
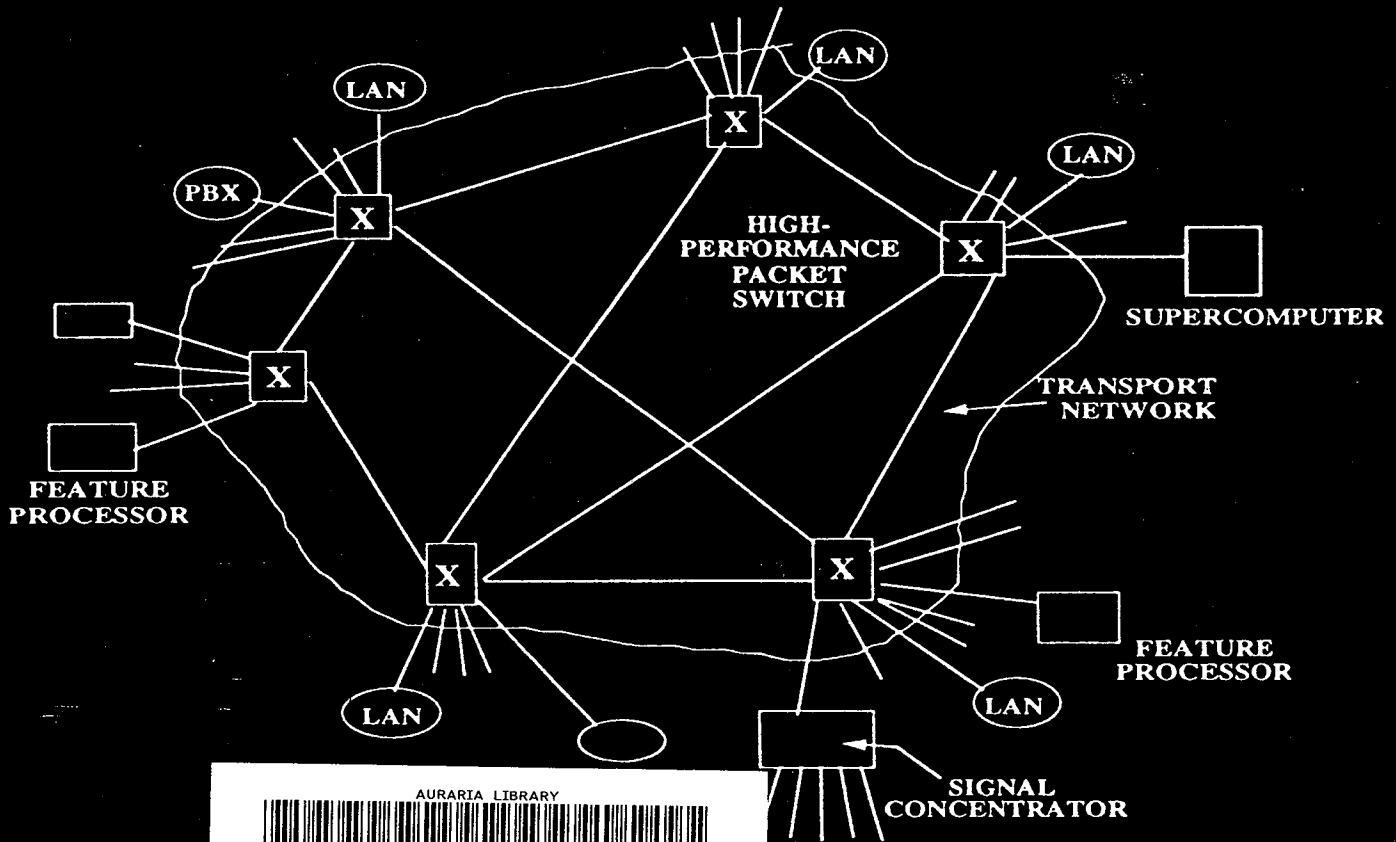


EXHIBIT B

Applications of Communications Theory
Series Editor: R. W. Lucky

An Introduction to Broadband Networks

LANs, MANs, ATM, B-ISDN, and
Optical Networks for Integrated
Multimedia Telecommunications



Acampora

An Introduction to Broadband Networks

LANs, MANs, ATM, B-ISDN, and Optical Networks for Integrated Multimedia Telecommunications

Anthony S. Acampora

Columbia University, New York, New York

The proliferation of high-speed local data networking and the potential of new multimedia service offerings have spurred worldwide research activities in broadband telecommunications networks.

A first-of-its-kind volume, *An Introduction to Broadband Networks* is a basic text that thoroughly explains local, metropolitan, and wide-area broadband networks for integrating voice, data, image, and video. Presenting important concepts in clear prose, the author examines modern telecommunications systems, such as current high-speed networks; describes newly emerging approaches; and offers possible foundations for the telecommunications infrastructure of the next decade. To encourage further study, the comprehensive discussions include highlighted sections on mathematical development and a selection of problem sets. With further insight provided through a combination of basic performance analysis and survey of applications, the text describes:

- The genesis of broadband networks
- The relationships between LANs, MANs (FDDI, IEEE 802.6 DQDB), WANs, asynchronous transfer mode (ATM), and broadband ISDN
- Several ATM switch architectures and a comparative study of their performances
- Traffic-related management strategies for wide-area ATM networks
- Emerging concepts for all-optical, or "passive," lightwave networks

Ideal as an introductory text for students and non-specialists new to this exciting field, *An Introduction to Broadband Networks* is also an important reference for experienced telecommunications professionals, including systems designers, hardware and software engineers, research and development managers, and market planners.

An Introduction to Broadband Networks

**LANs, MANs, ATM, B-ISDN,
and Optical Networks for
Integrated Multimedia Telecommunications**

Applications of Communications Theory

Series Editor: R. W. Lucky, *Bellcore*

Recent volumes in the series:

BASIC CONCEPTS IN INFORMATION THEORY AND CODING:

The Adventures of Secret Agent 00111

Solomon W. Golomb, Robert E. Peile, and Robert A Scholtz

COMPUTER COMMUNICATIONS AND NETWORKS

John R. Freer

COMPUTER NETWORK ARCHITECTURES AND PROTOCOLS

Second Edition • Edited by Carl A. Sunshine

DATA COMMUNICATIONS PRINCIPLES

Richard D. Gitlin, Jeremiah F. Hayes, and Stephen B. Weinstein

DATA TRANSPORTATION AND PROTECTION

John E. Hershey and R. K. Rao Yarlagadda

DIGITAL PHASE MODULATION

John B. Anderson, Tor Aulin, and Carl-Erik Sundberg

DIGITAL PICTURES: Representation and Compression

Arun N. Netravali and Barry G. Haskell

FIBER OPTICS: Technology and Applications

Stewart D. Personick

FUNDAMENTALS OF DIGITAL SWITCHING

Second Edition • Edited by John C. McDonald

AN INTRODUCTION TO BROADBAND NETWORKS: LANs, MANs, ATM, B-ISDN, and Optical Networks for Integrated Multimedia Telecommunications

Anthony S. Acampora

AN INTRODUCTION TO PHOTONIC SWITCHING FABRICS

H. Scott Hinton

OPTICAL CHANNELS: Fibers, Clouds, Water, and the Atmosphere

Sherman Karp, Robert M. Gagliardi, Steven E. Moran, and Larry B. Stotts

PRACTICAL COMPUTER DATA COMMUNICATIONS

William J. Barksdale

SIMULATION OF COMMUNICATION SYSTEMS

Michel C. Jeruchim, Philip Balaban, and K. Sam Shanmugan

A Continuation Order Plan is available for this series. A continuation order will bring delivery of each new volume immediately upon publication. Volumes are billed only upon actual shipment. For further information please contact the publisher.

An Introduction to Broadband Networks

**LANs, MANs, ATM, B-ISDN,
and Optical Networks for
Integrated Multimedia Telecommunications**

Anthony S. Acampora

*Columbia University
New York, New York*

PLENUM PRESS • NEW YORK AND LONDON

Library of Congress Cataloging-in-Publication Data

Acampora, Anthony S.

An introduction to broadband networks : LANs, MANs, ATM, B-ISDN,
and optical networks for integrated multimedia telecommunications /
Anthony S. Acampora.

p. cm. -- (Applications of communications theory)

Includes bibliographical references and index.

ISBN 0-306-44558-1

1. Broadband communication systems. I. Title. II. Series.

TK5102.5.A285 1994

621.382--dc20

93-41272

CIP

TK 5102.5 .A285 1994

Acampora, Anthony S.

An introduction to broadband
networks

ISBN 0-306-44558-1

© 1994 Plenum Press, New York
A Division of Plenum Publishing Corporation
233 Spring Street, New York, N.Y. 10013

All rights reserved

No part of this book may be reproduced, stored in a retrieval system, or transmitted
in any form or by any means, electronic, mechanical, photocopying, microfilming,
recording, or otherwise, without written permission from the Publisher

Printed in the United States of America

To my wife Margaret and my children Anthony and Rose, for the patience and understanding which they extended to me during the many hours which I consumed in preparing the original manuscript. Further, to my many friends and colleagues who have provided me with so much intellectual stimulation throughout my entire career.

Preface

This is an elementary textbook on an advanced topic: broadband telecommunication networks. I must declare at the outset that this book is not primarily intended for an audience of telecommunication specialists who are well versed in the concepts, system architectures, and underlying technologies of high-speed, multimedia, bandwidth-on-demand, packet-switching networks, although the technically sophisticated telecommunication practitioner may wish to use it as a reference. Nor is this book intended to be an advanced textbook on the subject of broadband networks. Rather, this book is primarily intended for those eager to learn more about this exciting frontier in the field of telecommunications, an audience that includes systems designers, hardware and software engineers, engineering students, R&D managers, and market planners who seek an understanding of local-, metropolitan-, and wide-area broadband networks for integrating voice, data, image, and video. Its primary audience also includes researchers and engineers from other disciplines or other branches of telecommunications who anticipate a future involvement in, or who would simply like to learn more about, the field of broadband networks, along with scientific researchers and corporate telecommunication and data communication managers whose increasingly sophisticated applications would benefit from (and drive the need for) broadband networks. Advanced topics are certainly not ignored (in fact, a plausible argument could be mounted that all of the material is advanced, given the infancy of the topic). However, the objective is to provide a gentle introduction for those new to the field. Throughout, concepts are developed mostly on an intuitive, physical basis, with further insight provided through a combination of performance curves and applications. Problem sets are provided for those seeking additional training, and the starred sections containing some basic mathematical development may be safely skipped with no loss of continuity by those seeking only a qualitative understanding.

Telecommunication networks have emerged as a strategic component of the worldwide infrastructure needed to support economic development, scientific discovery, educational opportunities, and social advancement. Driven by the pro-

liferation of high-speed local data networking and the potential of new multimedia service offerings, broadband telecommunication networks have become the focus of research, development, and standards activities worldwide. Underlying the field of broadband telecommunications are the concepts of universal interface and bandwidth-on-demand: all types of traffic are to be presented to the network in a common self-routing packet format, and high-bandwidth applications are distinguished from low-bandwidth applications only on the basis of the frequency with which the packets are generated by the users.

This book is focused on the principles of operation, architectures, performance management and traffic control strategies, protocols, emerging standards, and future directions for integrated broadband networks capable of supporting voice, data, image, and video telecommunication transport and services. The range of coverage spans from high-speed networks in use today (such as Local Area Networks) through those that will emerge over the next five years (Wide-Area ATM networks), up to and including those that might form the telecommunication infrastructure of the next decade (optical networks). The impact of rapidly advancing technologies, including VLSI, lightwave devices, and software, is addressed, and the genesis of broadband networks is explored. The relationship between LANs, MANs (FDDI, IEEE 802.6 DQDB), WANs, Asynchronous Transfer Mode (ATM), Broadband ISDN, and all-optical networks are developed, and several ATM switch architectures are presented and their performances compared. Traffic-related management strategies are described, including call setup procedures, signaling, congestion control, admission control, flow control, performance management, and network resource allocation. Traffic descriptors, important for developing call control procedures that guarantee variable/continuous bit rate service qualities, are discussed, along with possible rate-enforcement procedures. In an extension of conventional broadband network architectures, the concept of the "optical ether" is introduced, and possible second- or third-generation architectures for ultrabroadband all-optical or lightwave networks are presented, discussed, and studied. Finally, several categories of broadband multimedia applications are described.

Readers will gain an understanding of fundamental principles, underlying technologies, architectural alternatives, standards, and future directions for broadband networks. They will develop familiarity with LANs, MANs, WANs, FDDI, DQDB, ATM, and B-ISDN, as well as with traffic control and performance management issues and strategies. They will be exposed to concepts and approaches for self-routing ATM switches, and they will also be exposed to fiber optics principles, lightwave device technologies, and the potential of all-optical networks. Finally, they will be able to assess the opportunities for broadband multimedia service offerings.

The book contains eight chapters, starting with an introductory chapter that describes concepts and principles generic to all broadband networks. The origins

of congestion are explored, and basic congestion control procedures are established. Various types of telecommunication traffic are characterized, and the concepts of universal packet format and bandwidth-on-demand are introduced. The driving forces behind broadband networks are discussed, along with a description of the types of services and user applications enabled by broadband networks. The need for asynchronously multiplexed transmission is described, along with the basic concepts behind segmentation and reassembly of information-bearing telecommunication signals. The distinction behind connectionless and connection-oriented service is described, and communication protocol layering is briefly discussed.

Chapter 2 contains a review of LANs, including performance criteria and evaluation methodologies. Although LANs are well treated elsewhere and are, in many regards, distinct from broadband networks, their treatment has been included here to illustrate several relevant and fundamental concepts, to assist the student or practitioner having no prior exposure to this field, to introduce some of the mathematical techniques used to study packet-oriented transport networks of any type (LAN, MAN, broadband), and to present the principles behind the use of ATM (an important technology for broadband networks) for LAN service. This chapter may be skimmed or skipped by those already familiar with LANs, or only its starred sections may be omitted by those uninterested in the mathematical details.

Chapter 3 contains descriptions of Metropolitan Area subnetworks: Fiber-Distributed-Data-Interface and the IEEE 802.6 Distributed Queue Dual Bus, including operating principles, channel access protocols, performance, protocol data units, offered services, and integration into the broadband network environment.

Chapter 4 is devoted to principles and architectures of the ATM switch. Statistical smoothing at the input and output are distinguished, and several switch architectural alternatives are presented and compared, including the Banyan switch, the fully connected Knockout switch, the multistage Batcher-Banyan switch, the Tandem Banyan switch, and the Shared Memory switch.

Chapter 5 is devoted to ATM, the underlying technology for B-ISDN, including cell format, virtual paths and channels, the adaptation layer, signaling, and control.

Chapter 6 covers the principles of traffic control and performance management for broadband multimedia telecommunications. Congestion control that involves traffic descriptors, admission control at call setup time, and rate-enforcement procedures are contrasted with those that also involve flow control at the ATM cell level, buffer allocation and management, and selective cell loss. Quality-of-service metrics are developed, and the effects of time delay in a high-speed wide-area network are assessed.

Chapter 7 presents advanced concepts for all-optical or "passive" lightwave networks having an enormous capacity potential that is measured in units of

terabits per second. Capabilities of state-of-the-art lightwave devices such as tunable semiconductor lasers, tunable optical filters, wavelength-division multiplexers, and fiber-optic amplifiers are described, and principles of optical communications are presented. Constraints imposed on lightwave network architectures by lightwave device limitations are discussed, and candidate architectures that employ wavelength-selective routing with multihop access stations are presented. Techniques that exploit the relative independence between the physical topology of the optical medium and the logical connectivity among the access stations are presented, and logical reconfiguration via wavelength reassignment is presented as a technique to optimize overall capacity, adapt to network failures, and modularly expand the network.

Chapter 8 explores a range of possible applications enabled by broadband networks: a more natural mode of people-to-people communications; the access to (and manipulation of) multimedia electronic libraries by people and machines; wide-area distributed computing; delivery of entertainment HDTV signals; simplification of network management and control algorithms; execution of a rich set of diagnostic routines for improved network reliability; new educational delivery mechanisms; scientific computation and three-dimensional image rendering; and medical diagnostics and treatment programs.

Contents

Chapter 1. Introduction

1.1. Megatrends in Technology	1
1.2. The Evolving Architecture of Telecommunication Networks.....	3
1.3. Broadband Networks: Driving Forces	11
1.4. The Distinction between LANs and MANs.....	14
1.5. Some Possible Applications	16
1.6. A Word about Transmission Formats and Traffic Types	20
1.7. A Brief Word about Communication Protocols, Bridges, Routers, Gateways, Datagrams, and Virtual Circuits	26
1.8. Contention in Telecommunication Networks.....	33
1.9. Problems.....	36
References	40

Chapter 2. Review of Local Area Networks

2.1. Functional Elements.....	43
2.2. Functional and Physical Topologies.....	45
2.3. Random Access for a Functional Bus LAN.....	47
2.4. Throughput Analysis of a Slotted Aloha System.....	49
2.5. Throughput Analysis of an Unslotted Aloha System	51
2.6. Carrier-Sense-Multiple-Access	53
2.7. Throughput Analysis of Carrier-Sense-Multiple Access with Collision Detection	55
2.8. Token Rings.....	59
2.9. Throughput Analysis of Token Ring	61
2.10. Short Bus Architecture	63
2.11. Delay Performance.....	67
2.12. Appendix	81
2.13. Problems.....	81
References	85

Chapter 3. Packet Switch Interconnection Fabrics

3.1. Switch Functionality	87
3.2. The Need for Queuing in a Space Division Packet Switch	90
3.3. Analysis of Lost Packet Performance in the Absence of Smoothing Buffers	92
3.4. Input Queuing	94
3.5. Traffic-Handling Capability of an Input-Queued Packet Switch	96
3.6. Output Queuing	97
3.7. Mean Delay for a Packet Switch with Output Queuing	100
3.8. Switching Elements	102
3.9. A Modified Input-Queued Switch and the Required Number of β Elements	104
3.10. The Banyan Switch	111
3.11. The Knockout Switch	118
3.12. Analysis of Cell Loss for the Knockout Switch	131
3.13. The Batcher-Banyan Switch	132
3.14. The Tandem Banyan Switch	142
3.15. The Shared Memory Switch	146
3.16. Concluding Thoughts on Packet Switch Interconnecting Fabrics ..	149
3.17. Problems	149
References	155

Chapter 4. Metropolitan Area Networks

4.1. Distributed Queue Dual Bus	157
4.2. DQDB Segmentation, Reassembly, and Protocol Data Units	164
4.3. Delay Performance of DQDB for Bursty Data Traffic	173
4.4. Blocking Performance of DQDB for Circuit-Switched Traffic	182
4.5. Fiber-Distributed-Data-Interface (FDDI)	186
4.6. Delay and Blocking Performances of FDDI	194
4.7. Problems	195
References	197

Chapter 5. Broadband ISDN and Asynchronous Transfer Mode

5.1. Broadband ISDN and ATM: Preliminaries	199
5.2. B-ISDN Protocol Reference Model	203

5.3. Call Setup Procedures	205
5.4. Virtual Channels and Virtual Paths	207
5.5. Function of the ATM Adaptation Layer.....	212
5.6. ATM LAN	217
5.7. Problems.....	220
References	221

Chapter 6. Issues in Traffic Control and Performance Management

6.1. The Importance of Traffic Control and Performance Management.....	223
6.2. Admission Control.....	225
6.3. Policing.....	232
6.4. Flow Control	233
6.5. Priority Control	235
6.6. Self-Learning Strategies.....	236
6.7. Conclusion	237
6.8. Problems.....	238
References	240

Chapter 7. Lightwave Networks

7.1. What Is a Lightwave Network?	243
7.2. Essentials of Lightwave Technology	252
7.3. Direct Detection and the Quantum Limit.....	255
7.4. Wavelength Division Multiplexing.....	257
7.5. Principles of Coherent Detection	260
7.6. A Coherent Lightwave System.....	266
7.7. Optical Devices	268
7.8. Broadcast-and-Select Lightwave Packet Networks	274
7.9. Multihop Lightwave Packet Networks	279
7.10. Capacity of a Multihop Network—Analysis	286
7.11. Capacity of a Multihop Network—Results	288
7.12. Delay and Lost Cell Performance of Multihop Networks.....	291
7.13. Rearrangeable Lightwave Packet Networks.....	294
7.14. Problem Definition and Reconfiguration Heuristic	298
7.15. Traffic-Handling Capability of a Rearrangeable Multihop Network	303

7.16. Rearrangeable Multihop Networks with Wavelength Reuse	306
7.17. Conclusion	310
7.18. Problems	311
References	315

Chapter 8. Broadband Applications

8.1. Network Services Applications	317
8.2. End-User Applications	320
References	327

Index	329
--------------------	------------

Introduction

1.1. Megatrends in Technology

The field of modern telecommunications is being rapidly transformed by megatrends in three underlying core technologies: microelectronics, photonics, and software. In the field of microelectronics, advances in materials, design methodologies, high-resolution photolithography, and fabrication processes have produced high-gate-count very large scale integration (VLSI) circuitry characterized by ultrahigh reliability and capable of performing highly sophisticated functions at fast, real-time clock speeds. Submicron complementary metal oxide semiconductor (CMOS) technology permits single-chip custom integration approaching an equivalent count of one million gates if the circuit pattern is regular (e.g., memories, certain types of packet switches) and can be operated at a clock speed of 100–200 MHz. Emitter coupled logic (ECL), while not permitting as high a degree of integration as CMOS, can operate at clock speeds approaching 1 GHz with moderate functional complexity and even higher speeds with further-reduced functional complexity. Gallium arsenide (GaAs) technology, while not as mature as CMOS and ECL silicon technology, is capable of operating at clock speeds in excess of 10 Gbit/sec. These capabilities, in turn, have had profound impact on the fields of consumer and military electronics, computers, and telecommunications. For example, in the field of telecommunications, circuit boards populated by VLSI chips perform the routing, protocol processing, storage and media access control needed to enable low-cost Local Networking among desktop computers and workstations distributed throughout large office buildings and college campuses. Until quite recently, this type of functionality was feasible only in software running on a minicomputer at execution speeds that were but a tiny fraction of the speed at which these operations are performed today in VLSI-dominated hardware designed specifically for the intended application. Moreover, knowledge-based computer-aided design tools (or expert systems) permit workers to custom-design sophisticated circuitry, fabricate VLSI chips, and rapidly prototype, debug, and evaluate ideas for new telecommunication products and applications. Standard

programmable logic arrays provide even greater design flexibility, but do not offer the degree of functionality enabled by custom designs.

Similarly, advances in the field of photonics have had significant impact on telecommunications and, here, we may thus far have witnessed only the tip of the iceberg. The great appeal of lightwave technology as it applies to the field of telecommunications is the low loss, low dispersion, and extremely high bandwidth afforded by single-mode silica-based optical fiber and associated passive components (e.g., couplers, power splitters and combiners). When combined with narrow-linewidth single-frequency semiconductor lasers and low-noise optical receivers, it becomes possible for optical fiber links to support very high point-to-point digital data rates over very long unrepeated distances. The speed-distance product, long regarded to be the essential figure of merit of a point-to-point fiber-optic transmission system, has steadily increased over the past 15 years from the 2 Gbit-km/sec produced by multimode fiber systems using light-emitting diodes to about 200 Gbit-km/sec as produced by a single-mode fiber system using single-wavelength lasers. Substantially higher figures of merit are expected through deployment of optical amplifiers. As impressive as these gains have been, however, applications of lightwave technology to telecommunications have been largely limited to transmission systems, with optical fiber systems affording distinct technical, performance, and economic advantages relative to copper and radio-based systems. The architecture and very nature of the telecommunication network have not yet been fundamentally altered by photonics, but advances in optical amplifiers, wavelength-agile photonic components (lasers and optical filters), and passive optical components promise new applications to all-optical or lightwave *networks*, as opposed to lightwave transmission systems. A lightwave network allows the signals of many users geographically dispersed over some large service region to simultaneously share the bandwidth of a common optical medium (fiber, splitters, couplers, etc., which may be configured into a tree, bus, or some other arbitrary physical topology) in such a way that the identity of each signal is maintained and each signal is distributed by the tributaries of the medium to its correct receiver or set of receivers. Wavelength multiplexing is one technique that may accomplish these objectives. Such networks have the potential to unleash the vast bandwidth of the all-optical medium, three to four orders of magnitude greater than that occupied by a lightwave transmission system, thereby enabling the creation of an all-optical telecommunication infrastructure characterized by low-cost bandwidth and stimulating bandwidth-intensive end-user applications and other services needed to better support network operations (such as fault diagnostic routines).

In the software domain, powerful desktop workstations, file servers, object-oriented programming, distributed operating systems, distributed data bases, and distributed computing are creating the demand for telecommunication networks with features and capabilities far surpassing those intended for a voice and voice-

grade data-dominated traffic environment. New applications requiring the assembling and processing of vast amounts of digitally encoded information from geographically dispersed sources and data bases, the cooperation among computing resources separated by large distances, and the involvement of human workers networked by their workstations and collectively contributing to the resolution of scientific, business, political, and educational issues are transforming the functionality, architecture, and components comprising both citywide and worldwide telecommunication infrastructures. These same distributed software capabilities also afford a more thorough integration of computers and computing technologies into the control fabric of the telecommunication infrastructure.

1.2. The Evolving Architecture of Telecommunication Networks

This book is, most emphatically, *not* about these core technologies but, rather, is focused exclusively on telecommunication *networks*: the new system-level concepts, operating principles, architectures, performances, applications, and opportunities made possible by rapid advances in the aforementioned underlying core technologies. Networks intended to provide service over a local area (single building or small campus), metropolitan area (medium- to large-size city), and wide area (spanning thousands of kilometers) will all be treated. In all cases, the networks to be considered can be classified by the descriptor *broadband* and exhibit some common characteristics:

1. The network supports bandwidth-on-demand at each access port or access station to which a generic "user" (e.g., telephone, workstation, personal computer, host computer, network feature processor, network control computer, network data base, gateway to another network) attaches. As shown in the example of Figure 1.1, each "user" is attached to the network via a high-speed access link operating at some specified data rate. All information presented by the user is sent to the network via these access links in some common packet format. In the example of Figure 1.1, we assume that this common format consists of a sequence of short, fixed-length "cells" of information [as we shall see, this is the basis of asynchronous transfer mode (ATM)], and the user is responsible for transforming the signals generated by the actual terminating equipment (e.g., workstations) into this common format. In such an approach, large data blocks or continuous bit rate information streams are segmented and delivered to the network as a sequence of fixed-length cells; the larger data blocks are then reassembled from the cells arriving at the receiver. The terminating equipment controls its *effective bandwidth* by means of the *frequency* with which cells are generated, accessing more or fewer of the

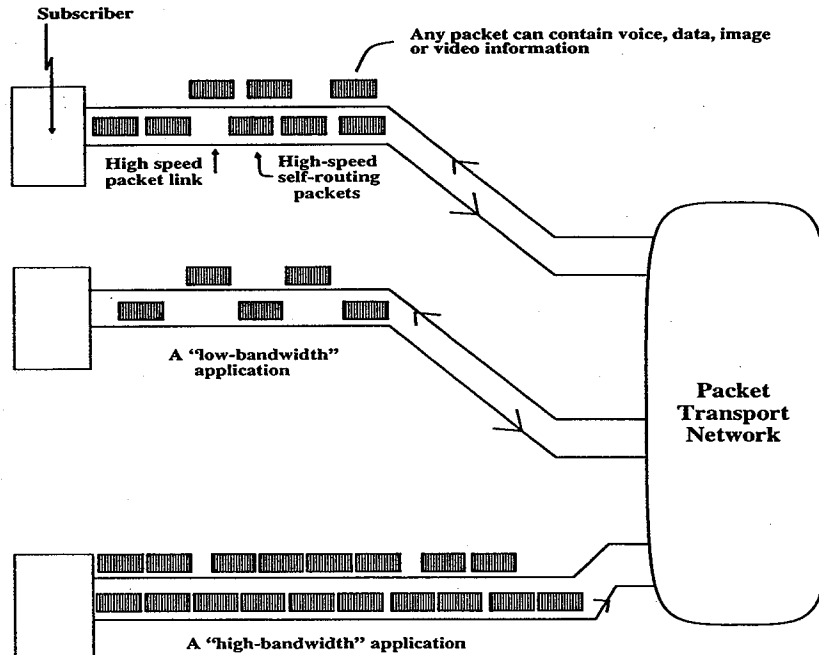


Figure 1.1. An illustration of bandwidth-on-demand.

fixed-length time slots appearing on the access links. The only feature that distinguishes a low-bandwidth application from a high-bandwidth application is the frequency of cell generation, and users can access bandwidth-on-demand at any effective data rate up to the maximum speed permitted by the access link. Within the network, all cells are handled autonomously. The network operates only on the cell *header* which contains the information needed by the network to deliver the cell to its intended receiver. The network disregards the nature or contents of the cell payload, which may contain voice, data, image, or video information (except, possibly, to satisfy the different quality-of-service requirements of different traffic classes, as will be discussed in Chapter 6).

2. In general, a given user attaches through one of a limited number of types of access links, each characterized by a unique format. Each type of access link therefore appears as a *universal port* to the class of users connected through that type of link. The link format may be a low-level specification including, for example, bit rate, packet or cell size, and location and contents of the payload and header information fields. Alternatively, a link format may also include higher-level functionality such as flow control. For a given format, bandwidth can be accessed on demand up to the

maximum bit rate specified for that format. Several universal formats or interfaces are provided to avoid encumbering a simple termination with the sophistication demanded by a complex termination. (For ATM, at least two types of universal ports are permitted, both of low-level functionality: one operates at 155 Mbit/sec, the other at 622 Mbit/sec).

3. A generic "user" can simultaneously generate independent multimedia information signals, each signal intended for a distinct receiver. The packets or cells created from these signals arrive at the network time-multiplexed through that user's access port, and the network uses the information contained in the cell header fields to rapidly direct each cell from the time-multiplexed sequence to its intended destination. Thus, a user chooses the receiver for each of its signals merely by placing the correct information into the header fields of the corresponding cells. This allows a user to rapidly direct each message of a message string to that message's intended receiver, or to concurrently maintain a communication link or connection with each of several other users, rapidly accessing a given connection and its respective receiver by again placing the correct information into the corresponding cell headers. Similarly, a given user may concurrently receive information signals from a multiplicity of senders. Each user may generate or terminate varying combinations of voice, data, image, and full-motion video traffic, and the network access stations must accept such fully integrated traffic, support the quality of service guaranteed to each traffic type, maintain the distinguishability among the messages multiplexed through a common access port, and maintain synchronization among the diverse traffic types that may be destined for (or arrive from) any one receiver (or sender). The integration of bursty traffic of variable bit rate with various types of continuous bit rate traffic, and the ability to rapidly forward each message of a sequence to a distinct receiver (essential for distributed processing, data communications, and other types of highly multiplexed applications including multiplexed voice and video transport), are hallmarks of emerging broadband networks.
4. Some universal ports may support circuit-switched connections for which the required network transport resources are reserved, or scheduled, at call setup time. Other universal ports support a packet-oriented transport in which all traffic types are segmented at the source into a stream of cells, with each cell being autonomously relayed through the network prior to reassembly at the destination (this type of port was demonstrated in Figure 1.1). In the latter case, either datagram (connectionless) or virtual circuit (connection-oriented) services may be provided to the user, as will be explained in Section 1.6.
5. The telecommunication network can be separated into two components:
 - (a) a transport network responsible for moving information among geo-

graphically remote user and equipment locations, relaying packets through switching nodes, and reliably delivering applied information to the intended receivers; and (b) a service network containing the feature processors, distributed data bases, and applications that are accessed by subscribers through the transport network, along with the distributed network control and management environment. The transport network is hardware-based in that the real-time access speeds are too fast to permit header processing, packet relaying, and protocol-related decisions to be made by software-driven processors. The overall supervision, management, and control of the network are, however, enabled by a system of software-driven processors that are interconnected via the transport network but do not directly process the high-speed, real-time, user-to-user information.

Today's telecommunication infrastructure is characterized by two distinct elements: transmission systems and switching systems. Transmission systems are responsible for the movement of properly formatted information on a point-to-point basis. Typically, as shown in Figure 1.2a, information arising from a multitude of sources are multiplexed onto a common transmission medium, which carries the information to a geographically remote destination where it is detected and demultiplexed into its constituent components. Here, multiplexing may be done in the time-division mode, in which the identity of each source of information is maintained by assigning that information to a particular, periodically recurring time slot of a larger transmission frame (more will be said of time division multiplexing later). In a variant of traditional synchronous time division multiplexing, the identity of the individual sources might be maintained by segmenting the informational content of each source into fixed or variable length packets, each of which carries a unique identifier within its header field. The packets associated with different sources are asynchronously time multiplexed via a device known as a statistical multiplexer (again, more later; this is the approach taken by ATM). Frequency division multiplexing is an alternative to time division multiplexing in which the identity of each source is maintained by modulating that source's information onto a unique carrier frequency, or channel, of a multichannel carrier transmission system (typically done for radio and satellite transmission links).

The end points of the transmission system might be either remote terminals which gather and distribute information to end-users, or switching nodes which terminate some multitude of transmission links, as shown in Figure 1.2b. It is the responsibility of the switching system to route the information corresponding to each end-user source to the correct outbound transmission system by forming the correct interconnection patterns among its input and output ports. For example, in a wide-area network, the connections $s_{1,1}$, $s_{2,1}$, \dots , $s_{N,1}$ in some given locality might be carried to the switch by a common transmission system and then de-

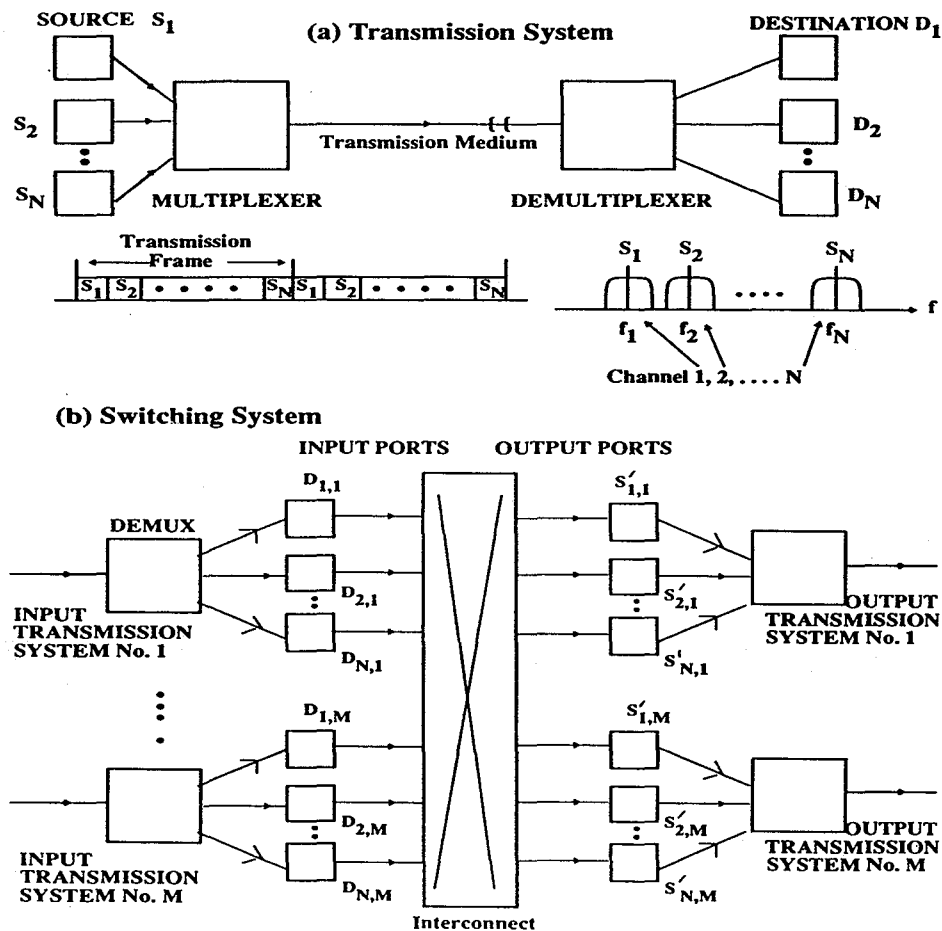


Figure 1.2. Transmission and switching system block diagrams.

multiplexed, switched, and remultiplexed for delivery, each to a destination in a different locality. The switch is reconfigurable; that is, the interconnection patterns among its input and output ports can be changed in response to changing patterns among source–destination pairs.

In today's telecommunication infrastructure, large, centralized switching systems of this type are interconnected by point-to-point transmission systems such that a connection originating in one locality might be multiplexed, transported, demultiplexed, switched, remultiplexed with other connections, transported, and so forth, many times before arriving at its ultimate destination in some geographically remote locality. Network services (such as administration and billing, diagnostics, and connection setup) and customer services (such as voice store-and-forward) are provided via dedicated feature processors, typically co-located with the switches within large switching offices. Control over these pro-

processors is typically exercised via a separate signaling network, an overlay to the communication network as embodied in the switching and transmission systems of Figure 1.2.

By contrast, the integrated broadband telecommunication infrastructure, drawing heavily on VLSI, lightwave, and software technologies, will be characterized by the following functional elements: transport, interfaces, and distributed services. The traditional boundary between switching and transmission systems is already becoming blurred, and, in the broadband era, the functionality of these systems will be replaced by the transport network which has exclusive responsibility for the movement of information from geographically dispersed input ports to the correct output ports. The transport network may contain traditional point-to-point transmission systems and centralized switching systems, supplemented by (or entirely displaced by) multipoint transmission systems. Such multipoint systems allow information-bearing signals to enter and leave the transmission medium from multiple, geographically distributed access stations, and hardware-based intelligence located within each access station is capable of making local routing decisions (accept a segment of information, ignore that information, regenerate and relay the information, etc.), effectively creating a distributed switch. Users (which may, themselves, be gateways to other networks) will be interconnected by the transport network and attach via bandwidth-on-demand universal ports. Services provided by the network attach via universal ports of the same design as those by means of which users attach. Control information is transported among the users and feature processors over the transport network, i.e., the transport network does not distinguish between control and communication signals. Similarly, the transport network does not distinguish a feature processor from any other user; it is responsible exclusively for delivering the information entering at the universal ports to the correct outputs, but is not concerned with the informational *content* of the signals. In this sense, as shown in Figure 1.3, the distributed transport network views all information as having originated in some generic distributed *application*, the components of which are attached through the universal access ports. Different applications might be distinguished by the amount of bandwidth drawn from each port, but even this is ignored by the transport network as long as the total demand per port (in either direction) does not exceed the capacity of the port. End-users derive enhanced services from the network by interacting with the applications provided by the network. Alternatively, large distributed end-user organizations may install their own shared applications, each of which may itself involve cooperation among geographically separated software routines. In this manner, the transport network might provide a "virtual private network" for large, geographically dispersed end-user organizations. The fact that multiple end-user organizations are, in fact, sharing a common transport network is totally transparent to the end-users. Administration and billing, connection setup, diagnostic routing, and other network services are sup-

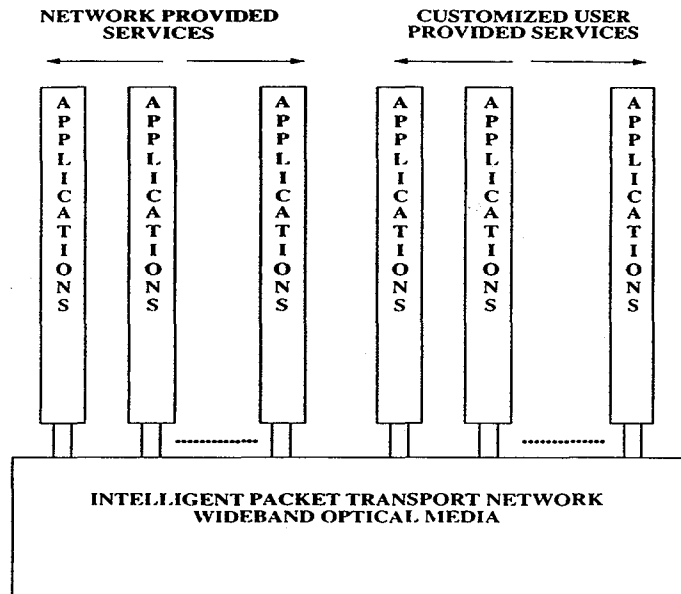


Figure 1.3. A conceptual representation of a distributed communication network.

ported as yet additional distributed applications running over the transport network.

The transport network may assume one of several forms. As shown in Figure 1.4, the transport network may contain some number of geographically dispersed self-routing packet switches. These switches, using custom-designed VLSI circuitry, read the header fields of each arriving packet, making the appropriate decision with regard to output link. Contention for the output links is handled via smoothing buffers contained within the switches (more will be said of this later). The switches are interconnected by point-to-point transmission facilities which serve exclusively to carry packets to the next relaying point, and, as shown in Figure 1.4, distributed access to transmission facilities is not permitted. The network is "transparent"; it does not process the user-applied data produced by the local area networks (LANs), private branch exchanges (PBXs), or other user devices attached to the network, but simply transports the information to the correct output. Higher-level protocol functions (retransmission requests, selection of flow control parameters, etc.) are performed externally to the network, using any protocol model agreed on by the two or more participants involved in any given connection (there is always at least one transmitting and one receiving participant). The transport network is totally devoid of software, insofar as the actual transport of applied data is concerned; software processing of packet headers as required for switching is simply too slow at the real-time data rates of the interconnecting transmission links. Software for non-real-time supervisory functions (e.g., maintenance, diagnostics, call setup) resides within feature processors

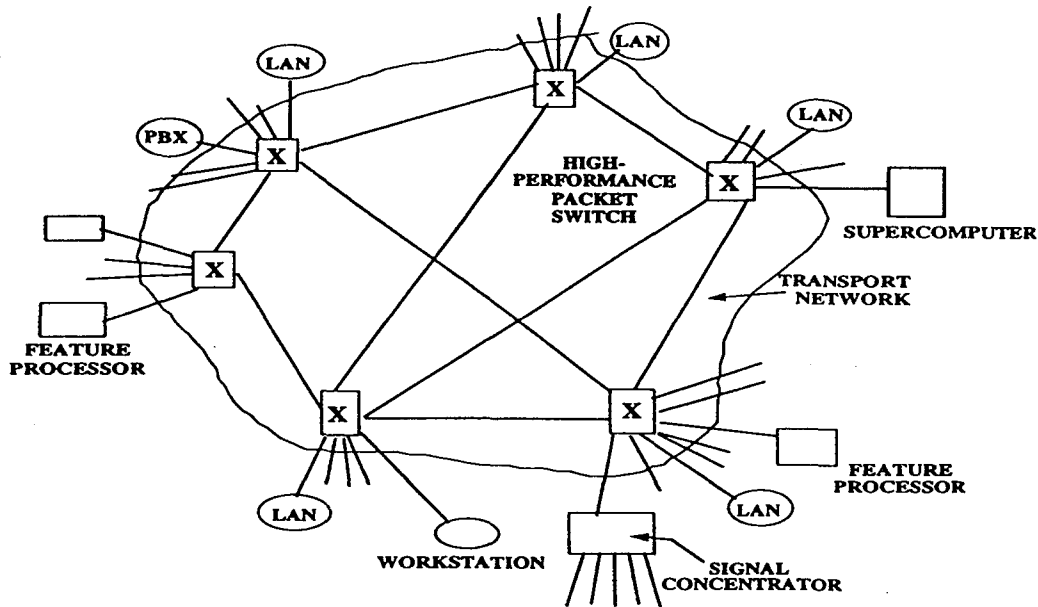


Figure 1.4. A transparent multimedia network using packet switches interconnected by fixed-point transmission links, surrounded by its supervisory feature processors and signal-supplying users.

which are external to the transport network and which connect to the transport network through the universal ports, appearing to the transport network like any other group of users.

An alternative architecture for the transport network appears in Figure 1.5. Here, the transport network consists of a passive optical medium containing fiber, optical couplers, splitters, power combiners, etc., and a system of geographically distributed access stations. All logical functions (packet switching, buffering, etc.) are performed exclusively within the access stations which contain all of the real-time electronics; the "optical ether" serves simply to move information among the access stations. For example, many high-speed channels might be wavelength multiplexed onto the "optical ether," each providing a dedicated path between two (one transmitting and one receiving) access stations. The user pair connected by a given optical channel may be changed with time by tuning the optical transmitters and receivers contained within each access station to different wavelengths. Traffic generated by a given user is transparently transported from that user's access station to the access station of the receiver. The access stations provide real-time transport, again devoid of software data-handling (software is too slow). Non-real-time supervisory functions are provided by software and processing elements distributed among the access stations to create a distributed management environment which does not directly process the real-time high-speed traffic. Enhanced services (data base access, shared computing facilities,

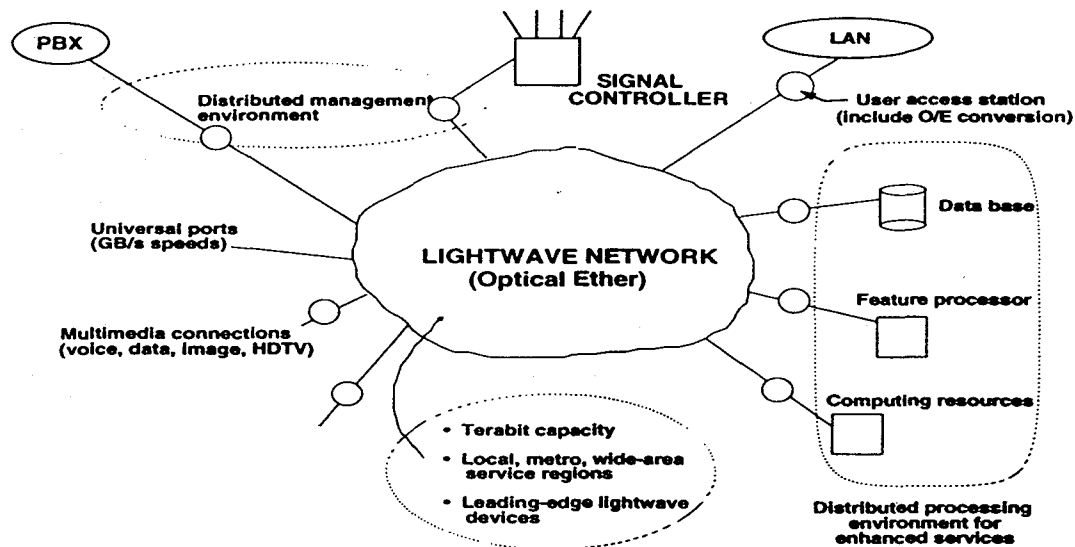


Figure 1.5. A distributed telecommunication network based on the notion of the "optical ether."

etc.) are provided via equipment, external to the transport network, each component of which connects to the transport network via one of the distributed access stations.

1.3. Broadband Networks: Driving Forces

Interest in broadband networks first began to develop in the early 1980s. At that time, LANs were finding their way into offices, universities, and small industrial parks, and the need developed to extend the intrapremises, data-oriented services provided by LANs to communities of users who might be geographically located throughout cities (metropolitan area networks) or larger regions (wide-area networks). The intrapremises data services provided by LANs include, for example, terminal-to-host access, host-to-host file transfers, distributed processing, shared data base access, and electronic mail. The LAN equipment used to support these services is well matched to the bursty data nature of the telecommunication traffic: a single, high-speed channel, time-shared by user access stations, which implements the rules for accessing the channel, with information presented in a format suitable for packet switching. Packets generated by the data device (e.g., terminal, host, PC, file server) attached to a particular access station are sequentially addressed to different destinations by placing the appropriate information in the packet header, thereby avoiding the unacceptable time delay associated with reserving channel time for each packet by means of a centralized setup

procedure. Each access station reads the header of every packet broadcasted onto the shared channel and accepts those packets determined to be intended for local reception. The ability of each access port to support multiple simultaneous sessions and the high peak data rate afforded by the common shared channel are, perhaps, the primary distinctions between data services offered by a LAN and other types of data services such as might be offered, for example, by a digital PBX.

In this context, the concept of the metropolitan area network (MAN) began to develop as a means to enable interpremises extension of LAN services. As shown in Figure 1.6, the MAN consists of a larger, interpremises version of the LAN shared channel, with one access station of each LAN serving as the concentrator for interpremises traffic. In this way, hosts and data bases connected by a LAN in one building may be accessed by terminals and workstations connected to a LAN in another building. The MAN would provide the interpremises inter-




Figure 1.6. Use of a metropolitan area network to extend the geographical service range of local area networks.

connection among LANs, thereby creating the appearance of a "super-LAN" to extend LAN data services over metropolitan area geographies. This "super-LAN" would be a private network: owned, installed, operated, and maintained by a single entity and providing service to users affiliated with that entity.

Toward the mid-1980s, the common carriers began to develop an interest, first in MANs and soon thereafter in broadband networks to enable switched, high-speed public offerings. The interest of the carriers was motivated by (1) the increasing number of LANs installed within their service regions and the concurrent demand for LAN interconnects; (2) the concern that large corporate customers, demanding MAN-like service to interconnect their rapidly proliferating LANs, would install their own private MANs, thereby depriving the carriers of a potentially important source of new revenue; and (3) the opportunity to offer new broadband services, based on newly emerging technologies, to both business and residential customers. This latter consideration has since caused a marked departure from the earlier broadband driving force: the broadband network would no longer be limited to LAN interconnects and data-oriented services, but instead would be repositioned on a more pervasive basis as the enabler of integrated voice, data, image, and full-motion video services to large and small customers over wide service areas. In this context, networks intended for use either exclusively as metropolitan area LAN interconnects or as the point of entry to a pervasive, wide-area broadband infrastructure are referred to as metropolitan area subnetworks, MAN subnetworks, or simply as MANs.

Central to this change of emphasis is the notion of bandwidth-on-demand. As previously described, each subscriber would be provided with an access link capable of transporting digital information at a peak data rate commensurate with the bandwidth of that link. Multimedia information (voice, data, image, video) would be presented to the link in some specified self-routing packet format, and the frequency with which packets are generated and loaded onto the link would determine the effective "bandwidth" derived from the link.

The maximum rate of packet generation could be no greater than that which would cause the link to be loaded 100% of the time. For example, the packet stream generated by a 64 kbit/sec digital voice connection would load a 155 Mbit/sec link (one of the standard rates for ATM) only for a very small fraction of the time; a digital video connection, however, might load the same link for a significant fraction of the time. Voice, data, image, and video packets would all be multiplexed onto the link serving a particular subscriber and the packets corresponding to different types of traffic, or different packets of the same traffic class, could be forwarded to different network destinations. The only constraint is that the average packet loading does not exceed 100% on the link. (In practice, the average packet loading might typically be limited to approximately 80% of the link capacity in order to avoid excessively large queue buildup for bursty traffic.) The bandwidth-on-demand port to the network would accept and deliver properly

* formatted packets, independent of the packet content or traffic type, thereby assuming the characteristics of a "universal" information outlet, much analogous to an electric power outlet: the power outlet delivers electricity to any type of appliance (e.g., toaster, dishwasher), as long as the current delivered does not exceed the rating of the circuit; the universal information outlet can support the packets of any digital device (e.g., host, digital telephone) as long as the average rate of packet generation does not cause depletion of the bandwidth of the link.

The notions of bandwidth-on-demand and broadband multimedia networks are critically dependent on the provisioning of fiber-optic facilities to the home and office. For reasons to be described later, the bandwidth potential of a fiber-optic link is enormous, and the actual capacity of the link is limited by speed constraints of the terminating electronics and electro-optic components (e.g., semiconductor laser, photodetector, optical receiver). Unlike the twisted copper pairs which are wired to most homes and offices today and which are inherently limited to operating speeds of several megabits per second, the terminating equipment of a fiber-optic link could readily support speeds as high as several gigabits per second, more than adequate for broadband networks; and even this is but a small fraction of the inherent bandwidth of the medium.

Another important aspect of public switched broadband networks is the provisioning of enhanced features and services for multimedia traffic to complement the transport network for such traffic. For example, as an enhanced feature, a high-resolution facsimile service node might be provided which could accept a high-resolution document transmitted by a subscriber, continually retry delivery to busy facsimile machines, and duplicate and deliver the document to all recipients on some specified mailing list. More will be said later of other possible (and considerably more exciting) broadband applications.

In this book, we will develop the principles of broadband networks both for switched public services and, as a subset, for private LAN interconnects. In so doing, we shall develop the concepts of, and demonstrate the interrelationships among LANs, MANs, self-routing packet switching, ATM, Broadband Integrated Services Digital Network (B-ISDN), Distributed Queue Dual Bus (DQDB, the emerging IEEE 802.6 standard), Fiber-Distributed-Data-Interface (FDDI), and later-generation very-high-speed all-optical networks. Performance analysis will be summarized in graphical form to permit comparison among alternative approaches and architectures. Protocol details for emerging B-ISDN, DQDB, and FDDI will be presented as appropriate.

1.4. The Distinction between LANs and MANs

The distinction between LANs and metropolitan area subnetworks is primarily quantitative. In general, a metropolitan area subnetwork is quite similar to a local area network in that both are based on distributed switching and distributed

transmission systems which, collectively, comprise the distributed transport network. For both, there is the notion of a shared transmission medium which combines and carries signals among a multitude of user pairs. For both, information is carried in a self-routing packet format, and the identity of the information which flows between each source-destination pair is maintained by appending to each packet some appropriate header identification field.

However, unlike the LAN, which is intended to provide data-oriented services among users who are spread out within a building or over a small campus, a MAN typically involves much larger distances: the coverage region of a LAN is typically under several kilometers, while the MAN typically covers a region spanning several tens of kilometers or even higher. Also, the transmission speeds exhibit order-of-magnitude differences, with the LAN operating, typically, at a rate of under 10 Mbit/sec and the MAN, typically, at a rate of greater than 100 Mbit/sec. These distinctions, longer distance and higher speed, collectively require that different strategies be used to enable the transmission medium to be shared among a plurality of users.

The aggregate capacity offered by a network is, by definition, the largest permissible value for the total volume of information carried by the network such that some specified quality-of-service objectives (e.g., delivery delay, fraction of lost or misrouted information) can be maintained, that is, the aggregate capacity is the largest possible value for the total average traffic presented by the individual sources. By virtue of the higher transmission speeds employed and the fact that multiple shared-media MAN subnetworks might be combined with remote statistical concentrators and packet switching nodes to form a wide-area broadband network, the aggregate capacity provided by the broadband network is enormous (tens or hundreds of gigabits per second). Although the aggregate capacity of a broadband network may be quite high, no one user can present a traffic load exceeding the capacity of the universal port. This capacity is, in general, slightly lower than the transmission speed of a single LAN or MAN subnetwork's shared channel, to accommodate inefficiencies associated with media access and packet headers. Thus, for a MAN, the link access speed is typically greater than 100 Mbit/sec but rarely greater than several gigabits per second. (The electronics and electro-optics which would be needed to drive the medium at a rate greater than 2-3 Gbit/sec would be regarded, as of this writing, as available only in research laboratories. This is the so-called fundamental electro-optic bottleneck; even as this speed is gradually increased with advancing technology, there will still remain an enormous mismatch between electro-optic data rates and the bandwidth of optical media, which constrains the permissible architectures of lightwave networks; more will be said of this in Chapter 7.)

Another quantitative difference between LANs, MAN subnetworks, and wide-area broadband networks concerns the expected number of subscribers. A premises-based LAN typically interconnects fewer than several hundred access stations, although this number may grow to be as high as several thousand for a

particularly large installation. By contrast, a MAN subnetwork carries the inter-premises traffic of several or many LANs, and the total number of subscribers whose traffic must be managed is correspondingly higher.

Because larger geographical distances are involved, the user of the MAN subnetwork or the broadband network suffers a larger end-to-end propagation delay as compared with the LAN user. For the MAN subnetwork, this longer propagation delay requires that different media access schemes be used. For wide-area broadband networks, the longer propagation delay requires traffic management strategies which can tolerate appreciable signaling delay and which recognize that, because of the high capacities and long distances, much information is resident on the transport network at any given time. Longer propagation delay might also translate into longer latency insofar as cooperation among distributed computers is concerned, but this effect might be partially offset by the higher link access speeds involved.

Finally, unlike the LAN, which is always operated as a single-owner private network, the MAN subnetwork and wide-area broadband network might be either publicly or privately owned and operated. A public network, as previously noted, might provide virtual private network services for large users.

1.5. Some Possible Applications

As mentioned, the original application intended for broadband networks was to extend the on-premises data services enjoyed by LAN users across large geographies. Later, the notion of a universal information outlet began to emerge. Although these remain important, a sampling of other possible applications, spanning the end-user and network operations domains, might include:

1. *Provisioning of fiber-to-the-home.* In this variant of the classic chicken-and-egg dilemma, the advent of broadband networks is dependent on penetration of fiber-optic facilities to large offices, homes, and small offices. At the same time, economical feasibility of fiber-to-the-home, based on conventional star-on-star wiring topologies (one link per home/office, hubbed at a remote terminal, with local loops further hubbed at the central office) has yet to be demonstrated, especially if voice and voiceband data are the only types of traffic to be carried. Broadband networks may help to break this cycle by stimulating the market for new revenue-generating multimedia services while at the same time providing high-speed transport over a physical plant characterized by lower-cost distributed topologies. Physically, the plant might look much like the distribution network for cable TV signals but would borrow heavily from network notions of packet access, distributed ports, and shared media to provide two-way bandwidth-on-demand connections between many pairs of access ports rather than merely the delivery of

a limited number of signals broadcasted from some head-end to downstream receive-only stations.

2. *Enhanced network reliability.* Broadband is often viewed as a technology with the potential to offer multimedia applications to large and small users alike, by virtue of the economical provisioning of link access speeds over two to three orders of magnitude faster than narrowband, voice-oriented networks. Often overlooked is the opportunity to reserve a significant amount of capacity, taken from the large aggregate pool, to enhance network reliability. By exploiting the unprecedented availability of bandwidth, not only might shared backup facilities be reserved to spare failed resources, but a rich set of capacity-consuming diagnostic routines might be developed and installed to continually test and monitor the status of network hardware and software elements. The greater the capacity devoted to running test routines, the more diagnostic data that can be generated, the better the ability to detect even subtle bugs and failures, and the faster the network can be reconfigured to route information around the point of failure.

3. *Lightweight communication protocols.* Historically, as communication applications grew more sophisticated, a set of protocols or rules governing the flow of information and various quality checks were developed to enable reliable delivery of information in a relatively hostile network environment characterized by noisy, bit error-prone channels and congested network links and switching nodes. Protocols intended to ensure the integrity of delivered data, control the flow of traffic over congested elements, and monitor for data lost by buffer overflow and other causes, consume processing resources and constrain the sustainable rate of information flow among terminating devices to the fastest that the protocol processor will permit. For example, even over a high-speed high-quality link, execution of a protocol intended to protect information over a lower-quality link might limit the effective rate of information flow between two machines (the throughput) to a rate of perhaps a few megabits per second. Furthermore, increasing the link speed alone, but leaving the protocol unchanged, cannot provide for a higher throughput. While the development of faster, more customized processors will certainly relieve this "protocol bottleneck," a major improvement will also result from the development of so-called "lightweight" protocols which exploit the relatively benign characteristics of broadband networks to eliminate the need for extensive layers of protection. Since the broadband network will offer enormous aggregate capacity, it is possible to operate the network with a comfortable margin of safety, allowing it to carry a volume of traffic very large by today's standards, while at the same time maintaining a comfortable distance from the capacity limits. In this relatively uncongested environment, the protocols needed to protect against traffic overload (flow control, delivery prioritization, buffer overflow, etc.) might be scaled back or eliminated entirely. Furthermore, the fiber-optic medium associated with the broadband network is designed to produce a bit error rate orders of magnitude lower than that provided by radio or copper alternatives (10^{-12}

versus 10^{-6} , for example), and the requirements for error detection/error correction protocols might be relaxed considerably. Merely by permitting the streamlining or elimination of multiple protocol layers, broadband networks may permit an order of magnitude or more increase in throughput among communicating processors and workstations (megabits or tens of megabits to hundreds of megabits per second). Such an increase is essential to produce the tight computer-to-computer coupling required for efficient distributed processing.

* 4. *Broadband multimedia connections.* Voice and voiceband data are the dominant forms of end-user traffic carried by today's telecommunication network. End-users are offered (a) circuit-switched 3-kHz analog channels, (b) circuit-switched 64 kbit/sec channels, and (c) for the Integrated Services Digital Network Basic Rate Service, two independently switchable 64 kbit/sec channels plus a 16 kbit/sec channel that can carry packet or circuit-switched data. By contrast, the broadband network will offer bandwidth-on-demand, multimedia connections, which integrate voice, data, image, and full-motion video traffic through universal ports operating at link access speeds in excess of 100 Mbit/sec. Traffic, in each of its forms, will be reduced to a sequence of addressed data packets, with a full-motion video connection distinguished from a voice connection, as far as the network is concerned, only by its significantly higher rate of packet generation. Again, the use of VLSI renders practical the formatting of information in all of its multimedia forms into self-routing packets which can be carried over a common transport network. All forms of traffic can be carried on a common connection, and a multitude of connections, destined for the same or different network access ports, can be multiplexed through a common universal port.

5. *Access to distributed data bases, pictorial archives, and electronically published information.* Delivery of vast, electronically published information libraries to a large number of subscribers is currently limited by (a) insufficient end-user bandwidth provided by a narrowband telecommunication infrastructure; and (b) lack of suitable data base software. A high-quality black-and-white image (400 dots/inch, $8\frac{1}{2}$ by 11-inch size) contains about 15 million bits. Such an image could be delivered, in uncompressed form, to a human reader in about one-tenth of a second over the 155 Mbit/sec user port of a broadband network, but would require over 50 min via a 4800 bit/sec voiceband data link. Time of delivery over the broadband network is consistent with the patience of most people and the human responsiveness for archival browsing: the voiceband connection does not come close. Furthermore, only images containing fewer than 480 bits of information could possibly be compressed and delivered within one-tenth of a second by the aforementioned voiceband data link (such an image would be very uninteresting indeed, consisting for the most part only of large black blotches over a white background, or vice versa). The insufficiencies of narrowband networks are even more dramatic if color images are involved. With a broadband infrastructure, one can envision a system of distributed multimedia data bases and a software en-

vironment which enables each of a large number of users to synthesize needed information by accessing and assembling segments from several sources. For example, a user might electronically browse through a favorite new magazine as naturally and conveniently (and with the same graphical quality) as if hard copy were at hand; for a particularly interesting article, video segments from the evening news might be added. Not only does the broadband network provide the required link access speed for natural interaction but, also, it enables high-speed cooperation among data bases and processors to collect, assemble, and display the requested information.

6. *Entertainment video on demand.* Over-the-air and cable broadcast systems offer viewers the ability to select stations from a limited menu, with the content of the programming and hour of viewing entirely controlled by the network and/or cable operators. Broadband networks will provide adequate aggregate capacity and link access speed, along with two-way interaction, to permit viewers to select individual programs rather than broadcasted channels. Furthermore, each viewer can request a particular program for viewing at a time of day convenient to that viewer. In addition to handling an unprecedented volume of end-user traffic, the broadband network must, once again, allow programs resident in geographically dispersed data bases to be rapidly accessed and locally cached since, undoubtedly, multiple viewers will want to view the same program but with overlapping start times. The degree of difficulty encountered is strongly dependent on the start-time granularity.

7. *Natural, lifelike video conferencing.* Video conferencing has been heralded as a much-needed productivity enhancer because, in principle, it tends to avoid the need for time-consuming business travel. Today's bandwidth-constrained video conferencing services are characterized by large studios at each of several locations, with the physical presence of each participant required at one of the studios. One or more video monitors are present in each studio, enabling participants to see and hear participants at other locations. The quality of the image is often lacking because of the compression algorithms currently used to remove redundancy from the video signals and allow the signals to be carried over standard transmission facilities, typically at 1.5 Mbit/sec. To this is added a highly artificial environment in which the video camera at each remote location is steered toward the conference participant who is currently speaking or is speaking the loudest. The ability to judge the reaction of other participants and to engage in side discussion, a natural part of dynamic interaction, is totally lacking. Broadband networks will overcome many of these video conferencing shortcomings. Within a broadband environment, conference participants might not need to leave their individual offices. Each participant might select the video signal from any other participant (or, concurrently, from a group of geographically separated participants) for local display on a high-resolution monitor. Images of all participants might appear within a bar at the top or bottom of the display. The current speaker

might be highlighted by the display. Windows could be added so that a participant could watch the current speaker, while simultaneously conducting a side conversation with another participant. Other windows might contain electronic blackboards, and documents might be distributed by high-speed scanners and printers. In short, the abundant capacity of the broadband environment might be applied to enhance the naturalness of a multiparticipant video conference.

1.6. A Word about Transmission Formats and Traffic Types

For traditional byte-interleaved time division multiplexing, as illustrated in Figure 1.7, time on a transmission link is divided up into repetitive, fixed-length intervals known as transmission or switching frames, each of which is further subdivided into time slots capable of carrying one byte of information. The time slots are sequentially numbered throughout the frame. A time-multiplexed connection over such a transmission link consists of an allocation of one particular time slot from the switching frame, repetitively recurring on all successive frames. Information from the connection is broken down into byte-size units which are then carried in the assigned time slot on repetitive frames. Different time slots within the frame carry information corresponding to different connections. The identity of a particular connection is uniquely determined by the position of the assigned time slot within the switching frame. With such an approach, a network controller, or scheduler, is always needed to assign each newly requested connection to a particular time slot, and to return time slots to the unused pool, for possible reassignment, upon termination of a particular connection. Since each connection is uniquely determined by its assigned time slot, a switch or demultiplexer which might terminate the transmission line can be instructed to route the connection on the basis of position within the switching frame. The switching




Figure 1.7. A typical byte-interleaved transmission format. Here, a given connection is uniquely determined by the fixed time slot assigned for the connection in each frame. A scheduler is always required to manage the time slot allocation.

sequence repeats frame after frame until a new connection is added or an existing one is dropped. Connections are assigned by the network controller to the time slots of two or more transmission links terminating on the same switch such that at no time are the time slots on more than one inbound link destined for the same outbound link. This may require the use of time slot interleaving within the switch to permit conversion of an inbound time slot assignment to a different outbound assignment. In this way, the network controller manages the competition among connections by means of reservation and time slot scheduling.

The number of time slots per frame, and the period of the switching frame, collectively determine both the data rate of the transmission line and the data rate of each connection. For example, in the DS1 transmission system, each 125 μ sec switching frame contains 24 time slots, each one byte in duration, plus one additional bit, not associated with any connection, used for signaling and to delineate the frame boundaries. Each connection therefore consists of one byte every 125 μ sec for a data rate per connection of 64 kbit/sec. The frame-mark bit, regularly occurring every 125 μ sec, itself comprises an 8 kbit/sec signaling channel. Since there are 24 time slots per frame, each capable of supporting one 64 kbit/sec connection, the data rate of the transmission link is equal to 24×64 kbit/sec (= 1.536 Mbit/sec), plus 8 kbit/sec (the signaling channel), for a total of 1.544 Mbit/sec.

An alternative, hypothetical transmission system format, such as shown in Figure 1.8, might again use repetitive switching frames, each subdivided into time slots containing multiple bytes. Each assigned time slot might contain, for example, 53 bytes. Again, using a network scheduler, each requested connection would be assigned to a particular time slot of successive frames. In such a scheme, the 53-byte segments or cells corresponding to the various existing connections are said to be synchronously time multiplexed since the time slot assigned to a particular connection occurs regularly (synchronously), in the same position of




Figure 1.8. A hypothetical packet-interleaved transmission format. Here, a connection is identified either by a specifically numbered time slot reserved in each frame or by a connection identification number contained in the packet header.

successive frames. However, since each time slot can support a 53-byte segment, rather than a single individual byte, an alternative strategy becomes possible. Suppose that each 53-byte cell is divided into two fields, a data field containing user-supplied information and a header field containing network-supplied control information. Included in the header field is a number unique to each connection carried over the transmission link. For example, the data field might contain 48 bytes of user information, and the header might contain 5 bytes of control information, 3 of which are used to identify the connection. Then, it is no longer necessary to assign a given connection to a particular time slot within a switching frame; the time slot has become large enough to efficiently carry the information used to identify a given connection, and the need for both a synchronous time-multiplexed frame and a network time-slot scheduler disappears. The switch or demultiplexer terminating the transmission line would then “read” the connection identifier contained within the header, and make appropriate internal routing decisions in real time on each cell. The cells associated with a particular connection need no longer occur synchronously, but, rather, asynchronous patterns may develop in which the cells corresponding to a given connection are randomly interdispersed with the cells corresponding to other connections. In such a scheme, the transmission time slots occur regularly, or synchronously, but the payloads or cells associated with a particular connection occur asynchronously. In the above example, we have used parameters consistent with ATM, the emerging transport technology for broadband ISDN: 53-byte cells, each containing 48 bytes of “payload” plus a 5-byte cell header. Because cells corresponding to different connections occur asynchronously in time, and since no time slot scheduler is present, a terminating switch *must* contain smoothing buffers to temporarily store cells that simultaneously arrive on two or more inbound transmission links which are destined for the same outbound links, since only one cell can be present on any outbound link at any given time. Such competing cells are sequentially read onto the time slots of the outbound links.

Interest in asynchronous time multiplexing arises by virtue of the great variety of traffic types anticipated in the integrated broadband environment. With synchronous time multiplexing, each time slot of the switching frame supports a fixed data rate connection between two users, one transmitting and one receiving. This approach served well in an era when most telecommunication traffic was telephone-generated voice or modem-produced voice-grade data. Telephone traffic exhibits three unique characteristics. First, each telephone is involved in one and only one conversation at a time. Most telephone conversations involve two participants, and even when multiple participants are involved in a conference call, there is only one “conversation” in progress; a single telephone will not permit an individual to “time share” himself or herself among two or more conversations involving different groups of participants (such “time sharing” would, if possible, be quite unnatural for most of us). In somewhat more technical terms, the tele-

phone does not support multiple simultaneous connections, a limitation completely compatible with natural human dialogue. Second, telephone traffic is *persistent*, that is, the typical conversation does not involve long idle periods punctuated by an occasional outburst of speech. Rather, one or more parties to a telephone conversation are usually speaking. Third, from a cold, technical perspective, all voice conversations look the same: continuous 64 kbit/sec information streams for digital voice (the telephone does not distinguish a beautiful, inspirational poem from the sound of a jackhammer breaking through the pavement!). Today's telecommunication infrastructure has evolved to reflect these unique properties of voice traffic: synchronous time multiplexing provides a persistent standard data rate connection between two end points. Where possible, data traffic (which has distinctly different characteristics from voice) is made to look like voice by means of a conditioning data modem. The volume of data handled in this fashion is sufficiently low that it can readily be carried over a network initially intended to support voice, although the resulting cost to the data user is unnecessarily high for the relatively low amount of bursty traffic carried over the network (data bursts are nonpersistent, but the network connection is nonetheless maintained even during long idle periods).

By contrast, the variety of traffic types to be supported in the broadband telecommunication environment mandates a different format. Some types of traffic are persistent, others are not. Among the persistent types (voice, full-motion video), the required connection data rates are vastly different (64 kbit/sec for uncompressed voice; greater than 100 Mbit/sec for uncompressed video, although codecs might compress this to 45 Mbit/sec, 1.5 Mbit/sec, or even lower, sacrificing quality and escalating cost along the way). The peak data rates required by nonpersistent types of traffic (typically called bursty traffic with connections characterized by occasional, randomly occurring short periods of activity) span several orders of magnitude. Telemetry traffic (temperature sensors, security monitors, etc.) may require a peak data rate during their active periods ranging from a few bits per second up to several hundred bits per second. Data rates for connections from remote terminals to a time-shared host range from a few hundred bits per second up to several hundred kilobits per second. Facsimile transmission may require anywhere from a few kilobits per second up to the megabit per second range, and interactions between mainframe computers may require peak data rates ranging from several tens of kilobits per second up to as high as 1 Gbit/sec. High-resolution image storage, retrieval, and display systems may require peak rates exceeding one Gbit/sec.

Moreover, in addition to spanning a peak rate range greater than nine orders of magnitude, the sources on nonpersistent traffic are characterized by the need to maintain multiple simultaneous connections from each terminating device. For example, unlike the telephone, a mainframe computer may need to concurrently maintain time-shared dialogues with many remote terminals and several other

mainframes with which it is sharing tasks, memory, and data files. The frequency with which such a mainframe may generate time-multiplexed bursts of data (data packet), each intended for a different receiver with which it is in active dialogue, may be quite high (hundreds of kilopackets per second). To accommodate the needs of bursty data sources, two alternatives may be considered. First, a persistent, synchronous connection of the appropriate data rate can be established between each data device and each other device to which it is in active dialogue. Considering the range of data rates involved, and the infrequency of activity for each separate dialogue, this is not a practical solution. The second approach requires network equipment which can read packet headers in real time to effect the correct packet relaying decisions at each of the possibly several intermediate packet switching nodes along the various connection routes.

Carrying the packet switching approach one step further, the actual format of each packet, at the network transport level, can be standardized; this is the basis of ATM. A standard cell size is defined. Its format consists of a fixed-length data field and fixed-length header. The header contains, at a minimum, sufficient information to allow the network to deliver the packet to the correct receiver and for the receiver to identify the source. Typically, an activity bit is also included (or something equivalent) such that the network knows if the time slot contains an active or empty cell. Large data packets are divided or segmented into multiple fixed-size cells, each with the appropriate routing header. As shown in Figure 1.9a, the destination address (or its equivalent in the form of a connection identifier) must be carried in each segment produced from a given packet. Since the large packet undergoing segmentation may not contain an integral number of cells, the last cell produced from the packet may be only partially filled. Similarly, packets smaller than the cell size will only partially fill one cell. All cells generated by one data device will be discharged onto the network through a network access station at a peak data rate sufficiently high to satisfy peak rate needed by the highest speed packet generated anywhere within the network (alternatively, the access speed for all cells will be selected and standardized at a rate sufficiently high to satisfy most packets; those packets requesting a higher peak rate must either "throttle back" or use specifically provisioned facilities, or else not be supportable over the network). Once inside the network, cells from multiple sources are statistically multiplexed onto network-internal transmission facilities and relayed along their correct routes by packet switching nodes.

At the receiving end, a data device accepts the time-multiplexed cells produced by those sources with which it is in active dialogue, and reassembles the packets from which the cells were created. The cell header field is used to deliver each cell to its intended receiver. Since cells corresponding to different sources (and, therefore, different packets) may arrive intermingled in time, the receiver must sort the cells on the basis of the originating sender, as shown in Figure 1.9b. Thus, in general, each receiver must be capable of reassembling several packets

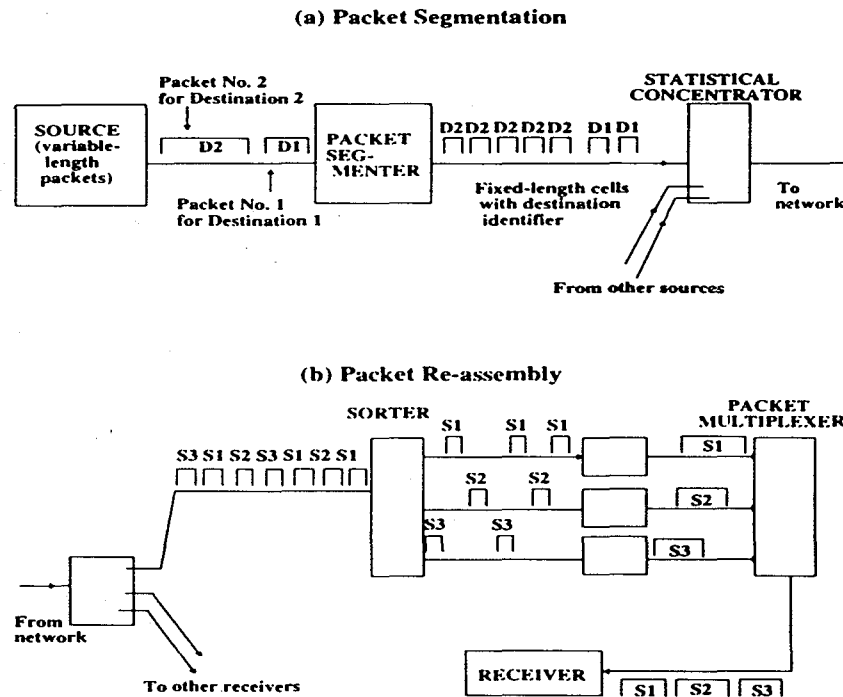


Figure 1.9. An example of packet segmentation and reassembly for Asynchronous Time Multiplexing.

in parallel, adding to one of the packets being reassembled each time a cell corresponding to that packet is received. Clearly, this sorting operation will require that each cell header must contain not only destination information but also source information or other information from which the identity of the cell sender can be deduced by the receiver. Reassembled packets are then time-sequenced for delivery to the the terminating device.

A possible format for the data cell is shown in Figure 1.10. We see that to enable packet segmentation, each cell corresponding to a given packet must contain a destination field, and to permit packet reassembly at the receiver, a source field must also be provided. In addition, the activity bit is useful to the network switches for identifying those time slots containing active cells. The activity bit and source and destination identifiers can be explicitly carried as shown in Figure 1.10 or, as an alternative, a connection identifier can be assigned at call setup time which uniquely identifies the sender and receiver. Each cell created from packets corresponding to a given connection will bear the same connection identifier, which can then be used by the network switching, multiplexing, and demultiplexing equipment for relaying purposes (remember that the connection identifier uniquely identifies the receiver) and by the receiver for reassembly purposes (remember that the connection identifier uniquely identifies the source). It is this latter strategy of assigning a connection identifier to each active con-

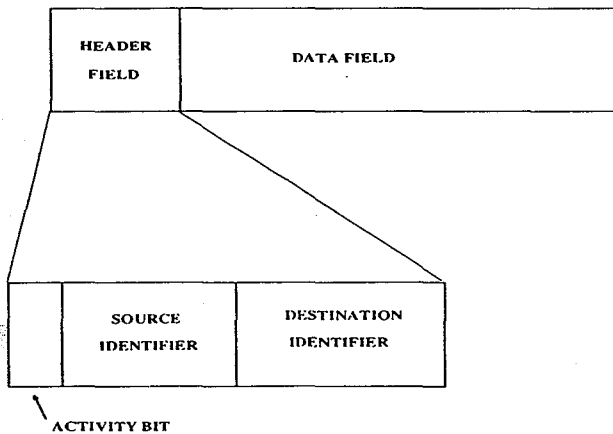


Figure 1.10. A possible format for the information cell or packet. (Note: The header field may not explicitly express source, destination, and activity, but contain other information from which these can be derived in real time using custom VLSI.)

nection which has been adopted for ATM; by using connection identifiers rather than explicit source/destination fields in the cell headers, it becomes possible for the network to support an arbitrarily large number of connections, even though the connection identifier field is of constant length (we shall see later that connection identifiers can be reused in geographically disjoint portions of the network). Also, as we shall see, use of a connection identifier greatly facilitates congestion control in a broadband network through a procedure known as admission control.

Finally, persistent types of traffic can be carried over the asynchronously time-multiplexed network by segmenting the traffic streams produced by the persistent sources into fixed-length cells, each with an appropriate header to permit relaying and reassembly; the segmentation and reassembly process are virtually identical to those needed for variable-length data packets.

* Thus, we see that asynchronous time multiplexing of small, self-routing cells provides a “universal” mechanism for integrating and transporting both persistent and nonpersistent traffic, and both low and high peak rates, over a common telecommunication network. The network does not know, or care, about the “payload” of each cell, nor does the network know or care whether a given cell contains voice, video, image, or computer data information. The network cares *only* about the cell header, which contains the route and reassembly instructions.

1.7. A Brief Word about Communication Protocols, Bridges, Routers, Gateways, Datagrams, and Virtual Circuits

Communications protocols are the sets of rules which users have agreed to follow when establishing services and transferring information. In addition to permitting the establishment and management of connections and keeping track of the multitude of connections which may be multiplexed through a common

physical network port, protocols are also needed to enable reliable communications in a relatively hostile telecommunication environment characterized by noisy transmission and network congestion. These impairments cause the appearance of random bit errors in delivered data, loss of segments of information from a delivered data packet or stream, and misrouting of information to unintended receivers. To achieve the desired reliability, protocols check for errors in delivered information, search for missing segments, and control the rate of information flow to contain congestion at some acceptably low level.

Unfortunately, not all vendors of information-related equipment have embraced one common agreed-upon set of communication protocols. As a result, the rules followed by one type of terminating equipment may differ from those followed by another. As a result, it is not always possible to establish a connection between any two devices.

A partial remedy for this situation was developed by the International Standards Organization (ISO). A model for communications, known as the Open Systems Interconnect (OSI) reference model, has been developed which describes the functions (but not their implementation) which are provided by communication protocols, and is shown in Figure 1.11. The model contains seven layers, with each layer providing service to the layer above it such that the higher layer can communicate with its "peer" in another termination in a manner which is transparent with regard to the communications processes occurring below it. For example, the N th layer in the protocol stack might be responsible for controlling the rate of information flow between the applications running on two terminals. To perform this function, the N layers of the two terminals trade real-time information between them concerning the amount of application information which has been sent, the amount received, and the additional amount which can be sent before a pause is needed to enable absorption and processing of the information. The N -layer communications assume a certain degree of integrity or quality of service




Figure 1.11. A representation of the ISO-OSI seven-layer protocol model.

with regard to the accuracy of the exchanged real-time information. It is the responsibility of layers $N-1, N-2, \dots, 2, 1$ to provide the required accuracy. Moreover, in performing its role, the N th layer is not interested in data rates on the physical links, voltage levels corresponding to logical zeros and logical ones, or whether the medium is copper, radio, or optical fiber. Again, it is the responsibility of layers $N-1, N-2, \dots, 2, 1$ to work out these details, so that layer N can do its job transparently with regard to that which is transpiring at the lower levels.

As shown in Figure 1.11, the seven-layer OSI model permits multiple applications to be multiplexed over one physical port; such multiplexing may be provided at layer 4 (transport layer). Among other things, it is the responsibility of layer 4 to merge information from different applications, apply headers to enable the applications to be distinguished, and route time-multiplexed arrivals to the correct terminating application. The OSI model admits sublayers for each layer contained in the model.

A much more thorough and complete description of the OSI model can be found in any of several excellent texts.

As is also shown in Figure 1.11, information generated by one transmitting “application” may need to be transported through several subnetworks to reach the intended receiver. These subnetworks may use different physical media, different transmission formats, different routing headers, different flow control strategies, etc. Whenever information leaves one “subnetwork,” it may become necessary to “close” its communications protocol to a certain layer in the protocol stack, and to “reopen” the protocol at that layer in the new “subnetwork.” For example, if subnetwork A uses a particular type of N -layer flow control protocol, and subnetwork B uses another, then the interface between the two subnetworks must (1) maintain the dialogue with the application on A using protocol A up to and including layer N , while (2) maintaining a dialogue with the application B using protocol B up to and including layer N . It is assumed, here, that the protocols are compatible at layers $(N + 1)$ and higher so that, again, the protocol conversions provided by layers N and below are totally transparent to layers $(N + 1)$ and higher, allowing the $(N + 1)$ st and higher layers from the terminating devices to communicate “peer to peer.” The layer, N , at which the interface between two subnetworks must “close” the protocol defines the type of interconnecting device required.

A *bridge* is a unit that interconnects two subnetworks that use a common Link Layer Control (LLC) procedure, but may use different Media Access Control (MAC) procedures. The interface between the LLC and the MAC occurs within layer 2 of the OSI model. The LLC sublayer consists of data packets fully formatted with prescribed header fields and data fields for transmission over a network, in need *only* of being placed onto the physical medium. The MAC sublayer is responsible for securing time on a shared transmission facility, and the physical layer (layer 1) is responsible for bit timing, voltage levels, clock recovery,

etc. An example of two different MAC procedures, both of which support a common LLC procedure, are the IEEE 802.3 standard for Carrier-Service-Multiple-Access with Collision Detection (CSMA/CD), used on a shared bus LAN, and the IEEE 802.5 standard for a token-based ring LAN; more will be said of these later. Two compatible applications, using identical communication protocols at the LLC sublayer and above, but with one resident on a CSMA/CD bus LAN and the other on a token ring LAN, can communicate if the two LANs are linked by a CSMA/CD-to-token passing bridge. The functions required of a bridge may be sufficiently simple so that real-time implementation in VLSI is possible.

Similarly, a *router* is a unit that interconnects two networks sharing common protocols at the network layer (layer 3) and above, but possibly using different data link layers (layer 2). Thus, for example, the format of the LLCs for the two networks may differ in some portion of the header. The router would then be responsible for closing the protocol at layer 2 on one network and reopening at layer 2 on the other, thereby enabling two compatible applications using a common communication protocol at layer 3 and higher, but resident on networks differing at layer 2 and below, to communicate. A router, being more complicated than a bridge, may require some degree of software processing, thereby presenting more of a bottleneck to the rate of exchange of information or information throughput. Finally, a *gateway* is a unit that interconnects two networks which differ at layer 3 or above. Since the two networks may differ all the way up to and including layer 7, a gateway may be required to close/open the protocol for each network through the application layer (layer 7), a complicated task requiring non-real-time software and often causing significant delay and degraded information throughput.

Among their other virtues, broadband networks seek to avoid some of the issues associated with communication protocols. First, the need for much of the protection provided by communication protocols is reduced since bandwidth is abundant, congestion is less frequent, and low-noise transmission facilities are used. In the nonbroadband environment, these issues are so serious that the same type of protection is often provided via different strategies at different layers of the protocol stack, e.g., error detection may occur at layer 2, guaranteeing a certain (low) probability of undetected error to layer 3, which may have its own error control procedures to further reduce the undetected error probability to layer 4, and so on. By circumventing these impairments, broadband telecommunications may permit very meaningful simplification of communication protocols.

Second, broadband networks are generally defined and standardized below layer 1.5 (the MAC layer arising with OSI layer 2). The broadband network is transparent to the higher-layer protocols which may be in use, and can be used to transparently transport information between two end-point applications which have agreed to use the same protocols at layer 1.5 and above. It is the responsibility of the end-point pairs to agree on use of compatible protocols at layer 1.5 and

above, or to use the services of a third end point to provide any necessary protocol conversion at or above layer 1.5; the transport network *per se* is not responsible for bridges, routers, and gateways, although the protocol translation services normally provided by a bridge, router, or gateway module may be offered by connecting an appropriate module to one of the transport network's access ports where the unit would appear to the network like any other terminal. This transparency is achieved by means of an access port adaptation layer, responsible for converting the format of the user-supplied information into the universal format required by the network. There are as many different types of adaptation layers appearing among the network's access ports as there are different types of protocols chosen by the users. More will be said later about the adaptation layer. In this way, each pair of end-users is free to use any protocol at layer 2.0 and above, as long as the needs of these protocols are met by the grade of service provided by the broadband network universal interface. Alternatively, each user may use some standard protocol, and enjoy connectivity with another user which uses a different protocol by requesting a connection via a router or gateway. The router or gateway service might be an offering of the network provider or might be offered by a "fourth party" (the three participants always present are the two users and the network service provider). The participants may all belong to a common enterprise, with network transport provided by a separate enterprise (possibly a public network); the users and router/gateway provider may belong to a common enterprise different from that of the network transport provider (virtual private network); all participants may belong to different enterprises, etc.

One final word about communication protocols. These are often classified by the descriptors "connection-oriented" and "connectionless." Connection-oriented protocols require a call setup procedure, despite the fact that information flows in packet format with header fields containing the routing information. The call setup procedure selects a path or route to be used by all packets associated with a given connection, and the traffic intensity appearing on each physical network link is controlled by limiting the number of connections sharing that link. Path selection statistically spreads the total applied load equitably among all links and packet switching nodes, generally trying to avoid congestion. Since the statistical load to be presented by each requested connection is known by the call processor, the path can be chosen to avoid unacceptable degradation to the quality of service previously guaranteed to earlier connections; if a suitable path cannot be found, the connection can be blocked (similar to a "busy signal" in today's voice network). If a new connection is admitted, a "virtual connection" number is assigned to that connection, and appears in the header field of all packets belonging to that connection. The virtual connection number implicitly identifies both the source and destination for each packet upon call establishment. Each switch along the selected path is informed of the assigned virtual connection number, and is provided with routing instructions to be followed whenever a packet containing

that virtual connection number arrives. The connection is termed "virtual" because, unlike synchronous time multiplexing, no network resources (e.g., time slots) are reserved for the exclusive use of that connection. Rather, the user perceives the appearance of a connection, along with many of the services expected from a connection (e.g., packets arrive at the receiver in the sequence generated), but the network resources are, in reality, statistically shared among some multitude of connections. The connection is, therefore, "virtual." Among the services guaranteed by a real connection is bounded latency or delay in the network; a virtual channel cannot guarantee bounded latency since the network load is statistical but, through a combination of admission control, flow control, congestion control, and virtual connection prioritization, it can provide guarantees such as "no more than X% of packets will be delivered with a delay in excess of Y milliseconds." The values for X and Y can be selected by the user at call setup time, depending on willingness to pay the presumed greater charges for higher service quality.

The three phases associated with connection-oriented protocols are (1) call establishment, (2) information transfer, and (3) call release. Call establishment and information transfer procedures are illustrated in Figure 1.12a-c; call release procedures are similar to those used for call establishment and are not illustrated. Shown in Figure 1.12 are generic "access ports" (open circles), switching nodes (closed circles), and interconnecting transmission links. Each access port has a permanent virtual channel number assigned for communication with the call processor, which is attached to the transport network through an ordinary access port and appears to the transport network as any other user or application; call processing may, in fact, be the shared responsibility of processors attached to geographically distributed access ports but, for simplicity, we shall assume the existence of one centralized processor when explaining virtual connection procedures. Using this permanent virtual connection, a user may request a connection to some desired destination; the requested connection might be two-way to enable full duplex operation. The call processor, using a permanent virtual connection from itself to the destination, asks the destination if it wishes to accept the connection. If so, the call processor will attempt to find a path which, when loaded with an additional virtual connection, will not cause the quality of service enjoyed by other connections to fall below the guaranteed minimum level; if the destination refuses to accept the connection, or if a suitable path cannot be found, the connection is blocked and the originator is so informed via a permanent virtual connection from the call processor. If the connection can be established, all switching nodes along the selected path are informed by the call processor of the new virtual connection number, and are provided with appropriate routing instructions; permanent virtual connections from the call processor to the switching nodes are used for this purpose. During the information transfer phase, the call processor drops out of the picture and the enabled access ports exchange informa-

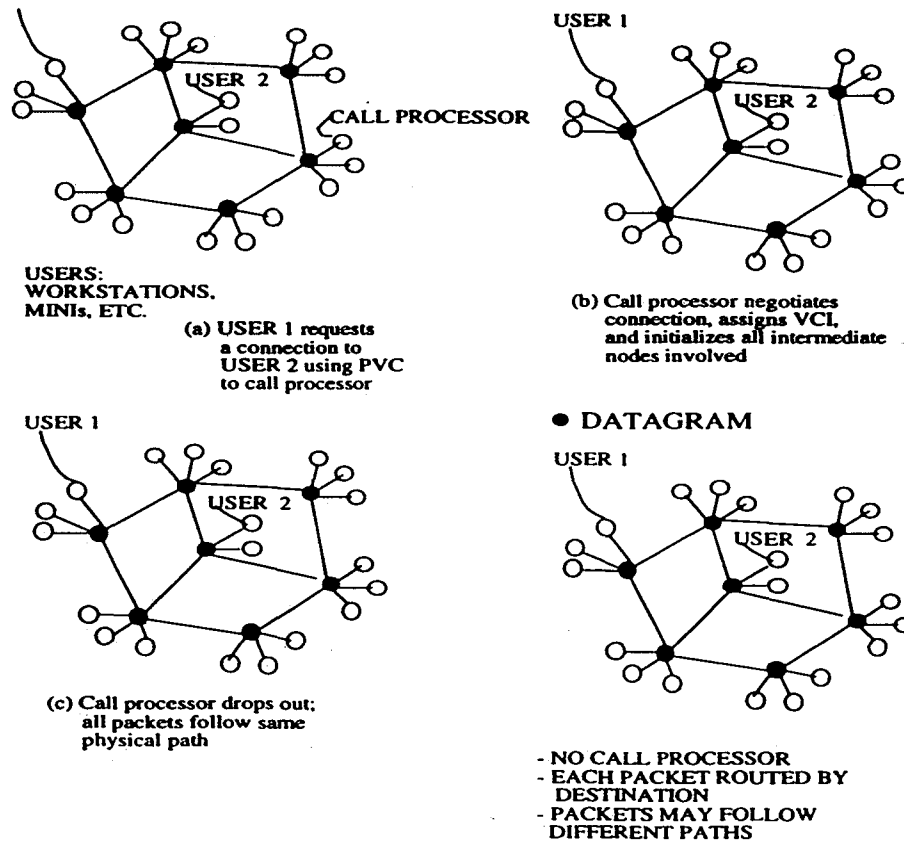


Figure 1.12. Information transfer via virtual circuits and datagrams.

tion over the assigned path (or paths, for for a duplex connection). Each packet contains the virtual connection number in its header, and only the header is processed (in real time, with custom VLSI circuitry) by the switching nodes to effect routing decisions. Since all packets affiliated with a given virtual connection follow the same route, they are delivered in the same sequence in which they were generated. A process similar to call initiation is used when either access port wishes to disconnect; again, permanent virtual connections are used to signal between the access ports and call processor, and between the call processor and the switching nodes.

Connectionless protocols are somewhat more basic. No call setup procedure or call processor is involved, each packet is treated by the transport network as an independent entity, and each packet contains source and destination information within its header. The destination field is used by each encountered switching node to effect a routing decision. However, since more than one switch output port may permit the packet to ultimately arrive at its correct destination, the decision of the

switch might be influenced by local congestion conditions. For example, if the switch has routed more packets (which may be associated with different access ports) to one of two outputs acceptable for some given packet, then the switch may choose to select the relatively underutilized output for that given packet. In this fashion, packets will eventually arrive at their intended destinations, but the delivery sequence may be different from that with which the packets were generated by the source since, for a given source-destination pair, not all packets will traverse the same path through the network. Also, since there are no call setup procedures, admission control cannot be used and alternative techniques may be needed to guarantee quality of service.

Transport of information in broadband networks tends to be via connection-oriented protocols since a wide variety of traffic types are to be integrated and carried. Since quality-of-service needs may be different for different traffic types, traffic control is more difficult than for homogeneous data-only types of networks where the simplicity of connectionless protocols can be exploited. In the broadband domain, all traffic control techniques, including admission control at call setup time, must be invoked to permit traffic integration since, for example, the resources needed by a bandwidth-intensive image or full-motion video dialogue between two end points might adversely affect the quality of service expected by hundreds of lower-speed data users. A network manager is needed to ensure fairness in the use of network resources, given the enormous differences among the types of traffic to be integrated onto the broadband network.

The fact that the transport network employs connection-oriented protocols does *not* imply that specific end-users cannot perceive a connectionless service. For example, many connectionless LANs can be interconnected by means of a connection-oriented broadband network. To provide such a service, the LAN access ports to the broadband network might be fully interconnected over the broadband network by means of full-duplex permanent virtual connections. Then, a sequence of time-multiplexed connectionless packets (datagrams) generated by one LAN, each intended for a different receiving LAN, would arrive at the input to the broadband network which serves that LAN. Here, the packets would be demultiplexed and routed to their correct receivers by means of the permanent virtual connections among the LAN access ports. In this way, users attached to connectionless LANs could send datagrams to users attached to remote LANs without first invoking a call setup procedure.

1.8. Contention in Telecommunication Networks

Stripped to its barest essentials, a telecommunication network can be viewed as a multiport system responsible for (1) transmission of information over remote distances, (2) routing of information applied at the various input ports to the

correct output ports, and (3) resolution of any contention that may arise when two or more information processes simultaneously seek to use some common facility.

Application of information to a telecommunication network may be viewed as a sequence of random, uncoordinated events. An attempt to initiate a single telephone connection between two users, say, in New York and Los Angeles typically occurs without the coordination or the cooperation of two other users who are also attempting a telephone connection between, say, Philadelphia and San Diego, although both connections may seek to use a common switch in Omaha. Thus, contention may arise at the Omaha switch. Similarly, two computers may simultaneously seek to send a datagram to the same destination computer, in which case contention arises over the use of the network output port serving that destination computer.

These two examples serve to illustrate the two types of contention which may arise in telecommunications. For the first, some multitude of requests may contend for a common or intersecting set of network resources (e.g., transmission, routing, storage facilities). Such contention affects the quality of service seen by the end-user, relative to that which would be experienced if adequate resources were always available to satisfy all requests, since some requested connections may not be satisfied or information delivery may be delayed. The effects of this type of contention on the quality of service experienced by the network user can be controlled by network design: deployment of additional network resources to satisfy the anticipated traffic load will provide quality of service equal to the design objective. However, economic considerations mitigate against deployment of resources sufficient to meet every possible pattern of requests. A network designed, for example, to handle the demand corresponding to *every* telephone simultaneously requesting a connection to *any other* telephone would require a prohibitively large quantity of telecommunication equipment, most of which would sit idle during normal, less stressing conditions. Conversely, a network designed to handle an "average" traffic load might perform poorly when slightly stressed. Thus, the amount of equipment deployed is chosen to provide some desired quality of service under greater-than-expected, but not worst-case, traffic loading conditions.

The second type of contention is more fundamental. Here, multiple information sources simultaneously try to access a common network output port, and quality of service cannot be improved by network design alone (additional telecommunication equipment will not prevent a busy signal when a parent tries to dial home and one family member is a talkative teenager). A goal of network design, then, is to achieve a quality of service no poorer (or only slightly poorer) than that produced by unavoidable "output contention." The network should be designed such that the effects of network-internal congestion are small compared with the effects of output contention, and perceived quality of service should be dominated by output contention.

Although output contention is unavoidable, there are some limited steps that

the network might take to offset its effects. Using the "talkative teenager" example once again, the network might, as a service, alert the teenager that a second connection is being attempted; should the teenager wish to respond, the first connection is placed on hold, and the teenager then decides whether or not to accept the parent's call. As a second example, the network can be instructed to deliver contending packets to a common receiver in a sequence which respects the priorities of the contending packets for real-time delivery. While not avoiding output contention, a combination of network design and enhanced services can often soften its effect on service quality.

While network-internal contention can be controlled by design and output port contention is unavoidable, neither type of contention can be ignored, and the network must resolve contention however it may arise. There are only two ways of resolving contention in a telecommunication network. The first technique involves reservation: at the time that a connection is requested, the network "manager" determines whether or not network resources exist to be assigned for the exclusive use of that connection, and whether or not the output port can accept the connection. For example, the network manager may seek a circuit on the various synchronously time-multiplexed transmission systems separating the input and output ports (both directions for a full-duplex connection). Among the many paths that the connection might take, the network manager needs to find only one with the required resources available on each link. If the resources are available, and the output port can accept the connection, then the network manager will make an allocation of the identified network resources for the *exclusive* use of the new connection for the entire duration of the connection. Since resources have been allocated on an exclusive basis, no subsequent requests can interfere with the newly established connection. If the output port cannot accept the connection, or if available network resources cannot be allocated for the exclusive use of the requested connection, then the requesting party will be informed that the connection cannot be made (in telephony, this would correspond to a busy signal). This first approach, in which network resources are *reserved* for the exclusive use of the new connection, is referred to as circuit switching, and quality of service is typically defined by the blocking probability, i.e., the likelihood that a newly requested connection cannot be made.

For a second approach, the network does not seek to allocate resources on an exclusive basis. Rather, the network accepts all applied input signals, temporarily storing or delaying their delivery until the resources required to do so become available, or until the output port becomes free. This second approach is dependent on having sufficient storage capability within the network to "smooth" over the statistical variations in arrivals from the various input ports. Here, the network has been designed with adequate resources such that under stressful (but again, not worst case) traffic loading, the typical delay suffered by applied messages is acceptably small most of the time, and the fraction of messages "lost" by the network whenever the information to be stored exceeds the storage capacity

is acceptably low. This type of contention resolution is known as packet switching, and the quality of service is typically defined by several parameters: average message delay, fraction of messages with delay greater than some extreme limit of acceptability, and probability of message loss. For connectionless (datagram) packet switching service, network scheduling such as is needed for circuit switching is avoided since all applied traffic is accepted, and quality of service is maintained through a combination of network design and internal management of congestion (i.e., the network manager, sensing that the instantaneous demand for an internal link is too high, may begin to alternately route some messages to their respective destinations over different paths). For connection-oriented (virtual connection) packet switching service, a scheduler is reintroduced, not to reserve resources on an exclusive-use basis but, rather, to determine whether the additional traffic represented by some newly requested connection will statistically overload the network, thereby causing a degradation in the quality of service enjoyed by previously established connections. In a sense, within a virtual connection-oriented network, the network resources and the right to be connected to the desired output port are reserved on a statistical rather than on an exclusive-use basis: the network transmission links, switching nodes, and smoothing buffers are statistically shared among all admitted connections on an instantaneous, as-needed packet-by-packet basis.

1.9. Problems

- P1.1. A new transmission format uses a frame length of 1 msec which is subdivided into 100 time slots. Each time slot contains 100 bits. (a) What is the data rate of the transmission line? (b) If each time slot corresponds to one fixed-rate channel, what is the data rate of each such channel? (c) How many time slots must be assigned to a connection requiring a data rate of 2 Mbits/sec?
- P1.2. A set of black-and-white photographs is to be stored and delivered electronically. Each raster scan contains 100 horizontal lines, and each line is sampled 1000 times. Each sample is quantized to one of 256 gray levels. (a) How many bits of information are contained in each uncompressed image? (b) Assuming a browsing rate of 2 images/sec, a transmission rate of 1 Mbit/sec, and negligible time for image compression and expansion, by what factor must each image be compressed?
- P1.3. The packet format for some particular data network contains a connection number field of length equal to 20 bits. (a) Assuming that each packet is to be delivered to a single receiver, to how many distinct destinations can each transmitter address its packets? (b) Repeat part a, assuming that ten connection numbers are reserved for multicasting, and each multicast connection can involve, at most, 100 receivers.
- P1.4. A Time Slot Interchange (TSI) stores the information contained in each transmission frame arriving at its input such that an arbitrary permutation of the frame's time slots can be produced at the TSI's output. For example, if the frame contains

three time slots sequentially numbered 1,2,3, then, at the output of the TSI, it is possible to reorder the information contained within the time slots such that the output time slots contain, in time sequence, information from input time slots 3, 2, 1. TSIs can be used to resolve output port contention in time-multiplexed circuit switches. Consider a 3-input, 3-output switch. Each of the three input transmission frames contains 10 time slots. (a) Suppose the time slots for inputs 1, 2, and 3 are sequentially addressed to switch outputs as follows:

Input 1: 1, 2, 1, 1, 3, 3, 2, 3, 2, 2
 Input 2: 2, 2, 3, 3, 2, 1, 1, 2, 3, 3
 Input 3: 1, 2, 1, 1, 3, 3, 1, 1, 2, 3

Note that conflicts arise whenever two or more concurrent input time slots are to be switched to the same output. By means of three TSIs, resequence each of the inputs such that all conflicts are eliminated. (b) Next, consider the following three input sequences:

Input 1: 2, 3, 3, 2, 2, 3, 2, 1, 1, 2
 Input 2: 3, 3, 2, 2, 3, 1, 1, 2, 2, 2
 Input 3: 1, 3, 2, 2, 3, 1, 2, 3, 3, 2

Can these be resequenced to eliminate conflicts? If not, why not?

- P1.5. A data source produces packets, each with an average length of 1 kbyte. The source is active, on average, 10% of the time and, when active, delivers information at a peak rate of 10 Mbit/sec. (a) What is the average time duration of each packet? (b) What is the average data rate of the source? (c) How many sources can be statistically multiplexed onto a 5 Mbit/sec transmission link without causing the average traffic demand on that link to exceed its capacity?
- P1.6. A Metropolitan Area Network is used to interconnect several Local Area Networks. Traffic appearing on each LAN consists of a data packet containing 2 kbyte every 10 msec, and 20% of each LAN's locally generated traffic terminates at a remote LAN. Traffic patterns are uniform. Find the minimum data rate needed on the MAN such that 20 such LANs can be interconnected.
- P1.7. A statistical concentrator time-multiplexes packets arriving from three sources onto a common output link. At the concentrator, each source terminates at its own dedicated buffer. The three buffers are cyclically served in the sequence 1-2-3, and a buffer is skipped over if and only if it is empty. All packets are of fixed length, and packet arrivals from all sources have time-aligned boundaries. Packet arrivals are as shown in the following table.

Packet present on:

Time slot No.	Input 1	Input 2	Input 3
1	yes	no	yes
2	no	yes	yes

(cont.)

Packet present on:

Time slot No.	Input 1	Input 2	Input 3
3	no	no	yes
4	no	yes	no
5	no	no	no
6	yes	no	no
7	no	no	no
8	no	no	no
9	yes	yes	no
10	no	no	no
11	no	no	no

- Draw a time sequence showing the packets appearing on the output link, with each packet labeled by its source number. In the first time slot, the cyclical server is positioned to accept a packet, if present, from source number 1.
- P1.8. A retransmission protocol requires the receiver to acknowledge each arriving packet determined to be error free. The link bit error rate is 10^{-6} , each packet contains 1000 bits (including the error-detecting code bits), and the error-detecting code can detect all possible bit error patterns which may occur among the 1000 bits. In addition, the probability that an applied packet is lost by the network (e.g., buffer overflow) is 10^{-4} . Any packet not acknowledged within three transmission attempts is abandoned. Find the probability of abandonment.
- P1.9. A data communication protocol uses window flow control: each received packet is acknowledged and, for each packet source, at most five unacknowledged packets can be outstanding. The round-trip network propagation delay (elapsed time from packet generation to receipt of its acknowledgment) is 40 msec, each packet of length 100 bytes, and the transmission data rate is 56 kbit/sec. Each source immediately generates a new packet upon receipt of an acknowledgment. (a) Draw a timing diagram illustrating the time sequence of events (packet generation and acknowledgment) corresponding to the first 20 packets generated by a given source. (b) What is the average source throughput, that is, the average data rate in bits per second at which information is supplied by a given source? (c) What is the effective utilization efficiency of the transmission link? (d) Repeat c, but for a window size of 100 packets.
- P1.10. Repeat Problem 1.9a–d except that, now, let the transmission data rate be 800 Mbit/sec. (e) What is the window size needed to produce a utilization efficiency of 100%? (f) If forward-link congestion should develop at the receiver under the conditions of part (e), how much new information will enter the network before the interruption in acknowledgments is detected by the transmitter? (g) What conclusion do you draw concerning the usefulness of window flow control for broadband networks?
- P1.11. Consider an integrated services network in which voice and data connections share a common transmission link of capacity 45 Mbit/sec. Each voice connection

requires a dedicated capacity of 64 kbit/sec and each data connection generates fixed-length 1-kbyte packets at an average rate of 100 packets/sec. To achieve its quality-of-service objective, each data connection requires an "equivalent" capacity which is three times as high as its average data rate. Write an inequality defining the transmission link capacity region, that is, defining the set of permissible combinations of voice and data connections which may simultaneously share the transmission link.

- P1.12. In some hypothetical integrated bandwidth-on-demand environment, each voice connection requires a capacity of 50 kbit/sec, each video connection requires 5 Mbit/sec, and each data connection produces a 100-msec burst of information at a peak data rate of 10 Mbit/sec every second. Plot the peak bandwidth needed between $t = 0$ and $t = 10$ seconds for an "information outlet" which supports the following traffic profile:

Time (sec)	# voice connections	# video connections	# data connections
0	0	0	0
1	1	0	0
2	2	0	0
3	2	1	0
4	2	1	1
5	2	1	1
6	2	2	1
7	1	2	1
8	1	1	1
9	0	0	1
10	0	0	0

- P1.13. A high-speed network containing 500,000 user access ports has a per-port capacity of 1 Mbit/sec and an aggregate capacity of 100 Gbit/sec. To support the network fault diagnostic routine, each port generates a 1000-bit word every second, receipt of which must be acknowledged by the intended receiver by means of a 1000-bit status reply. Assuming normal operation, what fraction of aggregate network capacity is committed to fault diagnostics?
- P1.14. A data link operating at a rate of 56 kbit/sec exhibits a bit error probability of 10^{-6} . Each data packet contains 1 kbyte, including an error-detecting field capable of detecting any pattern of bit errors. Each time a packet is received containing a detected error, the protocol requires that it be retransmitted along with each of the nine prior packets (Go-back-10 retransmission scheme). What is the maximum link rate of information transfer across the data link?
- P1.15. An entertainment video signal is developed by means of a raster scan containing 500 scan lines for each of three primary colors. The entire scan is repeated every 10 msec, and each scan line is sampled 500 times. Each sample is quantized to one of 256 voltage levels. Assuming a video codec which removes signal redundancy

and achieves a compression factor of four, find the data rate required to transmit the resulting digital video signal (ignore the associated audio signals).

- P1.16. Consider the video signal of Problem 1.15, except that, now, the codec achieves a compression factor of 100. A video conference call uses such a codec to compress the signal generated at each of 10 sites, and each site is equipped with nine monitors such that it can concurrently view all the remote sites. (a) What is the total traffic demand, in bits per second, generated by the video conference? (b) What is the required receiving capacity per conference site (again, ignore the associated audio signals)?
- P1.17. Consider a multiprocessor containing 100 computing elements. When active, each computing element generates information at a peak rate of 100 Mbit/sec. Suppose that synchronous time division-multiplexed circuit switching is used to fully interconnect the processing elements, and that each circuit operates at the per-computing element peak rate of information generation. (a) What is the required data rate of the transmission link serving each computing element? (b) Assuming that each computing element sends to only one other at any given time, what is the maximum link utilization efficiency?

References

1. R. W. Lucky, *Silicon Dreams: Information, Man, and Machine*, St. Martin's Press, New York, 1989.
2. R. L. Geiger, P. E. Allen, and N. R. Stoader, *VLSI Design Techniques for Analog and Digital Circuitry*, McGraw-Hill, New York, 1990.
3. D. A. Hodges and H. G. Jackson, *Analysis and Design of Digital Integrated Circuits*, 2nd ed., McGraw-Hill, New York, 1988.
4. O. Wing, *Gallium Arsenide Digital Circuits*, Kluwer Academic, 1990.
5. A. Tanenbaum, *Computer Networks*, 2nd ed., Prentice-Hall, Englewood Cliffs, N.J., 1988.
6. T. Li, Advances in optical fiber communications: An historical perspective, *IEEE J. Selected Areas Commun. SAC-1*, April 1983.
7. P. S. Henry, Introduction to lightwave transmission, *IEEE Commun. Mag.* 23(5), May 1985.
8. S. E. Miller and I. Kaminow, *Optical Fiber Telecommunications II*, Academic Press, New York, 1988.
9. E. Desurvire, Lightwave communications: The fifth generation, *Sci. Am.* Jan. 1992.
10. C. A. Brackett, Dense wavelength division multiplexing networks: Principles and applications, *IEEE J. Selected Areas Commun. SAC-8*(6), Aug. 1990.
11. K. Murakami and M. Katoh, Control architecture for next generation communication networks based on distributed databases, *IEEE J. Selected Areas Commun. SAC-7*(3), April 1989.
12. J. J. Garrahan, P. A. Russo, K. Kitami, and R. Kung, Intelligent network overview, *IEEE Commun. Mag.* 31(3), March 1993.
13. S. R. Schach, *Software Engineering*, Aksen Associates.
14. F. Halsall, *Data Communications, Computer Networks, and OSI*, Addison-Wesley, Reading, Mass., 1988.
15. Members of the Technical Staff, Bell Telephone Laboratories, *Transmission Systems for Communications*, Bell Telephone Laboratories, Inc., 1982.
16. *IEEE Trans. Commun.*, special issue on Digital Switching, COM-27(7), July 1979.
17. *AT&T Tech. J.*, special issue on the 5ESS™ Switch, 64(6, Part 2), July-Aug. 1985.

18. A. R. Modarressi and R. A. Skoog, Signaling system No. 7: A tutorial, *IEEE Commun. Mag.* **28**(7), July 1990.
19. J.-Y. Le Boudec, The asynchronous transfer mode: A tutorial, *Comput. Net. ISDN Syst.* **24**, 1992.
20. A. S. Acampora and M. J. Karol, An overview of lightwave packet networks, *IEEE Net. Mag.* **3**(1), Jan. 1989.
21. C. D. Tsao, A local area network architecture overview, *IEEE Commun. Mag.* **22**(8), Aug. 1984.
22. *IEEE Commun. Mag.*, special issue on Metropolitan Area Networks, April 1988.
23. *IEEE Commun. Mag.*, special issue on Multimedia Communications, **30**(5), May 1992.
24. *IEEE J. Selected Areas Commun.*, issue on Medical Communications, **10**(7), Sept. 1992.
25. *IEEE J. Selected Areas Commun.*, issue on Multimedia Communications, **8**(3), April 1990.
26. *IEEE J. Selected Areas Commun.*, issues on Packet Speech and Video **7**(5), June 1989.
27. J. Walrand, *Communication Networks: A First Course*, Aksen Associates, 1991.
28. W. Stallings, *Data and Computer Communications*, 2nd ed., Macmillan Co., New York, 1988.
29. M. DePrycker, ATM switching on demand, *IEEE Net. Mag.* **6**(2), March 1992.
30. *Switching Systems*, American Telephone and Telegraph Co., 1961.
31. J. Auerbach, TACT: A protocol conversion toolkit, *IEEE J. Selected Areas Commun.* **SAG-8**(1), Jan. 1990.
32. E. Ball, N. Linge, P. Kammar, and R. Tasker, Local area network bridges, *Comput. Commun.* **11**(3), June 1988.
33. C. A. Sunshine, Network interconnection and gateways, *IEEE J. Selected Areas Commun.* **SAC-8**(1), Jan. 1990.
34. C. V. Bachman and P. Mondain-Monval, Design principles for communication gateways, *IEEE J. Selected Areas Commun.* **SAC-8**(1), Jan. 1990.
35. L. Roberts, The evolution of packet switching, *Proc. IEEE* **6**(11), 1978.
36. J. D. Day and H. Zimmerman, The OSI reference model, *Proc. IEEE* **71**(12), Dec. 1983.
37. W. F. Emmons and A. S. Chandler, OSI session layer: Services and protocols, *Proc. IEEE* **71**(12), Dec. 1983.
38. L. L. Hollis, OSI presentation layer activities, *Proc. IEEE* **71**(12), Dec. 1983.
39. J. C. McDonald, ed., *Fundamentals of Digital Switching*, Plenum Press, New York, 1983.
40. M. M. Buchner, Jr., M. Iwama, T. J. Herr, R. E. Staekler, and L. J. Gitten, Evolution to a universal information services environment, ISSLS '86, Tokyo.
41. M. Decina, Evolution toward digital wideband communication networks, 1987 Int. Switching Symp., Phoenix.
42. P. E. White, The role of the broadband integrated services digital network, *IEEE Commun. Mag.* **29**(3), March 1991.

