3 function. So we can say that although the network provides services up through layer 3, the characteristics of the network allow

these services to be implemented on two OSI layers. We shall explore this topic more fully in Chapter 8. For the purpose of this chapter it is sufficient to understand that the minimum essential communications functions that must be performed by the LAN correspond to layers 1 and 2 of the OSI model.

With the points above in mind, let us now think about the functional requirements for controlling a local network and examine these from the top down. We follow the reasoning, illustrated in Figure 5.1, used by the IEEE 802 committee.

At the highest level are the functions associated with accepting transmissions from and delivering receptions to attached stations. These functions include:

- Provide one or more service access points. A service access point (SAP), recall, is a logical interface between two adjacent layers.
- On transmission, assemble data into a frame with address and CRC fields.
- On reception, disassemble frame, perform address recognition and CRC validation.
- Govern access to the link.

These are the functions typically associated with layer 2, the data link layer. The first function and related functions are grouped into a logical link control (LLC) layer by IEEE 802. The last three functions are treated as a separate layer, called *medium access control* (MAC). This is done for the following reasons:

- The logic required to manage access to a multiple-source, multiple-destination link is not found in traditional layer 2 link control.

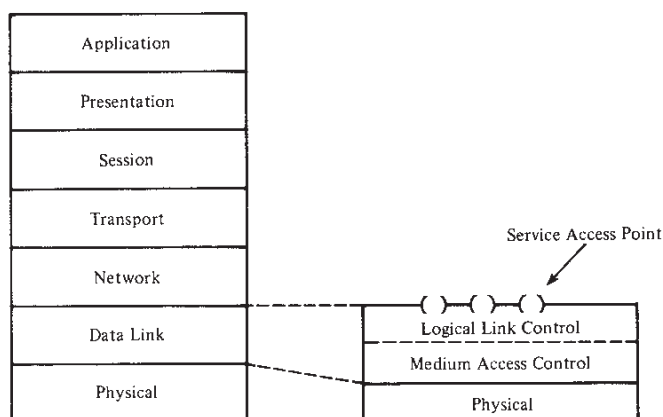


FIGURE 5.1 LAN Protocol Layers Compared to OSI

- For the same LLC, several MAC options may be provided, as we shall see.

Finally, at the lowest layer, are the functions generally associated with the physical layer. These include:

- Encoding/decoding of signals
- Preamble generation/removal (for synchronization)
- Bit transmission/reception

As with the OSI model, these functions are assigned to a physical layer in the IEEE 802 standard.

In the remainder of this section, we touch briefly on two aspects of LAN protocols. First, since the MAC layer is not found in the traditional OSI model, and to provide a context for later discussions, the characteristics and types of medium access control techniques are discussed. Then the structure for LAN frames is discussed briefly, using the IEEE 802 standard as an example.

We are then prepared to get more specific about LAN protocols. Section 5.2 discusses link control. Sections 5.3 and 5.4 provide details for various LAN medium access control techniques. Physical layer functions were discussed in Chapter 4.

Medium Access Control for Local Networks

All local networks (LAN, MAN, circuit-switched local network) consist of collections of devices that must share the network's transmission capacity. Some means of controlling access to the transmission medium is needed so that two particular devices can exchange data when required.

The key parameters in any medium access control technique are where and how. *Where* refers to whether control is exercised in a centralized or distributed fashion. In a centralized scheme, a controller is designated that has the authority to grant access to the network. A station wishing to transmit must wait until it receives permission from the controller. In a decentralized network, the stations collectively perform a medium access control function to dynamically determine the order in which stations transmit. A centralized scheme has certain advantages, such as:

- It may afford greater control over access for providing such things as priorities, overrides, and guaranteed bandwidth.
- It allows the logic at each station to be as simple as possible.
- It avoids problems of coordination.

Its principal disadvantages include:

- It results in a single point of failure.
- It may act as a bottleneck, reducing efficiency.

The pros and cons for distributed control are mirror images of the points made above.

The second parameter, *how*, is constrained by the topology and is a trade-off among competing factors: cost, performance, and complexity. In general, we can categorize access control techniques as being either synchronous or asynchronous. With synchronous techniques, a specific capacity is dedicated to a connection. We will see this in the circuit-switched local networks. Such techniques are not optimal in LANs and MANs because the needs of the stations are generally unpredictable. It is preferable to be able to allocate capacity in an asynchronous (dynamic) fashion, more or less in response to immediate needs. The asynchronous approach can be further subdivided into three categories: round robin, reservation, and contention.

Round Robin. Round robin techniques are conceptually simple, being based on the philosophy of "give everybody a turn." Each station in turn is given an opportunity to transmit. During that opportunity, the station may decline to transmit or may transmit subject to a certain upper bound, usually expressed as a maximum amount of data or time for this opportunity. In any case, the station, when it is finished, must relinquish its turn, and the right to transmit passes to the next station in logical sequence. Control of turns may be centralized or distributed. Polling on a multidrop line is an example of a centralized technique.

When many stations have data to transmit over an extended period of time, round robin techniques can be very efficient. If only a few stations have data to transmit at any given time, other techniques may be preferable, largely depending on whether the data traffic is stream or bursty. *Stream traffic* is characterized by lengthy and fairly continuous transmissions. Examples are voice communication, telemetry, and bulk file transfer. *Bursty traffic* is characterized by short, sporadic transmissions. Interactive terminal-host traffic fits this description.

Reservation. For stream traffic, reservation techniques are well suited. In general, for these techniques, time on the medium is divided into slots, much as with synchronous TDM. A station wishing to transmit reserves future slots for an extended or indefinite period. Again, reservations may be made in either a centralized or distributed fashion.

Contention. For bursty traffic, contention techniques are usually appropriate. With these techniques, no control is exercised to determine whose turn it is; all stations contend for time in a way that can be, as we shall see, rather rough and tumble. These techniques are of necessity distributed in nature. Their principal advantage is that they are simple to implement and, under light to to moderate load, efficient. For some

TABLE 5.1 Medium Access Control Techniques

	Centralized	Distributed
Round robin	Polling	Token bus Token ring Delay scheduling Implicit token
Reservation	Centralized reservation	Distributed reservation
Contention		CSMA/CD Slotted ring Register insertion

of these techniques, however, performance tends to collapse under heavy load.

Although both centralized and distributed reservation techniques have been implemented in some LAN products, round robin and contention techniques are the most common.

The discussion above has been somewhat abstract and should become clearer as specific techniques are discussed in this chapter and the next. For future reference, Table 5.1 places the techniques that will be discussed into the classification just outlined. Table 5.2 lists the MAC protocols that are defined in the LAN and MAN standards.

IEEE 802 Frame Format

This section presents the formats used for frames in the IEEE 802 standard. These formats are similar to those used by most proprietary networks. They are the basis for the LLC, MAC, and physical layer functionality.

At this point it is worth reviewing the HDLC format presented in Chapter 2. The requirements for a local network frame are very similar.

TABLE 5.2 Standardized Medium Access Control Techniques

	Bus Topology	Ring Topology
Round Robin	Token Bus (IEEE 802.4)	Token Ring (IEEE 802.5, FDDI)
Reservation	DQDB (IEEE 802.6)	FDDI-II
Contention	CSMA/CD (IEEE 802.3)	

Note: The DQDB and FDDI-II protocols for circuit-switched traffic are not fully specified in the standards and may be either distributed or centralized. All other standardized MAC protocols are distributed.

There must, of course, be a data or information field. A control field is needed to pass control bits and identify frame type. Starting and ending patterns are usually required to serve as delimiters. Addressing is required. Here is the main difference. Because LAN links are multiple-source, multiple-destination, both source and destination addresses are required. Further, unlike HDLC and virtually all other layer 2 protocols, the IEEE 802 LAN protocols support a form of multiplexing common in layer 3 protocols. As we shall see, this is accomplished in IEEE 802 by identifying service access points at each station.

Figure 5.2 shows the IEEE 802 formats. As can be seen, a separate format is used at the LLC level, and this is then embedded in the appropriate MAC frame. IEEE 802 supports three MAC alternatives: CSMA/CD, token bus, and token ring.

5.2

LINK LAYER PROTOCOL FOR LANs

In this section we look first at the general link level requirements for a local area network, then examine the IEEE 802 specification.

Principles

The link layers for LANs should bear some resemblance to the more common link layers extant. Like all link layers, the LAN link layer is concerned with the transmission of a frame of data between two stations, with no intermediate switching nodes.

It differs from traditional link layers in three ways:


1. It must support the multiaccess nature of the link (this differs from multidrop in that there is no primary node).
2. It is relieved of some details of link access by the MAC layer.
3. It must provide some layer 3 functions.

Figure 5.3 will help clarify the requirements for the link layer. We consider two stations or systems that communicate via a LAN (bus or ring). Higher layers (the equivalent of transport and above) provide end-to-end services between the stations. Below the link layer, a medium MAC layer provides the necessary logic for gaining access to the network for frame transmission and reception.


At a minimum, the link layer should perform those functions normally associated with that layer:

- *Error control*: End-to-end error control and acknowledgment. The link layer should guarantee error-free transmission across the LAN.
- *Flow control*: End-to-end flow control.


Logical Link Control (LLC)





CSMA/CD



Token Bus



Token Ring



- AC = Access Control
- DA = Destination Address
- DSAP = Destination Service Access Point
- ED = Ending Delimiter
- FC = Frame Control
- FCS = Frame Check Sequence
- FS = Frame Status
- SA = Source Address
- SD = Starting Delimiter
- SFD = Start Frame Delimiter
- SSAP = Source Service Access Point

FIGURE 5.2 IEEE 802 Frame Formats




FIGURE 5.3 LAN Communication Architecture

These functions can be provided in much the same way as for HDLC and other point-to-point link protocols—by the use of sequence numbers (N(S), N(R)).

It has already been mentioned that because of the lack of intermediate switching nodes, a LAN does not require a separate layer 3; rather, the essential layer 3 functions can be incorporated into layer 2:

- *Connectionless*: A service that does not require the overhead of establishing a logical connection is needed for efficient support of highly interactive traffic.
- *Connection-oriented*: A connection-oriented service is also usually needed.
- *Multiplexing*: Generally, a single physical link attaches a station to a LAN; it should be possible to provide data transfer with multiple end points over that link.

Because there is no need for routing, the above functions are easily provided. The connectionless service simply requires the use of source and destination address fields, as discussed previously. The station sending the frame must designate the destination address, so that the frame is delivered properly. The source address must also be indicated so that the recipient knows where the frame came from.

Both the connection-oriented and multiplexing capabilities can be supported with the concept of the service access point (SAP), introduced in Chapter 2. An example may make this clear. Figure 5.4 shows three stations attached to a LAN. Each station has an address. Further, the link layer supports multiple SAPs, each with its own address. The link layer provides communication between SAPs. Assume that a process or application X in station A wishes to send a message to a process in station C. X may be a report generator program in minicomputer A. C may be a printer and a simple printer driver. X attaches itself to SAP 1 and requests a connection to station C, SAP 1 (station C may have only one SAP if it is a single printer). Station A's link layer then sends to the LAN a connection-request frame that includes the source address (A, 1), the destination address (C, 1), and some control bits indicating that this is a connection request. The LAN delivers this frame to C, which, if it is free, returns a connection-accepted frame. Henceforth, all data from X will be assembled into a frame by A's LLC, which includes source (A,1) and destination (C,1) addresses. Incoming frames addressed to (A,1) will be rejected unless they are from (C, 1); these might be acknowledgment frames, for example. Similarly, station C's printer is declared busy and C will accept frames only from (A,1).

Thus a connection-oriented service is provided. At the same time, process Y could attach to (A,2) and exchange data with (B,1). This is an example of multiplexing. In addition, various other processes in A could use (A,3) to send datagrams to various destinations.

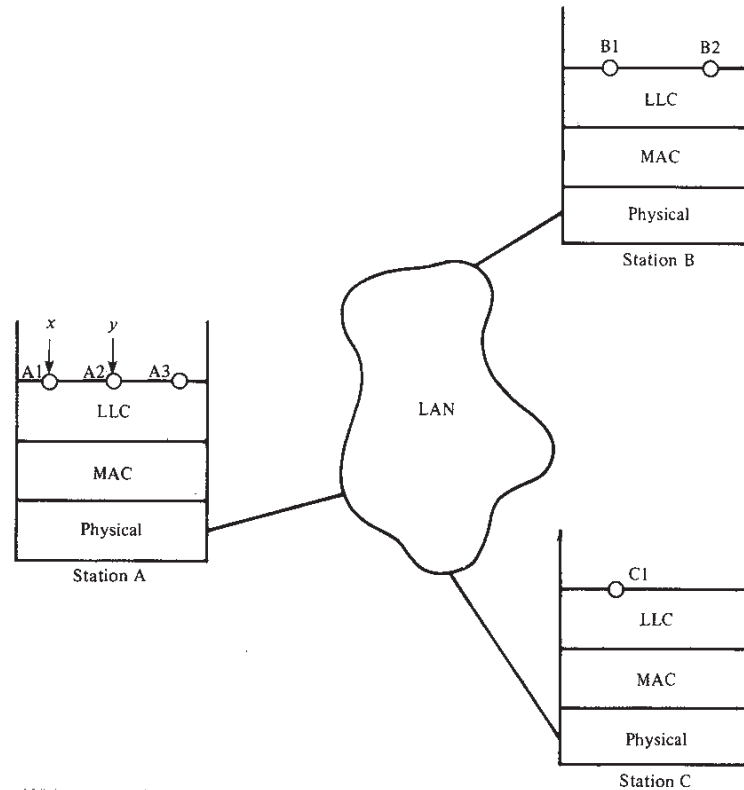


FIGURE 5.4 LAN Link Control Scenario

One final function of the link layer should be included, to take advantage of the multiple access nature of the LAN:

- *Multicast, broadcast:* The link layer should provide a service of sending a message to multiple stations or all stations.

Addressing

The preceding discussion referred to both station and LLC addresses. A further elaboration of this point is warranted. To understand the function of addressing, we need to consider the requirements for exchanging data.

In very general terms, communication can be said to involve three agents: processes, stations, and networks. *Processes* are the fundamental entities that communicate. One example is a file transfer operation. In

this case, a file transfer process in one station exchanges data with a file transfer process in another station. Another example is remote terminal access. In this case, a user terminal is attached to one station and controlled by a terminal-handling process in that station. The user, through the terminal-handling process, is remotely connected to a time-sharing system; data are exchanged between the terminal-handling process and the time-sharing process. Processes execute on *stations*, which can often support multiple simultaneous processes. Stations are connected by a *network*, and the data to be exchanged are transmitted by the network from one station to another. From this point of view, the transfer of data from one process to another involves first getting the data to the station in which the process resides and then getting the data to the process within the station.

These concepts suggest the need for two levels of addressing. To see this, consider Figure 5.5, which shows the overall format of data transmitted using the LLC and MAC protocols (compare Figure 2.18). User data to be sent are passed down to LLC, which appends a header. This header contains control information that is used to manage the protocol between the local LLC entity and the remote LLC entity. The combination of user data and LLC header is referred to as an LLC *protocol data unit* (PDU). After the sending LLC has prepared a PDU, the PDU is then passed as a block of data down to the MAC entity. The MAC entity appends both a header and a trailer, to manage the MAC protocol. The result is a MAC-level PDU. To avoid confusion with an LLC-level PDU, the MAC-level PDU is typically referred to as a *frame*.




FIGURE 5.5 LAN Protocol Data Units




FIGURE 5.6 LAN Addressing

Now, the MAC header must contain a destination address that uniquely identifies a station on the local network. This is needed since each station on the local network will read the destination address field to determine if it should capture the MAC frame. When a MAC frame is captured, the MAC entity strips off the MAC header and trailer and passes the resulting LLC PDU up to the LLC entity. The LLC header must contain a destination SAP address so the LLC can determine to whom the data are to be delivered. Hence, two levels of addressing are needed:

1. *MAC address*: identifies a station on the local network
2. *LLC address*: identifies an LLC user

Figure 5.6 illustrates the two levels of addresses. The MAC address is associated with a physical attachment point on the network. The LLC SAP is associated with a particular user within a station. In some cases, the SAP corresponds to a host process. Another case relates to a common type of attached equipment, referred to as a network interface unit (NIU).¹ Often, an NIU is used as a terminal concentration device. In this case, each terminal port on the NIU has a unique SAP.

So far, we have discussed the use of addresses that identify unique entities. In addition to these **individual addresses**, group addresses are also employed. A **group address** specifies a collection of one or more entities. For example, one might wish to send a message to all terminal users attached to a particular NIU, or all terminal users on the entire LAN. Two types of group addresses are used. A **broadcast address** refers to all entities within some context; this is also referred to as an all-stations address. A **multicast address** refers to some subset of entities within some context.

Table 5.3 depicts the possible combinations. The first five combinations are straightforward. A specific user can be addressed. A group of

¹NIUs are examined in detail in Chapter 8.

TABLE 5.3 LAN Addressing

MAC Address	LLC User Address (Service Access Point)
Individual	Individual
Individual	Multicast
Individual	Broadcast
Multicast	Broadcast
Broadcast	Broadcast
Multicast	Individual
Multicast	Multicast
Broadcast	Individual
Broadcast	Multicast

users or all users at a specific station can be addressed. And all users on some stations or all users on all stations can be addressed.

The last four combinations in the table are less obvious. It should be clear that LLC addresses are unique only within a single station. It is only the LLC entity within a station that examines the LLC header and determines the user. However, it is possible to assign LLC addresses uniquely across all stations; this is undesirable for the following reasons:

- The total number of users on all stations would be limited by the SAP field length in the LLC header.
- Central management of SAP assignment would be required, no matter how large and heterogeneous the user population.

On the other hand, it may be desirable to assign the same SAP value to entities in different stations. For example, a station management entity in a station may always be given an SAP value of 1, to facilitate network management. Or a group of management and control entities within a station may always be given the same multicast SAP address. When such a convention is followed, then it becomes possible to address data to one SAP address or a multicast SAP address in a group of stations or all stations.

IEEE 802 Logical Link Control

The IEEE 802 LLC standard [IEEE89a] is a good example of a LAN link control layer. It is well thought out and offers a variety of services. This section summarizes the features of LLC.

Figure 5.2 depicts the LLC frame. As can be seen, it specifies the source and destination service access points (thus allowing link multi-

plexing), a 1- or 2-byte control field, and a data field. The source and destination address fields are needed by the LLC, but are also used by MAC, and are included in the outer MAC frame. The LLC can be specified in three parts:

1. The interface with the station, specifying the services that LLC (and hence the LAN) provides to the network subscriber
2. The LLC protocol, specifying the LLC functions
3. The interface with MAC, specifying the services that LLC requires to perform its function

A variety of functions were mentioned in the previous section. Not all of these functions are needed in all environments. Accordingly, the 802 standard defines two general categories of data link control operation. The first is a connectionless operation that provides minimum service with minimum protocol complexity. This is useful and efficient when higher layers (e.g., network, transport) provide error control, flow control, and sequencing functions. It is also useful when the guaranteed delivery of data is not required. The second category is connection-oriented operation that provides the functions referred to above using a protocol similar to HDLC. These two types of operations are reflected in the specifications of both the LLC services and the LLC protocol.

LLC Services. LLC provides three services:

1. *Unacknowledged connectionless service:* This is a datagram service that simply allows for sending and receiving frames. It supports point-to-point, multipoint, and broadcast.
2. *Connection-oriented service:* This provides a logical connection between service access points. It provides flow control, sequencing, and error recovery.
3. *Acknowledged connectionless service:* This is also a connectionless service, but provides for acknowledgment, relieving higher layers of this burden. It supports point-to-point transfers.

These services are specified in terms of primitives that can be viewed as commands or procedure calls with parameters.² Table 5.4 summarizes the LLC primitives.

The **Unacknowledged Connectionless Service** is a datagram style of service that simply allows for sending and receiving LLC frames, with no form of acknowledgment to assure delivery. It supports point-to-point, multipoint, and broadcast addressing.

²These primitives always include one of four standard modifiers: request, indication, response, confirm. The interpretation of these primitives is discussed in Appendix 5B at the end of this chapter.

TABLE 5.4 Logical Link Control Primitives

UNACKNOWLEDGED CONNECTIONLESS SERVICE	
DL-UNITDATA.request	(source-address, destination-address, data, priority)
DL-UNITDATA.indication	(source-address, destination-address, data, priority)
CONNECTION-MODE SERVICE	
DL-CONNECT.request	(source-address, destination-address, priority)
DL-CONNECT.indication	(source-address, destination-address, priority)
DL-CONNECT.response	(source-address, destination-address, priority)
DL-CONNECT.confirm	(source-address, destination-address, priority)
DL-DATA.request	(source-address, destination-address, data)
DL-DATA.indication	(source-address, destination-address, data)
DL-DISCONNECT.request	(source-address, destination-address)
DL-DISCONNECT.indication	(source-address, destination-address, reason)
DL-RESET.request	(source-address, destination-address)
DL-RESET.indication	(source-address, destination-address, reason)
DL-RESET.response	(source-address, destination-address)
DL-RESET.confirm	(source-address, destination-address)
DL-CONNECTION-FLOWCONTROL.request	(source-address, destination-address, amount)
DL-CONNECTION-FLOWCONTROL.indication	(source-address, destination-address, amount)
ACKNOWLEDGED CONNECTIONLESS SERVICE	
DL-DATA-ACK.request	(source-address, destination-address, data, priority, service-class)
DL-DATA-ACK.indication	(source-address, destination-address, data, priority, service-class)
DL-DATA-ACK-STATUS.indication	(source-address, destination-address, priority, service-class, status)
DL-REPLY.request	(source-address, destination-address, data, priority, service-class)
DL-REPLY.indication	(source-address, destination address, data, priority, service-class)
DL-REPLY-STATUS.indication	(source-address, destination-address, data, priority, service-class, status)
DL-REPLY-UPDATE.request	(source-address, data)
DL-REPLY-UPDATE-STATUS.indication	(source-address, status)

This service provides for only two primitives across the interface between the next higher layer and LLC. DL-UNITDATA.request is used to pass a block of data down to LLC for transmission. DL-UNITDATA.indication is used to pass that block of data up to the destination user from LLC upon reception. The source-address and destination-address parameters specify the local and remote LLC users, respectively. Each of these parameters actually is a combination of LLC service access point and the MAC address. The data parameter is the block of data transmitted from one LLC user to another. The priority parameter spec-

ifies the desired priority. This (together with the MAC portion of the address) is passed down through the LLC entity to the MAC entity, which has the responsibility of implementing a priority mechanism. As we shall see, token bus and token ring are capable of this, but the 802.3 CSMA/CD system is not.

The **Connection-Oriented Service** provides a virtual-circuit style connection between service access points (between users). It provides a means by which a user can request or be notified of the establishment or termination of a logical connection. It also provides flow control, sequencing, and error recovery. It supports point-to-point addressing.

This service includes the DL-CONNECT set of primitives (request, indication, response, confirm) to establish a logical connection between SAPs. Once the connection is established, blocks of data are exchanged using DL-DATA.request and DL-DATA.indication. Because the existence of a logical connection guarantees that all blocks of data will be delivered reliably, there is no need for an acknowledgment (via indication and confirm primitives) of individual blocks of data. At any point, either side may terminate the connection with a DL-DISCONNECT.request; the other side is informed with a DL-DISCONNECT.indication.

The DL-RESET primitives are used to reset a logical connection to an initial state. Sequence numbers are reset and the connection is reinitialized. Finally, the two flow control primitives regulate the flow of data across the SAP. The flow can be controlled in either direction. This is a local flow control mechanism that specifies the amount of data that may be passed across the SAP.

The **Acknowledged Connectionless Service** provides a mechanism by which a user can send a unit of data and receive an acknowledgment that the data were delivered, without the necessity of setting up a connection.

This service includes DL-DATA-ACK.request and DL-DATA-ACK.indication with meanings analogous to those for the Unacknowledged Connectionless Service, plus DL-DATA-ACK-STATUS.indication to provide acknowledgment to the sending user. The DL-REPLY primitives provide a data exchange service. It allows a user to request that data be returned from a remote station or that data units be exchanged with a remote station. Associated with these primitives are the DL-REPLY-UPDATE primitives. These primitives allow a user to pass data to LLC to be held and sent out at a later time when requested to do so (by a DL-REPLY primitive) by some other station.

The specification of three types of service is intended to allow LLC to be used to support a variety of user requirements and to enable implementors to implement subsets of LLC to meet their specific needs and to optimize the implementation to those needs. The Unacknowledged Connectionless Service is the simplest and requires the minimum im-

plementation. In cases where higher layer protocols (usually transport) provide end-to-end error control and flow control, this minimum service is all that is needed. On the other hand, when the supported devices are very simple (e.g., terminals), it might make sense to forgo elaborate upper layers and rely on LLC to provide end-to-end control. Finally, the Acknowledged Connectionless Service may be useful in some real-time environments, such as factory LANs. For example, certain alarm or control signals may be very important and time-critical. Because of their importance, an acknowledgment is needed so that the sender can be assured that the signal got through. Because of the urgency of a signal, the user might not want to take the time to first establish a logical connection and then send the data.

LLC Protocol

The basic LLC protocol is modeled after the HDLC balanced mode, and it has similar formats and functions. These are summarized briefly in this section. The reader should be able to see how this protocol supports the LLC services defined above.

The format of an LLC protocol data unit is shown in Figure 5.2. First are the address fields. Both the DSAP and SSAP fields actually contain 7-bit addresses. The least significant bit of DSAP indicates whether this is an individual or group address. The least significant bit of SSAP indicates whether this is a command or response frame.

Figure 5.7 shows the format for the LLC control field (compare Figure 2.19). It is identical to that of HDLC and the functioning is the same, with four exceptions:

1. LLC makes use of only the asynchronous balanced mode of operation and does not employ HDLC's normal response mode or asynchronous response mode. This mode is used to support connection-oriented service. The set asynchronous balanced mode (SABME) command is used to establish a connection, and disconnect (DISC) is used to terminate the connection.
2. LLC supports a connectionless (datagram) service by using the unnumbered information (UI) frame.
3. LLC permits multiplexing by the use of SAPs.
4. LLC supports an acknowledged connectionless service by using two new unnumbered frames.

A brief summary follows.

As with HDLC, three frame formats are defined for LLC: information transfer, supervisory, and unnumbered. Their use depends on the type of operation employed. The types are Type 1 (connectionless), Type 2 (connection-oriented), and Type 3 (acknowledged connectionless).

	1	2	3	4	5	6	7	8	9	10-16
Information Transfer Command/Response (I-Format PDU)	0	N(S)							P/F	N(R)
Supervisory Commands/Responses (S-Format PDUs)	1	0	S	S	X	X	X	X	P/F	N(R)
Unnumbered Commands/Response (U-Format PDUs)	1	1	M	M	P/F	M	M	M		

Where

N(S)-Transmitter Send Sequence Number (Bit 2-Low-order Bit)

N(R)-Transmitter Receive Sequence Number (Bit 10-Low-order Bit)

S-Supervisory Function Bit

M-Modifier Function Bit

X-Reserved and Set to Zero

P/F-Poll Bit-Command LLC PDU Transmissions

Final Bit-Response LLC PDU Transmissions:
(1-Poll/Final)

FIGURE 5.7 IEEE 802 LLC Control Field Format

With Type 1 Operation, protocol data units (PDUs) are exchanged between LLC entities without the need to establish a logical connection. There is no acknowledgment, flow control, or error control. This type of operation supports the Unacknowledged Connectionless Service.

Three unnumbered frame formats are used. The UI (unnumbered information) frame is used to send a connectionless data frame, containing data from an LLC user. The XID (exchange identification) frame is used to convey station class (which operation types are supported). The TEST (test) frame is used to request a TEST frame in response, to test the LLC-to-LLC path.

With Type 2 Operation, a data link connection is established between two LLC entities prior to data exchange. This type of operation supports Connection-Oriented Service and uses all three frame formats. The information transfer frames are used to send data (as opposed to control information). N(S) and N(R) are frame sequence numbers that support error control and flow control. A station sending a sequence of frames will number them, modulo 128, and place the number in N(S). N(R) is a piggybacked acknowledgment. It enables the sending station to indicate which number frame it expects to receive next. These numbers support flow control since, after sending seven frames without an acknowledgment, a station can send no more. The numbers support error control, as explained below. The P/F field is set to 1 only on the last frame in a series, to indicate that the transmission is over.

The supervisory frame is used for acknowledgment and flow control. The 2-bit SS field is used to indicate one of three commands: Receive

Ready (RR), Receive Not Ready (RNR), and Reject (REJ). RR is used to acknowledge the last frame received by indicating in N(R) the next frame expected. The frame is used when there is no reverse traffic to carry to piggybacked acknowledgment. RNR acknowledges a frame, as with RR, but also asks the transmitting station to suspend transmission. When the receiving station is again ready it sends an RR frame. REJ is used to indicate that the frame with number N(R) is rejected and that it and any subsequently transmitted frames must be sent again.

Unnumbered frames are used for control purposes in Type 2 operation. The 5-bit MMMMM field specifies a particular command or response. The commands are:

- SABME (set asynchronous balanced mode extended): used by an LLC entity to request logical connection with another LLC entity.
- DISC (disconnect): used to terminate a logical connection; the sending station is announcing that it is suspending operations.

The foregoing frames are commands, initiated by a station at will. The following frames are responses:

- UA (unnumbered acknowledgment): used to acknowledge SABME and DISC commands
- DM (disconnected mode): used to respond to a frame in order to indicate that the station's LLC is logically disconnected
- FRMR (frame reject): used to indicate that an improper frame has arrived—one that somehow violates the protocol

The P/F bit is used to indicate that a response is requested to a command frame.

With Type 3 Operation, each transmitted frame is acknowledged. A new unnumbered frame, the Acknowledged Connectionless (AC) Information frame, is defined. Unlike the other frames used in LLC, this frame is not defined in HDLC. User data are sent in an AC command frame and must be acknowledged using an AC response frame. To guard against lost frames, a 1-bit sequence number is used. The sender alternates the use of 0 and 1 in its AC command frames, and the receiver responds with an AC frame with the corresponding number.

LLC-MAC Interface. The IEEE 802 LLC is intended to operate with any of the three MAC protocols (CSMA/CD, token bus, token ring). A single logical interface to any of the MAC layers is defined. The 802 standard does not define an explicit interface, but provides a model. The basic primitives are:

- MA-UNITDATA.request: to request transfer of an LLC frame from local LLC to destination LLC. This includes information transfer, supervisory, and unnumbered frames.

- MA-UNITDATA.indicate: to transfer incoming LLC frame from local MAC to local LLC.

5.3

MEDIUM ACCESS CONTROL—BUS/TREE

Of all the local network topologies, the bus/tree topologies present the greatest challenges and the most options for medium access control. This section will not attempt to survey the many techniques that have been proposed; good discussions can be found in [LUCZ78] and [FRAN81]. Rather, emphasis is placed on the two techniques that seem likely to dominate the marketplace: CSMA/CD and token bus. Standards for these techniques have been developed by the IEEE 802 committee.

A third technique, centralized reservation, is reviewed briefly. This is for the sake of completeness; virtually all access techniques for bus/tree are related to one of these three techniques.

Table 5.5 compares the three techniques on a number of characteristics. The ensuing discussion should clarify their significance.

CSMA/CD

The most commonly used medium access control technique for bus-tree topologies is carrier sense multiple access with collision detection (CSMA/CD). The original baseband version of this technique was developed and patented by Xerox [METC77] as part of its Ethernet local network [METC76]. The original broadband version was developed and patented by MITRE [HOPK80] as part of its MITREnet local network [HOPK79, HOPK77]. A baseband version inspired by Ethernet has been issued as an IEEE 802 standard [IEEE90b].

TABLE 5.5 Bus/Tree Access Methods

	CSMA/CD	Token Bus	Centralized Reservation
Access determination	Contention	Token	Reservation
Packet length restriction	Greater than twofold propagation delay	None	No greater than slot size
Principal advantage	Simplicity	Regulated/fair access	Regulated/fair access
Principal disadvantage	Performance under heavy load	Complexity	Required central controller

Before examining this technique, we look at some earlier schemes from which CSMA/CD evolved.

Precursors. All of the techniques discussed in this section, including CSMA/CD, can be termed *random access* or *contention* techniques. They are designed to address the problem of how to share a common broadcast transmission medium—the “Who goes next?” problem. The techniques are random access in the sense that there is no predictable or scheduled time for any station to transmit; station transmissions occur randomly. They are contention in the sense that no control is exercised to determine whose turn it is—all stations must contend for time on the network.

The earliest of these techniques, known as *ALOHA*, was developed for ground-based packet radio broadcasting networks [ABRA70]. However, it is applicable to any transmission medium shared by uncoordinated users. ALOHA, or *pure ALOHA* as it is sometimes called, is a true free-for-all. Whenever a station has a frame to send, it does so. The station then listens for an amount of time equal to the maximum possible round-trip propagation time on the network (twice the time it takes to send a frame between the two most widely separated stations). If the station hears an acknowledgment during that time, fine; otherwise, it resends the frame. After repeated failures, it gives up. A receiving station determines the correctness of an incoming frame by examining the check sum. If the frame is valid, the station acknowledges immediately. The frame may be invalid, due to noise on the channel or because another station transmitted a frame at about the same time. In the latter case, the two frames may interfere with each other so that neither gets through; this is known as a *collision*. In that case, the receiving station simply ignores the frame. ALOHA is as simple as can be, and pays a penalty for it. Because the number of collisions rises so rapidly with increased load, the maximum utilization of the channel is only about 18%.

To improve efficiency, a modification of ALOHA [ROBE75] was developed in which time on the channel is organized into uniform slots whose size equals the frame transmission time. Some central clock or other technique is needed to synchronize all stations. Transmission is permitted to begin only at a slot boundary. Thus frames that do overlap will do so totally. This increases the maximum utilization of the system to about 37%. The scheme is known as *slotted ALOHA*.

Both ALOHA and slotted ALOHA exhibit poor utilization. Both fail to take advantage of one of the key properties of both packet radio and local networks, which is that the propagation delay between stations is usually very small compared to frame transmission time. Consider the following observations. If the station-to-station propagation time is large

compared to the frame transmission time, then, after a station launches a frame, it will be a long time before other stations know about it. During that time, one of the other stations may transmit a frame; the two frames may interfere with each other and neither gets through. Indeed, if the distances are great enough, many stations may begin transmitting, one after the other, and none of their frames get through unscathed. Suppose, however, that the propagation time is extremely small compared to frame transmission time. In that case, when a station launches a frame, all the other stations know it almost immediately. So, if they had any sense, they would not try transmitting until the first station was done. Collisions would be rare since they would occur only when two stations began to transmit almost simultaneously. Another way to look at it is that the short delay time provides the stations with better feedback about the state of the system; this information can be used to improve efficiency.

The foregoing observations led to the development of a technique known as carrier sense multiple access (CSMA) or listen before talk (LBT). A station wishing to transmit first listens to the medium to determine if another transmission is in progress. If the medium is in use, the station backs off some period of time and tries again, using one of the algorithms explained below. If the medium is idle, the station may transmit. Now, it may happen that two or more stations attempt to transmit at about the same time. If this happens, there will be a collision. To account for this, a station waits a reasonable amount of time after transmitting for an acknowledgment, taking into account the maximum round-trip propagation delay, and the fact that the acknowledging station must also contend for the channel in order to respond. If there is no acknowledgment, the station assumes that a collision has occurred and retransmits.

One can see how this strategy would be effective for systems in which the frame transmission time is much longer than the propagation time. Collisions can occur only when more than one user begins transmitting within a short time (within the period of propagation delay). If a station begins to transmit, and there are no collisions during the time it takes for the leading edge of the frame to propagate to the farthest station, then the station has seized the channel and the remainder of the frame will be transmitted without collision.

The maximum utilization achievable using CSMA can far exceed that of ALOHA or slotted ALOHA. The maximum utilization depends on the length of the frame and on the propagation time; the longer the frames or the shorter the propagation time, the higher the utilization. This subject will be explored in Chapter 9.

With CSMA, an algorithm is needed to specify what a station should do if the medium is found to be busy. Three approaches are depicted in

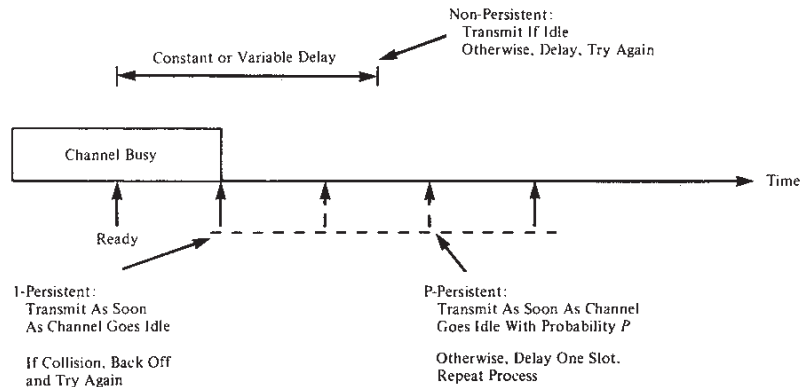


FIGURE 5.8 CSMA Persistence and Back-off

Figure 5.8. One algorithm is *nonpersistent* CSMA. A station wishing to transmit listens to the medium and obeys the following rules:

1. If the medium is idle, transmit.
2. If the medium is busy, wait an amount of time drawn from a probability distribution (the retransmission delay) and repeat step 1.

The use of random retransmission times reduces the probability of collisions. The drawback is that even if several stations have a frame to send, there is likely to be some wasted idle time following a prior transmission.

To avoid channel idle time, the *1-persistent protocol* can be used. A station wishing to transmit listens to the medium and obeys the following rules:

1. If the medium is idle, transmit.
2. If the medium is busy, continue to listen until the channel is sensed idle; then transmit immediately.
3. If there is a collision (determined by a lack of acknowledgment), wait a random amount of time and repeat step 1.

Whereas nonpersistent stations are deferential, 1-persistent stations are selfish. If two or more stations are waiting to transmit, a collision is guaranteed. Things get sorted out only after the collision.

A compromise that attempts to reduce collisions, like nonpersistent, and reduce idle time, like 1-persistent, is *p-persistent*. The rules are:

1. If the medium is idle, transmit with probability p , and delay one time unit with probability $(1 - p)$. The time unit is typically equal to the maximum propagation delay.

2. If the medium is busy, continue to listen until the channel is idle and repeat step 1.
3. If transmission is delayed one time unit, repeat step 1.

The question arises as to what is an effective value of p . The main problem to avoid is one of instability under heavy load. Consider the case in which n stations have frames to send while a transmission is taking place. At the end of that transmission, the expected number of stations that will attempt to transmit is np . If np is greater than 1, multiple stations will attempt to transmit and there will be a collision. What is more, as soon as all these stations realize that they did not get through, they will be back again, almost guaranteeing more collisions. Worse yet, these retries will compete with new transmissions from other stations, further increasing the probability of collision. Eventually, all stations will be trying to send, causing continuous collisions, with throughput dropping to zero. To avoid this catastrophe np must be less than one for the expected peaks of n . As p is made smaller, stations must wait longer to attempt transmission but collisions are reduced. At low loads, however, stations have unnecessarily long delays.

Description of CSMA/CD. All of the techniques described above could be used in a bus/tree topology with an electrical conductor medium or in a packet radio scheme. We now introduce *carrier sense multiple access with collision detection* (CSMA/CD), which, because of the CD part, is appropriate only for a bus/tree topology [it is also referred to as *listen while talk* (LWT)]. CSMA/CD can be used with either baseband or broadband systems. Where details differ between baseband and broadband, we will use IEEE 802 and MITREnet as examples for comparison.

CSMA, although more efficient than ALOHA or slotted ALOHA, still has one glaring inefficiency. When two frames collide, the medium remains unusable for the duration of transmission of both damaged frames. For long frames, compared to propagation time, the amount of wasted bandwidth can be considerable. This waste can be reduced if a station continues to listen to the medium while it is transmitting. In that case, these rules can be added to the CSMA rules:

1. If a collision is detected during transmission, immediately cease transmitting the frame, and transmit a brief jamming signal to assure that all stations know that there has been a collision.
2. After transmitting the jamming signal, wait a random amount of time, then attempt to transmit again using CSMA.

Now the amount of wasted bandwidth is reduced to the time it takes to detect a collision. Question: How long does that take? Figure 5.9 illustrates the answer for a baseband system. Consider the worst case of two stations that are as far apart as possible. As can be seen, the amount

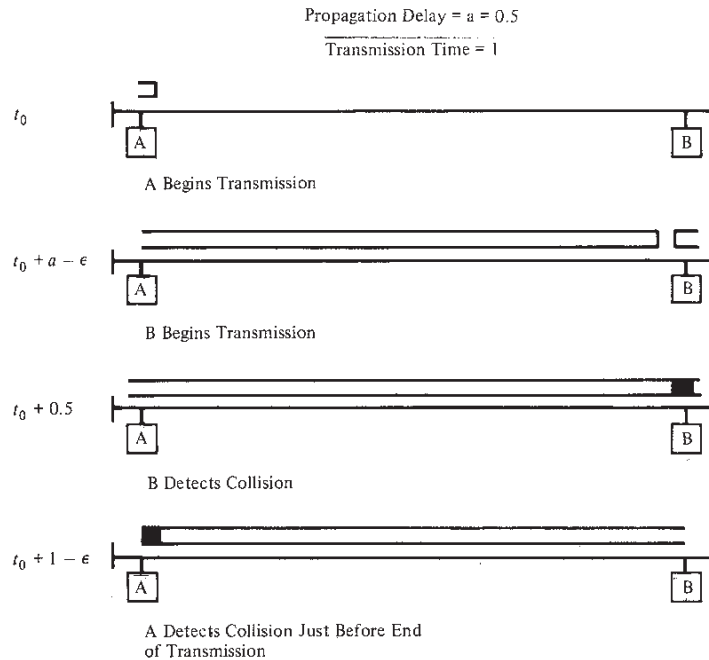


FIGURE 5.9 Baseband Collision Detection Timing

of time it takes to detect a collision is twice the propagation delay. For broadband bus, the wait is even longer. Figure 5.10 shows a dual-cable system. This time, the worst case is two stations close together and as far as possible from the headend. In this case, the time required to detect a collision is four times the propagation delay from the station to the headend. The results would be the same for a midsplit system.

Both figures indicate the use of frames long enough to allow CD prior to the end of transmission. In most systems that use CSMA/CD, it is required that all frames be at least this long. Otherwise, the performance of the system is the same as the less efficient CSMA protocol, since collisions are detected only after transmission is complete.

Now let us look at a few details of CSMA/CD. First, which persistence algorithm should we use: non-, 1-, or p-? You may be surprised to learn that the most common choice is 1-persistent. It is used by both Ethernet and MITREnet, and in the IEEE 802 standard. Recall that both nonpersistent and p-persistent have performance problems. In the nonpersistent case, capacity is wasted because the medium will generally remain idle following the end of a transmission even if there are stations waiting to send. In the p-persistent case, p must be set low enough to avoid




FIGURE 5.10 Broadband Collision Detection Timing

instability, with the result of sometimes atrocious delays under light load. The 1-persistent algorithm, which after all means $p = 1$, would seem to be even more unstable than p-persistent due to the greed of the stations. What saves the day is that the wasted time due to collisions is mercifully short (if the frames are long relative to propagation delay!), and with random back-off, the two stations involved in a collision are unlikely to collide on their next tries. To ensure that back-off maintains stability, IEEE 802 and Ethernet use a technique known as binary exponential back-off. A station will attempt to transmit repeatedly in the face of repeated collisions, but after each collision, the mean value of the random delay is doubled. After 16 unsuccessful attempts, the station gives up and reports an error.

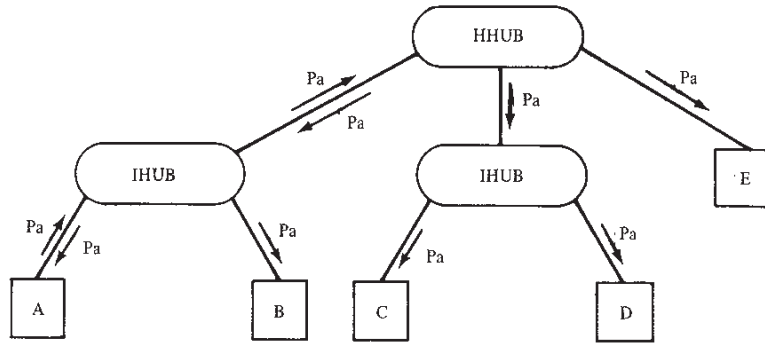
The beauty of the 1-persistent algorithm with binary exponential back-off is that it is efficient over a wide range of loads. At low loads, 1-

persistence guarantees that a station can seize the channel as soon as it goes idle, in contrast to the non- and p-persistent schemes. At high loads, it is at least as stable as the other techniques. However, one unfortunate effect of the back-off algorithm is that it has a last-in, first-out effect; stations with no or few collisions will have a chance to transmit before stations that have waited longer.


Although the implementation of CSMA/CD is substantially the same for baseband and broadband, there are differences. One example is the means for performing carrier sense. For baseband systems using Manchester encoding, carrier is conveniently sensed by detecting the presence of transitions on the channel. Strictly speaking, there is no carrier to sense digital signaling; the term was borrowed from the radio lexicon. With broadband, carrier sense is indeed performed. The station's receiver listens for the presence of a carrier on the outbound channel.

Collision detection also differs for the two systems. In a baseband system, a collision should produce substantially higher voltage swings than those produced by a single transmitter. Accordingly, Ethernet and the IEEE standard dictate that a transmitting transceiver will detect a collision if the signal on the cable at the transceiver exceeds the maximum that could be produced by the transceiver alone. Because a transmitted signal attenuates as it propagates, there is a potential problem with collision detection. If two stations far apart are transmitting, each station will receive a greatly attenuated signal from the other. The signal strength could be so small that when it is added to the transmitted signal at the transceiver, the combined signal does not exceed the CD threshold. For this reason, among others, IEEE 802 restricts the maximum length of cable to 500 m. Because frames may cross repeater boundaries, collisions must cross as well. Hence if a repeater detects a collision on either cable, it must transmit a jamming signal on the other side. Since the collision may not involve a transmission from the repeater, the CD threshold is different for a nontransmitting transceiver: a collision is detected if the signal strength exceeds that which could be produced by two transceiver outputs in the worst case.


A much simpler collision detection scheme is possible with the twisted-pair star-wiring approach (Figure 4.10). In this case, collision detection is based on logic rather than sensing voltage magnitudes. For any hub, if there is activity (signal) on more than one input, a collision is assumed. A special signal called the *collision presence* signal is generated. This signal is generated and sent out as long as activity is sensed on any of the input lines. This signal is interpreted by every node as an occurrence of collision. Figure 5.11 gives examples of the operation of a star-wired system with and without collisions. In the first example, a frame transmitted from station A propagates up to HHUB and is eventually received by all stations in the network. In the second example, a



(a) A Transmitting



(b) A and B Transmitting



(c) A, B, and C Transmitting

P_x = PDU from Station x
 CP = Collision Presence Signal

FIGURE 5.11 Operation of a Two-Level Star-Wired CSMA/CD Configuration

collision is detected by A's IHUB. The collision presence signal propagates up to HHUB and is rebroadcast down to all hubs and stations. The third example shows the result of a three-way collision.

There are several possible approaches to collision detection in broadband systems. The most common of these is to perform a bit-by-bit comparison between transmitted and received data. When a station transmits on the inbound channel, it begins to receive its own transmission on the outbound channel after a propagation delay to the headend and back. In the IEEE 802.3 specification, the bits up through the last bit of the source address field of the transmitted and received signals are compared, and a collision is assumed if they differ. There are several problems with this approach. The most serious is the danger that differences in signal level between colliding signals will cause the receiver to treat the weaker signal as noise and fail to detect a collision. The cable system, with its taps, splitters, and amplifiers, must be carefully tuned so that attenuation effects and differences in transmitter signal strength do not cause this problem. Another problem for dual-cable systems is that a station must simultaneously transmit and receive on the same frequency. Its two RF modems must be carefully shielded to prevent cross-talk.

An alternative approach for broadband is to perform the CD function at the headend. This is most appropriate for the split system, which has an active component at the headend anyway. This reduces the tuning problem to one of making sure that all stations produce approximately the same signal level at the headend. The headend would detect collisions by looking for garbled data or higher-than-expected signal strength.

IEEE 802 CSMA/CD. The IEEE 802 CSMA/CD standard [IEEE90b] is very close to that of Ethernet, and conforms to the preceding discussion. Figure 5.2 shows the MAC CSMA/CD frame structure. The individual fields are as follows:

- *Preamble*: a 7-byte pattern used by the receiver to establish bit synchronization and then locate the first bit of the frame.
- *Start frame delimiter (SFD)*: indicates the start of a frame.
- *Destination address (DA)*: specifies the station(s) for which the frame is intended. It must be a unique physical address (one destination transceiver), a multicast-group address (a group of stations), or a global address (all stations on the local network). The choice of a 16- or 48-bit address is an implementation decision and must be the same for all stations on a particular LAN.
- *Source address (SA)*: specifies the station that sent the frame. The SA size must equal the DA size.
- *Length*: Specifies the number of LLC bytes that follow.

- *LLC*: field prepared at the LLC level.
- *Pad*: a sequence of bytes added to assure that the frame is long enough for proper CD operation.
- *Frame check sequence (FCS)*: a 32-bit cyclic redundancy check value. Based on all fields, starting with destination address.

Token Bus

This is a relatively new technique for controlling access to a broadcast medium, inspired by the token ring technique discussed later. We will first provide a general description, then look at some of the IEEE 802 details.

Description. The token bus technique is more complex than CSMA/CD. For this technique, the stations on the bus or tree form a logical ring; that is, the stations are assigned logical positions in an ordered sequence, with the last member of the sequence followed by the first. Each station knows the identity of the stations preceding and following it. The physical ordering of the stations on the bus is irrelevant and independent of the logical ordering (Figure 5.12).

A control frame known as the *token* regulates the right of access. The token frame contains a destination address. The station receiving the token is granted control of the medium for a specified time. The station may transmit one or more frames and may poll stations and receive responses. When the station is done, or time has expired, it passes the token on to the next station in logical sequence. This station now has permission to transmit. Hence steady-state operation consists of alternating data transfer and token transfer phases. Nontoken-using stations are allowed on the bus. These stations can respond only to polls or requests for acknowledgment.

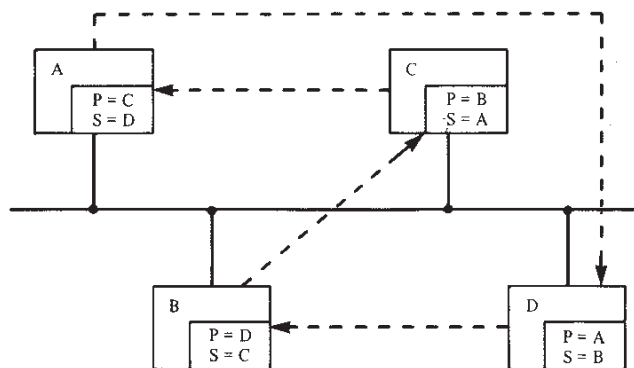


FIGURE 5.12 Token Bus

This scheme requires considerable maintenance. The following functions, at a minimum, must be performed by one or more stations on the bus:

- *Ring initialization*: When the network is started up, or after the logical ring has broken down, it must be initialized. Some cooperative, decentralized algorithm is needed to sort out who goes first, who goes second, and so on.
- *Addition to ring*: Periodically, nonparticipating stations must be granted the opportunity to insert themselves in the ring.
- *Deletion from ring*: A station must be able to remove itself from the ring by splicing together its predecessor and successor.
- *Recovery*: A number of errors can occur. These include duplicate address (two stations think it is their turn) and broken ring (no station thinks that it is its turn).

IEEE 802 Token Bus. The IEEE 802 token bus protocol follows the general principles outlined above [IEEE90c]. Figure 5.2 shows the MAC frame structure for token bus. The individual fields are as follows:

- *Preamble*: a one or more byte pattern used by receivers to establish bit synchronization and locate the first bit of the frame.
- *Start delimiter (SD)*: indicates start of frame.
- *Frame control (FC)*: indicates whether this is an LLC data frame. If not, bits in this field control operation of the token bus MAC protocol. An example is a token frame.
- *Destination address (DA)*: as with CSMA/CD.
- *Source address (SA)*: as with CSMA/CD.
- *LLC*: field prepared by LLC.
- *Frame check sequence (FCS)*: as with CSMA/CD.
- *End delimiter (ED)*: indicates end of frame.

The details of the protocol can be grouped into the following categories, which will be considered in turn:

- Addition of a node
- Deletion of a node
- Fault management by token holder
- Ring initialization
- Classes of service

First, let us consider how *addition of a node* is accomplished, using a controlled contention process called *response windows*. Each node in the ring has the responsibility of periodically granting an opportunity for new nodes to enter the ring. While holding the token, the node issues a *solicit-successor* frame, inviting nodes with an address between itself and the next node in logical sequence to demand entrance. The trans-

mitting node then waits for one response window or slot time (equal to twice the end-to-end propagation delay of the medium). One of four events can occur.

1. *No response*: Nobody wants in. The token holder transfers the token to its successor as usual.
2. *One response*: One node issues a *set-successor* frame. The token holder sets its successor node to be the requesting node and transmits the token to it. The requestor sets its linkages accordingly and proceeds.
3. *Multiple responses*: The token holder will detect a garbled response if more than one node demands entrance. The conflict is resolved by an address-based contention scheme. The token holder transmits a *resolve-contention* frame and waits four response windows. Each demander can respond in one of these windows based on the first 2 bits of its address. If a demander hears anything before its window comes up, it refrains from demanding. If the tokenholder receives a valid set-successor frame, it is in business. Otherwise, it tries again, and only those nodes that responded the first time are allowed to respond this time, based on the second pair of bits in their address. This process continues until a valid set-successor frame is received, no response is received, or a maximum retry count is reached. In the latter two cases, the token holder gives up and passes the token.
4. *Invalid response*: If the token holder hears a frame other than set-successor, it assumes that some other station thinks it holds the token. To avoid conflict, the station reverts to an idle or listen state.

Deletion of a node is much simpler. If a node wishes to drop out of the logical ring, it waits until it receives the token, and then sends a set-successor frame to its predecessor (the station that transmitted the token to it) containing the address of its successor. The existing station then sends the token as usual to its successor. On the next go-round, the former predecessor of the exited node will send the token to the former successor of the exited node. Each time that a station receives a token, it automatically sets its predecessor address to equal the source address of the token frame. Thus, the exited station is spliced out of the logical ring. If a node fails, it will not pick up the token when the token is passed to it, and this will be detected by the token sender, as explained below.

Fault management by the token holder covers a number of contingencies, listed in Table 5.6. First, while holding the token, a node may hear a frame indicating that another node has the token. If so, it immediately drops the token by reverting to listener mode. In this way, the number of token holders drops immediately to 1 or 0, thus overcoming the mul-

TABLE 5.6 Token Bus Error Handling

Condition	Action
Multiple token	Defer/drop to 1 or 0
Unaccepted token	Retry
Failed station	"Who follows" process
Failed receiver	Drop out of ring
No token	Initialize after time-out

multiple-token problem (which could be caused by two nodes having the same address). The next three conditions listed in the table are manifested during token passing. Upon completion of its turn, the token holder will issue a token frame to its successor. The successor should immediately issue a data or token frame. Therefore, after sending a token, the token issuer will listen for one slot time, to make sure that its successor is active. This precipitates a sequence of events:

1. If the successor node is active, the token issuer will hear a valid frame and revert to listener mode.
2. If the token issuer hears a garbled transmission, it waits four time slots. If it hears a valid frame, it assumes that its token got through. If it hears nothing, it assumes the token was garbled and reissues the token.
3. If the issuer does not hear a valid frame, it reissues the token to the same successor one more time.
4. After two failures, the issuer assumes that its successor has failed and issues a *who-follows* frame, asking for the identity of the node that follows the failed node. The issuer should get back a set-successor frame from the second node down the line. If so, the issuer adjusts its linkage and issues a token (back to step 1).
5. If the issuing node gets no response to its *who-follows* frame, it tries again.
6. If the *who-follows* tactic fails, the node issues a *solicit-successor* frame with the full address range (i.e., every node is invited to respond). If this process works, a two-node ring is established and life goes on.
7. If the *solicit-successor* tactic fails, it assumes that some major fault has occurred; either all other stations have failed, all stations have left the logical ring, the medium has broken, or the station's own receiver has failed. At this point, if the station has any more data to send, it sends that data and tries passing the token again. It then ceases transmission and listens to the bus.

Logical *ring initialization* occurs when one or more stations detect a lack of bus activity of duration longer than a time-out value: the token has been lost. This can be due to a number of causes, such as the network has just been powered up, or a token-holding station fails. Once its time-out expires, a node will issue a *claim-token* frame. Contending claimants are resolved in a manner similar to the response-window process. Each claimant issues a claim-token frame padded by 0, 2, 4, or 6 slots based on the first 2 bits of its address. After transmission, a claimant listens to the medium and if it hears anything, drops its claim. Otherwise, it tries again, using the second pair of its address bits. The process repeats. With each iteration, only those stations who transmitted the longest on the previous iteration try again, using successive pairs of address bits. When all address bits have been used, a node that succeeds on the last iteration considers itself the token holder. The ring can now be rebuilt by the response window process described previously.

As an option, a token bus system can include *classes of access* that provide a mechanism of prioritizing access to the bus. Four access classes are defined, in descending order: class 6, 4, 2, and 0.

Any station may have data to send in one or more of these classes. The object is to allocate network capacity to the higher-priority frames and send only lower-priority frames when there is sufficient capacity. To explain, let us define the following variables:

- THT = token holding time: the maximum time that a station can hold the token to transmit class 6 data
- TRT4 = token rotation time for class 4; maximum time that a token can take to circulate and still permit class 4 transmission
- TRT2 = token rotation time for class 2: as above
- TRT0 = token rotation time for class 0: as above

When a station receives the token, it can transmit classes of data according to the following rules (Figure 5.13):

1. It may transmit class 6 data for a time THT. Hence for an n -station ring, during one circulation of the token, the maximum amount of time available for class 6 transmission is $n \times \text{THT}$.
2. After transmitting class 6 data, or if there were no class 6 data to transmit, it may transmit class 4 data only if the amount of time for the last circulation of the token (including any class 6 data just sent) is less than TRT4.
3. The station may next send class 2 data only if the amount of time for the last circulation of the token (including any class 6 and 4 data just sent) is less than TRT2.
4. The station may next send class 0 data only if the amount of time for the last circulation of the token (including any class 6, 4, and 2 data just sent) is less than TRT0.

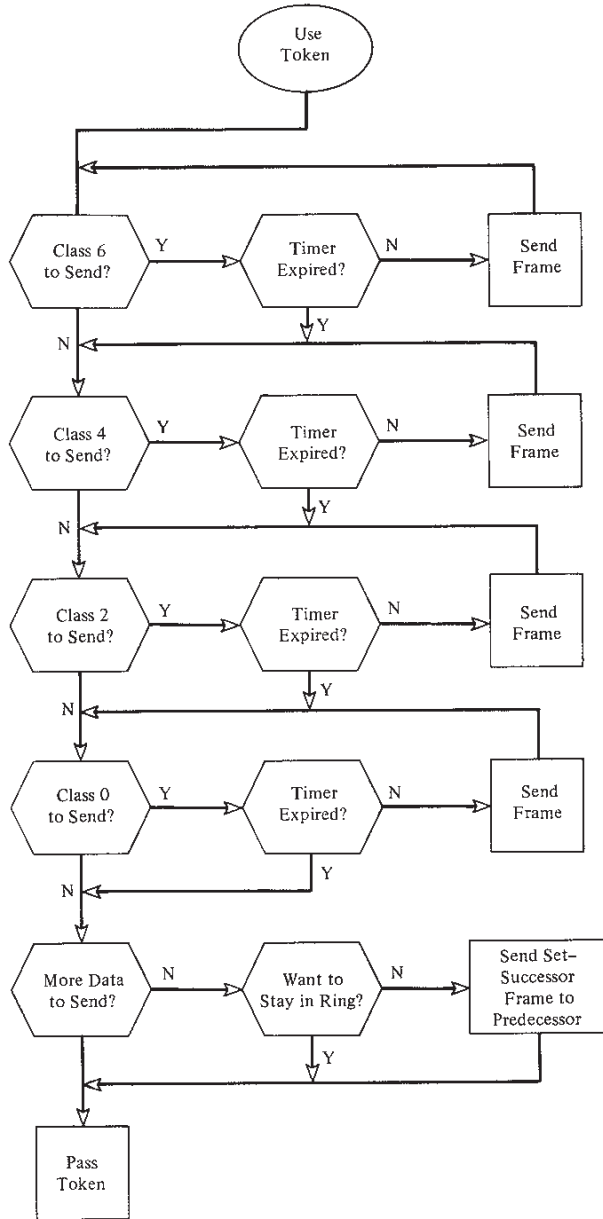


FIGURE 5.13 Token Bus Priority Scheme

