


SECOND EDITION

ATM

VOLUME I

Foundation for Broadband Networks



UYLESS BLACK



**PRENTICE HALL SERIES
IN ADVANCED COMMUNICATIONS TECHNOLOGIES**

IPR2023-00714
Cisco Systems, Inc. v. Orckit IP, LLC.
Cisco's Exhibit 1009
Ex. 1009, Page 1


ATM FOUNDATION FOR BROADBAND NETWORKS

VOLUME 1, SECOND EDITION





Prentice Hall Series In Advanced Communications Technologies





ATM FOUNDATION FOR BROADBAND NETWORKS

Volume I, Second Edition

UYLESS BLACK



Prentice Hall PTR
Upper Saddle River, New Jersey 07458
<http://www.phptr.com>

Library of Congress Cataloging-in-Publication Data

Black, Uyless D.

ATM—foundation for broadband networks, 2/E / Uyless Black.

p. cm.

Includes bibliographical references and index.

ISBN 0-13-083218-9

1. Asynchronous transfer mode. 2. Broadband communication systems. I. Title.

TK5105.35.B53 1995

621.382—dc20

95-5961

CIP

Acquisitions editor: Mary Franz

Cover designer: Talar Agasyan

Cover design director: Jerry Votta

Manufacturing manager: Alexis R. Heydt

Marketing manager: Miles Williams

Compositor/Production services: Pine Tree Composition, Inc.



© 1999 by Uyless Black

Published by Prentice Hall PTR

Prentice-Hall, Inc.

A Simon & Schuster Company

Upper Saddle River, New Jersey 07458

Prentice Hall books are widely used by corporations and government agencies for training, marketing, and resale.

The publisher offers discounts on this book when ordered in bulk quantities. For more information contact:

Corporate Sales Department

Phone: 800-382-3419

Fax: 201-236-7141

E-mail: corpsales@prenhall.com

Or write:

Prentice Hall PTR

Corp. Sales Dept.

One Lake Street

Upper Saddle River, New Jersey 07458

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

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN: 0-13-083218-9

Prentice-Hall International (UK) Limited, *London*

Prentice-Hall of Australia Pty. Limited, *Sydney*

Prentice-Hall Canada Inc., *Toronto*

Prentice-Hall Hispanoamericana, S.A., *Mexico*

Prentice-Hall of India Private Limited, *New Delhi*

Prentice-Hall of Japan, Inc., *Tokyo*

Simon & Schuster Asia Pte. Ltd., *Singapore*

Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*



I am told by one of my friends that a fictional creature in a novel was discussing the speed of the flight of birds and proclaimed that: perfect speed is being there.

While I do not know about the other thoughts of this creature, I have thought about this one idea many times. Aside from its physical and philosophical implications, it is appropriate to bring to mind that many of the emerging technologies (such as ATM) are attempting to close the gap (the latency) between being in one place, and then another. In this book, the issue is not the flight of birds but the movement of messages of information through a communications network.

Of course, with our current knowledge, we cannot achieve the bird's goal, due to the delay inherent in the physical aspects of travel. So, until we know about (and conquer) other dimensions (if we ever will), we are restricted to interpreting "being there" as the bird's physical flight from one point to another and its quest to make that flight of a shorter duration.

This goal has been an important aspect of communications networks since their inception, and the speed and efficiency of modern networks has surely had a profound effect on our personal and professional lives.

So, this book is dedicated to all of those in our industry who are striving to send the message faster.

At the same time, it is instructive to remember that a faster message does not necessarily make it any better.





Contents

PREFACE

xvii

CHAPTER 1	Introduction	1
	The Present Telecommunications Infrastructure	1
	Present Technologies for Voice, Video, and Data Networks	2
	Present and Future Requirements	5
	Downsizing and Outsourcing: Reliance on Telecommunications	5
	Present Systems: Too Much or Too Little	6
	Costs of Leased Lines	8
	Virtual Companies and Virtual Networks	8
	Fast Relay Networks and ATM	11
	Applications use of ATM	14
	Fast Relay Networks and SONET	15
	Broadband ISDN	16
	Principal specifications for atm	18
	The Anchorage Accord	20
	Summary	20

CHAPTER 2	The Nature of Analog: and Digital Systems	23
	Analog Systems 23	
	Cycles, Frequency, and Period 24	
	Bandwidth 25	
	Broadband and Baseband Signals 27	
	Other Definitions of Broadband 28	
	The Analog-to-Digital Conversion Process 28	
	Sampling, Quantizing, and Encoding 30	
	Other Coding Schemes ¹ 31	
	Timing and Synchronization in Digital Networks 32	
	Plesiochronous Networks 33	
	The Synchronous Clock Hierarchy 33	
	Clarification of Terms 34	
	Timing Variations 35	
	Slips—Controlled and Uncontrolled 35	
	Bit or Clock Slips 36	
	Summary 37	

CHAPTER 3	Layered Protocols the Model for ATM and SONET Networks	38
	Protocols and the OSI Model 38	
	OSI Layer Operations 40	
	Concept of a Service Provider 41	
	Encapsulation/Tunneling 45	
	ATM and the Model 47	
	Protocol Entities 47	
	Service Access Points (SAPs) 47	
	ATM and OSI Layers 48	
	The Internet Protocols (TCP/IP) 49	
	The Internet Layers 51	
	IP Functions 52	
	TCP Operations 55	
	The OSI Network and Transport Layer 56	
	Summary 57	

CHAPTER 4	Emerged Technologies	58
	Comparison of Switching Systems 58	
	The T1/E1 Systems 60	
	Purpose of T1 and E1 60	
	Typical Topology 61	

T1 and E1 Layers	62
T1/E1 PDUs	62
Conclusions on T1/E1	64
X.25	64
Purpose of X.25	64
Typical Topology	65
X.25 Layers	66
X.25 PDUs	66
Other Noteworthy Aspects of X.25	68
Conclusions on X.25	69
Integrated Services Digital Network (ISDN)	70
Purpose of ISDN	70
Typical Topology	70
ISDN Layers	73
ISDN PDUs	74
Conclusions on ISDN	78
Signaling System Number 7 (SS7)	79
Purpose of SS7	79
Typical Topology	79
SS7 Layers	81
SS7 PDUs	83
Conclusions on SS7	85
ATM and SONET: Reduction or Enhancement of Functions in Networks	85
Summary	85

CHAPTER 5	The Broadband Integrated Services Digital Network (B-ISDN) Model	87
	ISDN and B-ISDN	87
	B-ISDN Configurations	88
	ATM and the B-ISDN Model	89
	Examples of the Operations between Layers in the B-ISDN Planes	91
	B-ISDN Functions	93
	B-ISDN Service Aspects	95
	Summary	97

CHAPTER 6	Asynchronous Transfer Mode (ATM) Basics	98
	A Brief Review	98
	Why is ATM called "Asynchronous"?	99
	An ATM Topology	99

The ATM Interfaces	102
The VPI and VCI Labels	104
ATM Layers	106
ATM Layers and OSI Layers	107
Relationship of AAL, ATM, and the Network	109
Relationship of Layers to the OSI Layered Architecture	111
Where to Find Service Definitions and Primitives	112
Typical Protocol Stacks	112
ATM PDUS (cells)	114
Use of Two Identifiers	114
Metasignaling Cells and Other Cells	116
Rationale for the Cell Size	117
Network Transparency Operations	119
Errors and Error Rates	119
Error Correction and Detection	121
Probability of Discarding Cells	122
Overhead of the Cell Approach	124
Transmission Delay	126
ATM Labels	129
Multiplexing VCIs and VPIs	130
Cell Relay Bearer Service (CRBS)	130
Point-to-Multipoint and Multipoint-to-Multipoint Services	133
Summary	133

CHAPTER 7	The ATM Adaptation Layer (AAL)	134
	Principal Tasks of the AAL	134
	The AAL Sublayers	136
	Creating and Processing the AAL PDU	136
	Classes of Traffic	138
	Rationale for AAL Types	140
	Dividing CS into further sublayers	140
	AAL Naming Conventions	141
	AAL Type 1 (AAL1)	141
	The AAL1 PDU	142
	AAL 1 Modes of Operation	144
	Synchronization and Clock Recovery	145
	Running AAL 1 Traffic on a T1 Link	145
	AAL Type 2 (AAL2)	146
	The AAL 2 PDU	148
	Functional Model for AAL 2	150

Voice Packetization	152
Grouping Samples into Blocks	153
The Voice Packet	156
Packet Buildout at the Receiver	159
AAL Types 3, 4, 3/4, and 5 for Data	161
Pre-ATM Approach to Traffic Integrity Management	161
ATM Approach to Traffic Integrity Management	164
The Original AAL Types 3 and Type 4 (AAL3, AAL4)	166
AAL3/4	166
Naming conventions for AAL3/4	167
The AAL3/4 PDU	168
AAL3/4 Headers and Trailers	168
AAL3/4 Sequencing and Identification Operations	170
A Complete SAR-PDU and CPCS-PDU Example	171
Functional Model for AAL3/4	172
AAL Type 5 (AAL5)	175
Structure of AAL 5	175
The AAL5 PDU	176
Another Type—Available Bit Rate (ABR)	176
The AAL/ATM Primitives	177
Summary	180

CHAPTER 8	ATM Switching Operations	181
	ATM Switching	181
	Routing with the Cell Header	182
	Space and Time Switching	184
	Digital Cross Connects	185
	The Switching Fabric	187
	Point-to-Multipoint Operation	187
	Multiplexing, Label Switching, and Label Mapping	188
	Protection Switching	192
	Label Switching versus IP Address Routing	195
	Switching Technologies	196
	Shared Memory Switch	196
	Shared Bus Switch	197
	Crossbar Switch	197
	Multistage Switching	198
	Banyan and Delta Switching Networks	199
	VSLI and ASIC-based Switches	204
	Summary	206

CHAPTER 9	Traffic Management	207
	Traffic Management in an ATM Network	207
	Dealing with High Bandwidth Networks	208
	Low Delay Requirements for Processing Cells	209
	Challenge of Managing a Multiservice Network	211
	Control Mechanisms	212
	The Natural Bit Rate	212
	Potential Congestion Problems	214
	Traffic Control and Congestion Control	214
	Functions to Achieve Traffic Control and Congestion Control	215
	Allocation of Bandwidth	215
	Computing the Parameters for Queue Servicing	219
	Example of Queue Management Operations	219
	Dealing with Variable Delay	220
	Connection Admission Control (CAC) Procedures	222
	Usage Parameter Control (UPC)	223
	Traffic Management at the UNI—Basic Concepts	223
	Eckberg Scheme	224
	Multiplexing Traffic into the Cells	226
	Traffic Shaping Example	226
	Token Pools and Leaky Buckets	226
	Allocating Resources	230
	ATM Bearer Service Attributes at the UNI	231
	Traffic Control and Congestion Control	231
	Cell Arrival Rate and Cell Interval	232
	ATM Cell Transfer Performance Parameters	234
	ATM Layer Provisions for Quality of Service (QOS)	236
	ATM Forum and ITU-T Traffic Control and Congestion Control	238
	Generic Cell Rate Algorithm (GCRA)	238
	The Peak Cell Rate Reference Model	241
	Cell Delay Variation (CDV) Tolerance	242
	The ATM Services for LAN and Internet Traffic	244
	Examples of Feedback Operations	244
	Types of Feedback	247
	The ATM Service Architecture	249
	Examples of ABR Operations	251
	ABR Service Parameters	254
	The ABR Resource Management (RM) Cell	256
	Other Thoughts on the ABR Service Category	258

BuildOut Delay Procedures at the Receiving Endpoint	258
Work on the Guaranteed Frame Rate (GFR)	259
Summary	259

CHAPTER 10	Call and Connection Control	260
-------------------	------------------------------------	------------

ATM Connections on Demand	260
The ATM Address	263
The Address Registration MIB	265
Address Registration	265
The Connection Control Messages	266
Connection Setups and Clears	267
Q.2931 Timers and States	269
Connection Control Examples	270
Connection Setup	271
Connection Release	273
Restart Procedure	274
Status Inquiry	275
Add Party	275
Drop Party	278
Signaling AAL Reset and Failure	278
Functions of Q.2931 Messages and Information Elements	279
Messages for Call Control	279
Messages for Restart Operations	279
Messages for Adding and Dropping Parties	279
Descriptions of the Information Elements	282
Examples of Q.2931 Messages	287
Coding Conventions	287
AAL Parameters	288
User Traffic Descriptors	292
The PNNI Model	292
UNI Signaling, Version 4.0, and Voice and Telephony over ATM to the Desktop	295
Notification Procedures	295
Notification of Interworking Procedures (Progress)	295
AAL Negotiation	295
Call/Connection Alerting	296
Narrowband Bearer Capability Information Element	297

Narrowband High Layer Compatibility Information Element	297
Narrowband Low Layer Compatibility Information Element	297
UNI Signaling 4.0 Supplementary Services	297
Summary	299

CHAPTER 11	Internetworking with ATM Networks	300
-------------------	--	------------

The ATM Network as the Backbone for Other Networks	300
Using Q.2931 to Support Protocol Capability (Tunneling)	301
Broadband Low-Layer Information Element	304
THE NETWORK-TO-NETWORK INTERFACE	307
The ATM B-ISDN InterCarrier Interface (B-ICI)	308
Physical Layer Requirements at the B-ICI	311
Traffic Management at the B-ICI	311
Reference Traffic Loads	311
B-ICI Layer Management Operations	312
Specific Internetworking Services	312
PVC Cell Relay Service (CRS)	313
PVC Circuit Emulation Service (CES)	313
PVC Frame Relay Service (FRS)	316
SMDS Service	318
ATM LAN Emulation	321
RFC 1483 AND RFC 1577	322
The ATM Data Exchange Interface (DXI)	324
DXI Modes	325
DXI Support for Frame Relay	326
The Frame UNI (FUNI)	327
Are DXI and FUNI Needed?	328
Multiprotocol Over ATM (MPOA)	328
Companion Protocols to MPOA	329
Summary	330

CHAPTER 12	Synchronous Optical Network (SONET)	331
-------------------	--	------------

Purpose of SONET	331
Present Transport Systems and SONET	333
Foundations for SONET	333
Synchronous Networks	335
Optical Fiber—The Bedrock for SONET	336
Pertinent Standards	337

Typical SONET Topology	339
SONET Configuration	342
SONET Layers	343
Automatic Protection Switching (APS)	344
Payloads and Envelopes	346
Envelopes	346
Mapping ATM Cells into the SONET Envelope	350
Payload Pointers	350
Mapping and Multiplexing Operations	351
The Control Headers and Fields	355
SONET Equipment	356
Progress in SONET Penetration	362
Summary	362

CHAPTER 13	Operations, Administration, and Maintenance (OAM)	363
	THE NETWORK MANAGEMENT MODEL	363
	Operation and Maintenance (OAM) Operations	364
	ATM Functions at the U- and M-Planes	367
	U-Plane Operations	367
	M-Plane Operations	369
	End-to-End and Segment Flows	370
	The SONET OAM Functions	372
	Maintenance and Alarm Surveillance	372
	Failure States	373
	Alarm Indication Signals (AISs), FERF, and Yellow Signals	373
	Examples of Remedial Actions upon Entering a Failure State	375
	The OAM Headers	375
	Section Overhead	376
	Line Overhead	378
	STS Path Overhead (STS POH)	378
	ATM Use of the OAM Octets	379
	Using Payload Pointers for Troubleshooting Timing Problems	380
	OAM at the ATM Layer	381
	Fault Management	382
	Performance Management	385
	Activation/Deactivation	387
	The ATM Management Information Bases (MIBs)	388
	The Integrated Local Management Interface (ILMI)	389
	The ILMI MIB Tree Structure	389
	ILMI MIBs	393

MIB (RFC 1695)	397
The ATM MIB Groups	398
The ILMI MIB and the ATM MIB	403
The Layer Management/ATM Primitives	403
Other ATM Forum Specifications for Network Management	405
Summary	405

CHAPTER 14	Physical Layer Services for ATM	406
-------------------	--	------------

PHYSICAL LAYER OPTIONS FOR ATM	406
The ATM/Physical Layer Primitives	407
ATM Mapping into SONET STS-3c	408
ATM Mapping into DS3	410
Other Aspects of the DS3 Scheme	412
Circuit Emulation Service Interoperability (CES-IS) Specification	412
ATM Mapping into the 100 Mbit/s Multimode Fiber Interface	414
Functions of the U-Plane Physical Layer	414
ATM Mapping into the 155.52 Mbit/s Private UNI	416
Multimode Fiber Interface	416
Shielded Twisted Pair Interface	416
Private UNI for 51.84 Mbit/s and Subrates	418
Mapping DS1, DS3, and CEPT Payloads into SONET Frames	419
The VT/VC Structure	419
Floating and Locked VT Mode	422
Interworking ATM and SONET	423
The 2.4 Gbit/s Physical Layer Specification	425
Inverse Multiplexing for ATM (IMA)	425
Rules for IMA Operations	426
The IMA Sublayer in the Layered Model	429
Summary	430

CHAPTER 15	The ATM Market	431
-------------------	-----------------------	------------

A Soft Market, Initially	431
Recent Successes and Projections	431
ATM and IP	432
ATM in Residential Broadband	433
ATM's Future	433
In Conclusion	434



Preface

Some time ago, one of my clients, who is a systems engineer for a large telecommunications firm, told me that one of her major problems is staying abreast of the technologies that are embedded into her company's products. I hear this statement often. Like many others in our industry, this person does not have the time to read the technical specifications published by the standards organizations and the user forums. She must spend her professional time performing the day-to-day tasks of the job, essentially taking care of her clients and her accounts.

The common lament is that many professionals are barely ahead of their customers in their knowledge and ability to field their questions about not just the company products, but how they fit into an overall telecommunications architecture. Increasingly, this systems engineer has been forced to know about many diverse protocols, standards, and architectures. Her clients query her on topics such as the relationship of ATM and Frame Relay, why Novell uses IPX and not IP, why NetBIOS is not routable, etc.

It is to this person that this series is devoted (indeed, most of my books are so focused). It is my hope that I have provided a series that will meet this engineer's needs in the field.

This book is the second book of Prentice Hall's *Advanced Communications Technologies*, which serves as a complement to the flagship book, *Emerging Communications Technologies*.

I have included a chapter on existing technologies, titled “Emerging Technologies”. This chapter is a summary of a chapter of the same title from the flagship book for this series. I have added a section in this chapter on why functions and services of several of these technologies (for data networks) have been reduced or eliminated in an ATM network. I suggest the reader review this chapter for two reasons: (a) to make certain the ISDN, X.25, SS7, and T1/E1 systems are understood, and (b) to understand why ATM operations do not include many functions that are an integral part of current data networks.

The ATM story is far from complete. As of this writing, ATM systems are now being deployed, but some of the ATM standards are still being written. One cannot wait to write a book on emerging technologies until they have “emerged,” else there would be no book to write. So, this book represents the state of ATM as of the date of submittal of my work to my publisher.

I am happy to say that the first edition of this book has been received very well in the industry. This second edition reflects the changes that have taken place since the publication of the first edition. All chapters have been updated, with emphasis on these new or revised topics:

- Voice over ATM with AAL2
- The Available Bit Rate (ABR) and Unspecified Bit Rate (UBR) procedures
- Updated MIBs and network management
- Internetworking and tunneling with IP
- Inverse multiplexing
- Protection switching



1

Introduction

This chapter discusses the current telecommunications technologies in place today. It explains why the telecommunications industry is implementing new communications technologies to overcome the deficiencies of current systems. An analysis is made of the communications requirements of upcoming user applications.

As a prelude to subsequent chapters, a general overview is provided of the asynchronous transfer mode (ATM) and the synchronous optical network (SONET)/synchronous digital hierarchy (SDH) that are being developed to meet the needs of these applications.

THE PRESENT TELECOMMUNICATIONS INFRASTRUCTURE

Scores of books have been written about the impact of the computer on our society, and the flagship book for this series (*Emerging Communications Technologies*) examines this subject. Studies are cited as well about the role of telecommunications (and the recent growth of cordless telephony, optical fiber, cable television, and video-on-demand technologies) that are having a profound effect on our professional and personal lives. This book shall not reiterate the thoughts cited in these books, but shall concentrate on the technical underpinnings that must be in place for these technologies to flourish.

Present Technologies for Voice, Video, and Data Networks

The present communications infrastructure supporting voice, video, and data networks is founded on technology that is over 25 years old. In spite of their age, these networks have served the industry well, even in recent times, for they have provided a cohesive foundation on which to build the modern telecommunications infrastructure. Yet when we consider that the modern commercial computer is only about 35 years of age and the personal computer came into being in the early 1980s, a 25-year-old technology seems like a technical dinosaur.

In 1962 the U.S. Bell System (as it was known in the pre-divestiture days) installed the first commercial digital voice system in Chicago, Illinois. The system was known as T1, and carried 24 voice channels over copper wire between Bell's telephone offices. The European carriers followed shortly with a similar system, called E1. As technology improved, T1 was deployed in higher capacity systems. Shortly thereafter, T3 became a common carrier system for users who needed greater capacity than the T1 offerings. The T3 system can transport 28 T1 signals, which means one T3 link can support 672 voice calls. And the European carriers followed with the E3 technology.

T1 and T3 have become the foundations for the majority of voice networks systems provided by the North American telephone companies.¹ While these systems were designed originally for voice systems, they now can be configured to support data and video applications as well.

In the early 1970s, another technology was deployed to support data networks. This technology is called *packet switching*. Unlike the T1/E1 and T3/E3 networks, packet switching was designed for data applications, and packet switching networks have become the foundation for the majority of data networks.

At about the same time that packet switching networks were being deployed, the International Telecommunications Union-Telecommunication Standardization Sector (ITU-T, formerly the CCITT) published the X.25 specification. As the reader may know, X.25 defines the procedures for user computers to communicate with network machines (packet switches) and to transport data to another user computer. X.25 has become a widely used industry standard and has facilitated the building of standardized communications interfaces among different vendors' machines.

¹T1 and T3 are often used synonymously with DS1, and DS3. T1 and T3 describe the physical aspects, such as voltages and media specifications, and DS1 and DS3 describe the framing and formatting conventions. This book will use them interchangeably, in deference to common industry practice.

These communications technologies were designed to support fairly modest requirements for voice and data transmissions, at least when compared to modern applications needs. For example, the T1 systems support a transfer rate of 1.544 Mbit/s, and the T3 system operates at approximately 45 Mbit/s. These bit transfer rates may seem high to the reader, but remember that a 45 Mbit/s transport system like T3 only supports 672 voice calls—a lot of T3s have to be in operation to support the public telephone network.

Likewise, X.25 was designed for data systems that operate at only a few bit/s or a few hundred bit/s—typically 600 to 9600 bit/s. Although X.25 can be placed on very high-speed media and can operate quite efficiently at high speeds, a substantial amount of subscriber equipment and software has been designed for modest transfer rates—typically no greater than 19.2 kbit/s.

Once again, sending data at a rate of 19.2 kbit/s may seem fast. After all, this translates to a transfer rate of 2400 characters per second (19,200/8 bits per character), and no one can type in an email message that fast. However, for other applications, this speed is not sufficient. File transfers, database updates, and color graphics (to mention a few) need much greater transfer rates.

Typical Voice Networks. Table 1–1 provides some examples of typical voice carrier systems. The E1 system (also called CEPT1) is Europe's principal technology for carrier transport systems. Japan's basic technology is also based on T1, but Japan's higher capacity systems are not in alignment with either European or North American systems.

Nonetheless, because telephone networks have historically been highly regulated within a country and controlled by one enterprise (in the United States, AT&T before divestiture, and in other countries, the Postal Telephone and Telegraph Ministries [PTTs]), a national telephone network architecture uses common standards, conventions, and protocols. Thus, interworking different vendors' equipment is relatively simple

Table 1–1 Typical Voice Carrier Systems

Type	Digital Bit Rate	Voice Circuits	Age	Standard or Proprietary
DS1	1.544 Mbit/s	24	32 years	Standard*
E1	2.048 Mbit/s	30	32 years	Standard*
DS3	44.736 Mbit/s	672	31 years	Standard*

*Within national boundaries

(telephones, fax devices, and answering machines are common examples). This situation is not true with the architecture for data networks, as we shall see in the next section.

Typical Data Networks. In contrast to the voice world, data networks and protocols have evolved into an almost bewildering array of disparate and incompatible systems. Table 1–2 provides examples of the more widely known and used standards and vendor products.

Interconnecting some of these systems is almost impossible. When it is possible, the resulting systems are very complex and very expensive, due to the need to provide protocol converters between the systems. Yet these systems form the foundation for our current data networks.

As Table 1–2 shows, some systems are standardized while others are proprietary. Also, they may operate as local area networks (LANs), wide area networks (WANs), or both. Most of them were conceived over ten years ago, although all have been enhanced since their inception.

Why are so many incompatible systems in existence to do one thing: transport data between computers? The answer is simple. The data communications and computer industry, unlike the telephone industry, has had very little regulation imposed upon it. Additionally, this industry is quite young, and many systems and products were developed before standards were written by organizations such as the ITU-T, the ISO, and the Internet task forces.

Table 1–2 Typical Data Networks and Protocols

Vendor or Standard	Sponsor	Age ¹	Standard or Proprietary	WAN or LAN ²
X.25	ITU-T & ISO	25 years	Standard*	WAN
OSI	ITU-T & ISO	15 years	Standard*	Both
TCP/IP	Internet	13 years	Standard**	Both
SNA	IBM	22 years	Proprietary	Both
DECnet	Digital	22 years	Proprietary	Both
AppleTalk	Apple	12 years	Proprietary	LAN
Ethernet	Xerox, Digital, Intel	17 years	Standard***	LAN
Netware	Novell	12 years	Proprietary	LAN

¹Approximate ages; all have evolved and have been enhanced

²WAN is wide area network and LAN is local area network

*Recognized by international standards groups

**A de facto standard by virtue of its wide use

***Revised slightly to become the IEEE 802.3 standard

The current status of the evolution of the data communications industry is both good news and bad news. The good news is that the lack of a dominant player (such as “Ma Bell” in the telephone industry) has led to much competition and the availability of some extraordinary systems and products at reasonable prices. The bad news is that some customers are saddled with single-vendor systems, because a vendor-specific system cannot operate easily with any other vendor’s system.²

The industry is realizing that competition can continue but under an umbrella of standards. Frankly, many of the systems and products listed in Table 1–2 do just about the same things, but they do them differently.

PRESENT AND FUTURE REQUIREMENTS

In the past few years, the processing power of an ordinary personal computer (PC) has increased so rapidly that the terms high-speed workstations and mainframe computers are losing their meanings. Small machines are becoming as powerful the once “large machines.”

The reader need only take a brief glance at the daily newspaper to grasp the rapid increases in the power and functionality of the computer. The increased capacity of these machines means they can send and receive large amounts of information in a very short time. The end result is the need to upgrade communications systems to support the PC’s requirements. Future PCs will necessitate the development of even higher-speed and more powerful communications systems.

The communications infrastructure to support the connection of these computers and new applications must also be upgraded. While individual homes and workstations can continue to use conventional, existing media (the coaxial TV cable and twisted pair on the local loop), the service providers’ facilities, media, and networks need more bandwidth.

Downsizing and Outsourcing: Reliance on Telecommunications

In the last few years, several industrialized countries have witnessed a trend known as downsizing. Downsizing entails businesses (mostly large businesses) shedding employees, capital resources, and, in many instances, buildings. With this downsizing comes a new trend—outsourcing. Many of these firms are hiring outside contractors to pro-

²The situation is changing with the acceptance of the Internet and OSI-based protocols on a worldwide basis.

vide services—such as training, food operations, mailroom operations, and software programming—that were once provided by employees. While these companies must continue their ongoing operations, they are doing it increasingly with distributed computers and communications facilities linking their computers together.

It has long been a cliché that telecommunications is playing one of the most critical roles in our information society and, indeed, in our personal culture. This role will become even more important as more humans learn to interact with the computer and exploit its productivity potential. The trend, of course, will lead toward (this writer hopes) more open societies as the telecommunications infrastructure embeds itself into most people's lives. So, without going into a monologue on the benefits of the telecommunications infrastructure, the next section describes some of the problems faced in current communications architecture.

Present Systems: Too Much or Too Little

Voice. For today's voice transport systems, the present structure provides adequate capacity for many applications, but as stated earlier and explained in the flagship book for this series, the capacity is insufficient for others. In addition, these transport systems suffer from the asynchronous nature of their design. In this context, asynchronous means that the components of the network are not synchronized with a common clock. Consequently, it is not unusual for errors to occur between transmitting and receiving machines because the machines are using different timing schemes. An analogy would be a person talking too fast to a listener, who misses part of the speech.

Perhaps more serious, it is recognized increasingly that these systems have very limited operations, administration, and maintenance (OAM) capabilities, known by many people as network management capabilities. The supposition (30 years ago) of these simple designs made good sense, because the operations needed to support substantial OAM required more overhead than the limited-capacity network could bear. However, with the increased use of high-speed optical fiber and fast processors, building powerful OAM modules within new systems becomes feasible.

Wide Area Data Networks. For data communications networks, ironically, it is accepted that the current systems (especially WANs) may be doing too much, in that they are performing a number of redundant functions of marginal benefit. We shall have more to say about this idea

in subsequent chapters, but for the present, it can be stated that redundant functions are performed on the majority of user traffic. In X.25, for example, sequencing and flow control, as well as positive acknowledgments (ACKs) and negative acknowledgments (NAKs), are performed at least twice.

With the advent of relatively error-free, high-capacity networks, and with the concomitant implementation of very powerful end-user workstations, the new networks take the view that many of these operations are no longer needed in the network. Indeed, many of the functions are simply removed from the network and placed in the customer premises equipment (CPE), such as user workstations and personal computers.

We return to these important points several times in this book. The reader may refer to Chapter 6 (Errors and Error Rates) and Chapter 7 (Pre-ATM Approach to Traffic Integrity Management, and ATM Approach to Traffic Integrity Management) for immediate follow-up on these ideas.

Local Area Data Networks. It is recognized that the processing power of personal computers and workstations is doubling about every two years. Ten MHz processors were considered state of the art in 1990; 25 MHz processors were in use in 1992, and 350 MHz processors in 1998. The trend will continue and with it an associated need for more bandwidth to support the communications between these machines. A guideline, long accepted in the industry, is that a well-tuned computer system has one bit of I/O (input/output) for every instruction cycle. Workstations have processors that operate with a cycle time of about 1 nanosecond and a performance of 1 billion instructions per second. These workstations create enormous bandwidth demands on communications networks.

LANs have not kept pace with the progress of the CPUs, and with the exception of the fiber distributed data interface (FDDI), the technology has remained in the 4 to 16 Mbit/s range. As a consequence, the last few years have seen a decrease in the number of computers attached to a LAN segment [HERM93]. This trend cannot continue; it is too expensive and complex. Therefore, a new family of LANs is evolving to meet the increased needs of the user stations. The metropolitan area network (MAN) and fast Ethernet are among these solutions. So is ATM, the subject of this book.

Bottleneck at the Local Loop. One of the biggest problems facing the industry is the limited bandwidth of the local loop (the link between

the customer premisis equipment (CPE) and the voice, video, or data service providers. While this subject is beyond the scope of this book (but is covered in *Residential Broadband Networks: xDSL, HFC, & Fixed Wireless Access*, a companion book in this series), some new local loop technologies are deploying ATM to help solve the problem.

Costs of Leased Lines

In the past few years, LANs and WANs have been interconnected with bridges, routers, gateways, and packet-switched networks. These internetworking units connect to the LANs and WANs through dedicated communications channels (leased lines). As a general practice, leased lines have been “nailed up” end-to-end through the network to the user’s CPE. The user is provided with the leased line on a dedicated basis and the full transmission capacity is available 24 hours a day (with some exceptions). Therefore, the user pays for the circuit regardless of its utilization. Moreover, if a connection is needed to yet another location (say another city in the country), another leased line must be rented from the public telecommunications operators (PTOs), such as AT&T, Sprint, and MCI—once again on a end-to-end, continuous basis.

Even though leased lines are becoming less expensive, the use of these lines to connect internetworking units with LANs and WANs is still a very expensive process. Moreover, reliability problems occur because individual point-to-point leased lines have limited backup capability.

A better approach is to develop a LAN/WAN-carrier network that provides efficient switching technologies for backup purposes as well as high-speed circuits—a network that will allow users to share the expensive leased lines. This concept is called a *virtual network* or a *virtual private network* (VPN).

VIRTUAL COMPANIES AND VIRTUAL NETWORKS

During the past decade, smaller companies increasingly have been competing successfully against their larger rivals. Indeed, some large companies are so unwieldy and so fraught with bureaucracy and overhead that they cannot compete in many arenas with their smaller counterparts. Of course, the value of economy of scale still pertains; for example, my small company is not going to spend the enormous funds needed to develop an ATM switch.

However, the term virtual company is a useful description of a growing industry of small groups or individuals working out of rented busi-

ness centers or even at home (SOHO: small office/home office). This type of operation gives the illusion of the traditional private enterprise, but it usually has few or no employees, and may not even have a receptionist! Well, the answering service/ answering machine is a virtual receptionist.

A treatise on the reasons for this phenomenon is beyond the scope of this book. Nonetheless, it must be emphasized that the downsizing of companies and the distribution of workloads to remote offices (and homes) cannot occur without the accompanying supporting communications infrastructure. Indeed, the very premise of deconstruction, downsizing, and outsourcing is based on the idea that the smaller companies and their consultants (who act as virtual employees) will have access to each other through high-speed, reliable communications facilities. Increasingly, the facilities are being implemented with a virtual private network (VPN) (an old concept that has been renamed and taken off the shelf for use in today's networks).

The VPN is so named because an individual user or enterprise shares communications channels and facilities with other users. Switches are placed on these channels to allow an end user to have access to multiple end sites. Ideally, users do not perceive that they are sharing a network with each other, thus the term virtual private network—you think you have it, but you don't.

Figures 1–1 and 1–2 illustrate the concepts of VPNs and their advantage over dedicated systems. In Figure 1–1(a), four customer sites from company A are connected to each other through dedicated channels (leased lines). While effective, this approach is very expensive, and it is unlikely that these lines are used on a continuous basis. Company B in Figure 1–1(b) has a separate arrangement connecting its four offices with dedicated lines in the same cities as company A.

In contrast, through the use of a VPN (Figure 1–2), the two companies can share the communications facilities. The VPN provider provides a network for multiple users. This approach allows the traffic to be routed to various endpoints and does not require the end-user devices to “nail up” private leased lines. In some implementations, companies migrating to a VPN have reduced their costs by 30 percent vis-à-vis leased lines.

The difference between fully meshed leased lines and VPNs is even more dramatic when another location is added to a private network. If the user (say company A) wishes full connectivity to all sites, this approach requires the leasing of long-distance lines to all cities. Of course, with a fully meshed VPN, the same number of private lines are required between the switches, but the switches are relaying traffic from multiple users. So, if company A adds an office (say) in Dallas, then it would only require the leasing of one dedicated local loop to the most convenient





Figure 1-1 Leased lines.

VPN switch. The number of dedicated lines would remain the same as long as additional switches are not added to the network.

VPN is a relatively new term in the computer/communications industry. Yet, this new term describes an old concept; the ideas behind the VPN are not new at all. Public X.25 networks have offered VPN services for years, and switched T1 services also offer VPN-like features. However, we shall see that ATM offers more powerful VPNs than these older technologies.

The first part of this chapter has described some of the problems and challenges that exist in our present climate. The remainder of the chapter describes some solutions, notably ATM and SONET.³

³This book describes SONET. While SDH is quite similar to SONET, the two technologies use different terms and definitions.




Figure 1-2 Virtual private networks.

FAST RELAY NETWORKS AND ATM

Much of the emerging technology to support modern applications is based on the idea of relaying traffic as quickly as possible. This idea is often called fast packet relay or fast packet switching. These names are considered generic terms in this book and are used in a variety of ways in the industry. Therefore, we will use the term fast relay systems. Currently, fast relay comes in two forms: Frame Relay and cell relay.⁴ Figure 1-3 shows the relationships of these two forms of fast relay systems.

Frame Relay uses variable-sized protocol data units (PDUs), which are called *frames*. The technology is based on the link access procedure for the D channel (LAPD) that has long been used in integrated services digital network (ISDN) systems. Most frame-based implementations are using LAPD as the basic frame format for the relaying of the traffic across permanent virtual circuits (PVCs). A modified version of ISDN's Q.931 has been introduced into the industry for control signaling and for setting up connections between the user and the network. Q.931 can

⁴I could be called to task for this depiction because I have excluded Fast Ethernet, Gigabit Ethernet, and the emerging residential broadband (RBB) technologies. This depiction is a focus on wide area networks.




Figure 1-3 Types of relay systems.

form the basis for public switched offerings (switched virtual calls [SVCs]). The reader can refer to Chapter 4 for a tutorial on LAPD and Q.931.

In contrast to Frame Relay, cell relay uses a fixed length PDU, which is called a *cell*. The cell usually consists of a 48-octet payload with a 5-octet header, although some implementations use different cell sizes. This cell (with slight variations) is being used on both the asynchronous transfer mode (ATM) and the IEEE 802.6 standard, more commonly known as the Metropolitan Area Network (MAN) specification. In turn, the 802.6 standard is being used as a basis for the Switched Multi-megabit Data Service (SMDS).

The industry is still divided on the advantages and disadvantages of the use of cell relay versus Frame Relay. While there are arguments for both, the trend is toward the use of the cell relay technology. The reasons are many, but before they are explained, a brief summary of cell relay technology is in order.

At the CPE, such as a router or a PBX, a customer's traffic, which could be variable in length, is segmented into a smaller fixed length units called cells. As stated earlier, in most cell relay systems, the cells are only 53 octets in length with 5 octets devoted to a header and the remain-

ing 48 octets consisting of user information (the payload). The term cell is used to distinguish this PDU from variable length PDUs (frames and packets).

The term “asynchronous” in asynchronous transfer mode means the cells can be used without precise timing requirements and not necessarily on a synchronous, periodic basis. The cells may indeed be transported in a synchronous network (such as SONET), but they can be filled based on the needs of the application, which may be in a synchronous or asynchronous fashion. The ATM machines create a continuous stream of cells, and if no traffic is sent, the cells are empty and are called idle cells.

Cell relay is an integrated approach to networking in that it supports the transmission and reception of voice, video, data, and other applications. This capability is of particular interest to large companies that have developed multiple networks to handle different transmission schemes. As examples, common carriers, telephone companies and Postal Telephone and Telegraph Ministries (PTTs) must support many types of applications and historically have implemented a variety of networks to support them.

Why do many people prefer the cell technique over the frame technique? First, the use of fixed length cells provides for more predictable performance in the network than with variable length frames. Transmission delay is more predictable, as is queuing delay inside the switches. In addition, fixed-length buffers (with cell relay technology) are easier to manage than variable-length buffers. In essence, a fixed-length cell relay system is more deterministic than the use of a technology with variable length data units. Cell relay is also easier to implement in hardware than with variable length technology.

Some people have expressed concern about the high overhead of cell relay—because of the ratio of 5 octets of header to every 48 octets of user payload. But many people in the industry believe that the constant concern with efficient utilization of raw bandwidth is not a sound approach for the future. With the capacity of optical fiber and high-speed processors that are entering the marketplace, the approach of cell relay is to concentrate on superior quality-of-service features to the user. The philosophy is straightforward: Let the fast optical channels and the fast computers handle the transmission and processing of the overhead traffic.

Figure 1–3 shows several other aspects of the frame and cell relay technologies. Initial implementations of Frame Relay and cell relay have focused on PVCs (logical connections to the network that are available at any time). Recent enhancements have added SVCs (logical connections that are made available on demand, and not available at any time—similar to a dial-up telephone call).

Table 1–3 Problems that Existing Technologies Cannot Solve [ATM93b]

Problem	1st Place	2nd Place	3rd Place
1. LAN performance above 100 Mbit/s	5	4	2
2. Scalable WAN bandwidth	5	4	1
3. Integration of voice, video, and data	4	3	1
4. Network management and logistics	2	2	4
5. Uniform architecture in LANs, MANs, WANs	2	2	1
6. Bandwidth on demand (pay for use)	1	3	4
7. Network complexity	1	1	0
8. Support for multicast operations	0	1	1
9. Integration of multiple data applications	0	0	3
10. Support for synchronous applications	0	0	2

APPLICATIONS USE OF ATM

In 1993, the ATM Forum conducted a study of 200 companies to determine their plans for the future in relation to the use of ATM and other technologies [ATM93b]. The study entailed the companies' filling in questionnaires, and then consultant John McQuillan followed up with interviews. The study made no claim about its statistical validity, but purported to show accurate indications for the "near" future. The initial usage for ATM focused on data applications, but it was believed that voice and video would become more important in the future. Most respondents viewed that ATM would provide support for voice, video, and data by 1996.⁵

The survey asked the respondents to identify current problems that existing technologies cannot solve. The responses shown in Table 1–3 are similar to those concerns that this writer heard in 1993 (and continues to

⁵A few products in 1996 incorporated an integrated ATM switch to support voice, video, and data—but not many. The marketplace did not corroborate this part of the summary.

hear in this second edition) from clients as well. Certainly, the top seven answers in this survey are what I find, so I have left this survey in this new edition because it is still pertinent.

FAST RELAY NETWORKS AND SONET

At the beginning of this chapter, it was stated that for the first time in the history of the computer/communications industry, networks throughout the world are embracing a worldwide set of computer/communications standards. One of these standards deals with multiplexing and signaling hierarchies for digital carrier transport networks. Because of its importance, we shall introduce the subject here, and examine it more closely in later chapters.


During the past 35 years, three different digital multiplexing and signaling hierarchies have evolved throughout the world. These hierarchies were developed in the regions of Europe, Japan, and North America. Fortunately, all are based on the same pulse code modulation (PCM) signaling rate of 8000 samples a second, yielding 125 microsecond sampling slots ($1 \text{ second}/8000 \text{ samples} = .000125$). Therefore, the basic architectures interwork reasonably well.

But the regions vary in how the systems are implemented, which results in extensive and expensive conversion operations if traffic is exchanged between them (Figure 1–4). Moreover, the analog-to-digital conversion schemes also differ between North America and Europe, which further complicates interworking the disparate systems.

Japan and North America base their multiplexing hierarchies on the DS1 rate of 1.544 Mbit/s. Europe uses the E1 2.048 Mbit/s multiplexing scheme. Thereafter, the three approaches multiplex these payloads into larger multiplexed packages at higher bit rates and use different values for the multiplexing integer n .

As depicted in Figure 1–4, the synchronous digital hierarchy (SDH) and the synchronous optical network (SONET, North American term) support the schemes that have been in existence for many years, but specify a different multiplexing hierarchy.

The basic SDH/SONET rate is 155.52 Mbit/s. SDH/SONET then uses an x 155.52 multiplexing scheme, with x as the multiplexing factor (e.g., $155.52 * x$, where x is 3, 6, etc.). Rates smaller than 155.52 Mbit/s are available at 51.840 Mbit/s for SONET but not for SDH. We shall have more to say about Figure 1–4 in later chapters, and the other terms in this figure will be explained at that time.



Note: Unless noted otherwise, speeds in Mbit/s

Figure 1-4 SONET and SDH hierarchies.

At long last, worldwide agreement has been reached (with minor exceptions) on a common digital multiplexing transport scheme. This agreement can only foster new technologies and decrease the costs of implementing them.

BROADBAND ISDN

The ITU-T describes the Broadband Integrated Services Digital Network (B-ISDN) as a network built on the concepts of the ISDN model, and a network that is implemented with the ATM and SONET technologies. These two technologies are complementary to each other, as shown in Figure 1-5. In its simplest form, the SONET technology acts as the physical carrier transport system for the user payload. The user payload




Figure 1-5 ATM and SONET.

is carried in ATM cells. In other words, the SONET network acts as a service provider to the ATM traffic.

Figure 1-5 also shows that the ATM components can act as the user-to-network interface (UNI) with the user's CPE. This approach allows the user to negotiate a wide variety of services at the ATM node. In addition, a SONET system can be terminated at ATM switches. These switches relay the traffic onto outgoing SONET ports or onto local UNIs.

ATM and SONET form the "technical alliance" for B-ISDN. SONET provides extensive operations, administration, and maintenance (OAM) functions and the basic carrier transport services (including backup facilities), and ATM provides services and interfaces to the user machine, as well as switching operations between SONET communications links.

As explained in subsequent chapters, one component of ATM also allows the ATM node to process and support multiapplication systems such as voice, video, data, music, and fax. These application-specific operations remain transparent to SONET.

Table 1-4 summarizes several aspects of ATM and SONET. As explained earlier, both are international standards sponsored by the ITU-T

Table 1–4 Broadband ISDN (B-ISDN)

Vendor or Standard	Sponsor	Age	Standard or Proprietary	WAN or LAN
ATM	ITU-T	2–7 years	Standard	Both
SONET	ITU-T	5–9 years	Standard	Both ¹

¹Although most SONETs are deployed in WANs

and several other national standards bodies.⁶ Although the age columns in Table 1–4 suggest they are rather mature technologies, these numbers reflect the period in which research has been conducted and implementations have been occurring. Indeed, SONET and ATM are still evolving as of this writing, and research on the use of these technologies continues.


Also, SONET and ATM can operate on WANs or LANs, although most implementations of SONET have been in large public wide area carrier transport systems.

The reader might wish to compare Table 1–4 with Tables 1–1 and 1–2, which summarize the emerged technologies. The first thing to note is that ATM and SONET are considered to be single solutions to transport and switching systems in contrast to the multiplicity of standards and proprietary implementations that the industry has fostered in the past. This writer is not suggesting that ATM and SONET are replacements for all the technologies in Tables 1–1 and 1–2. Such is not the case at all. Rather, the design goal of the B-ISDN technologies is to provide for *backbone* transport systems for voice, data, and video applications. Notwithstanding this goal, ATM and SONET, if implemented properly, could result in the paring down or even elimination of some of the technologies listed in Tables 1–1 and 1–2. The rationale for this statement shall be supported in other chapters of this book.

PRINCIPAL SPECIFICATIONS FOR ATM

Figure 1–6 should prove helpful to you as you read the remainder of this book. We will examine each of the depicted operations/functions in this figure, because they provide the basis for the specifications and stan-

⁶SONET is not a formal ITU-T Standard. The ITU-T Standard is SDH, which was derived from SONET.



Where:

DXI	Data exchange interface
FUNI	Frame UNI
ICI	Intercarrier interface
ILMI	Integrated local management interface
ITU-T	International Telecommunication Union-Telecommunication Standardization Sector
LANE	LAN emulation
MIB	Management information base
MPOA	Multiprotocol over ATM
NNI	Network-to-network interface (also, network-node interface)
OAM	Operations, administration, and maintenance
PNNI	Private network-network interface
PVC	Permanent virtual circuit
RFC	Request for comments
SVC	Switched virtual circuit
UNI	User-network interface

Figure 1-6 Principal standards and specifications for ATM.

dards from which ATM products are developed and implemented. Specific parts of the book will also give you an update on the status of each of these areas.

As a general statement, the network-node interfaces (network-to-network) (NNI) were completed only last year by the ATM Forum and the ITU-T. The ITU-T operations, administration, and maintenance (OAM) standards are nearing completion. All others have been completed and are at various stages of implementation.

THE ANCHORAGE ACCORD

One of the ongoing complaints about some of the formal standards organizations is the time it takes for them to complete a technical specification. In addition, these specifications (standards) often are incomplete, and issues are left for “further study.”

This complaint certainly cannot be levied against the ATM Forum. It has published about 70 specifications. Indeed, there have been complaints that network managers and ATM vendors cannot form a cohesive ATM “outlook,” due to the profusion of information emanating from the ATM task forces. I prefer to have this overabundance of information, rather than a sparsity, none at all.

The ATM Forum reached an agreement (the Anchorage Accord) that is designed to address these problems. It contains a set of specifications that are stable and provides guidance on the applicability of the specifications to six specific network environments. These environments are shown Table 1–5 along with the applicable specifications to support each network environment.

In a nutshell, notable progress has been made by the ATM Forum in getting ATM started, and other issues are being addressed, such as voice over ATM.

It is the goal of the ATM Forum to complete the outstanding major specifications as soon as possible in order to facilitate the creation, implementation, and use of standardized ATM products.

SUMMARY

It is recognized that upcoming user applications are demanding greater throughput and lower delay from communications systems. ATM and SONET are designed to provide the capacity to support these applications.

Table 1-5 The Anchorage Accord

ATM Specification	Campus	Legacy LAN	Multimedia Desktop	Extended Campus	Legacy WAN	Multimedia WAN
ATM UNI v3.1	X	X	X	X	X	X
BICI v2.0					X	X
DXI v1.0		X				
ILMI v4.0	X	X	X	X	X	X
LANE over ATM 1.0						
Includes LANE						
Client Management	X	X				
Network Management:						
Customer Network						
Management for ATM						
Public Network Service			X			
Network Management:						
M4 Interface						
Requirements &						
Logical MIB						X
Network Management:						
CMIP for the M4						
Interface						X
Network Management:						
M4 Public Network						
View Requirements						
& Logical MIB						X
Network Management:						
M4 Logical MIB						
Addendum						X
Network Management:						
CES IW Requirements						
Logical CMIP MIB						X
Network Management:						
M4 Network View						
CMIP MIB 1.0						X
Network Management:						
AAL Management						
for the M4						X
PNNI v1.0				X	X	X
Frame UNI				X	X	
Circuit Emulation					X	X
Native ATM Services:						
Semantic Description			X			
Audio Visual Multimedia			X			
ATM Name System 1.0	X		X	X		X
UNI Signaling 4.0	X	X	X	X	X	X

From: *Business Communications Review* (BCR), June 1997, George Dobrowski and Marlis Humphrey.

Optical fiber systems are providing the foundation for the media support of these technologies and high-speed processors are providing the speed to process the traffic.

The movement toward worldwide standards for carrier transport technology, multiplexing techniques, and switching methodologies are at last providing the groundwork for a comprehensive and homogeneous approach to computer/communications networks.



8

ATM Switching Operations

This chapter examines ATM switching operations. It also explains how the VPI and VCI are used by a switch to make routing decisions. Since switching is not part of the ATM standards and vendors use a wide variety of techniques to build their switches, this chapter focuses on the prevalent methods in use today. In most implementations, this part of the ATM node is proprietary; vendors put a large amount of research and development into the switching fabric. In most instances, the author was able to obtain the general architectural descriptions from several vendors. Therefore, this chapter reflects examples of ATM switches as well as a description of current research on cell switching.

Several approaches to cell switching are examined in this chapter. The topic of ATM switching is worthy of an entire book, so the goal of the chapter is more modest and is designed to give the reader a general overview of cell switching architecture and the major components of an ATM switch. Be aware that there is no “best way” to build an ATM switch, and the ATM standards wisely do not address this aspect of the technology.

ATM SWITCHING

One of the most important aspects of ATM is the switch and how fast it relays cells (without cell loss). ATM must accept asynchronous and

synchronous traffic, as well as connection-oriented and connectionless traffic. Queuing delay and switching delay must be minimized if the ATM network is to perform satisfactorily in its support of user applications.

One goal of the ATM network is to adapt gracefully to changing network traffic profiles with the ability to adjust to increased or decreased input traffic (within limits, however). Since an ATM network must receive traffic that is somewhat unpredictable, the network must be able to adjust to changing network conditions. All these problems and their solutions are not defined in the ATM standards. Vendors are free to implement the techniques they deem appropriate for their product.

Routing with the Cell Header


As described in earlier chapters, the virtual channels are aggregated through multiplexers and switched (routed) through the network. Two types of switches are used in an ATM network for the multiplexing/demultiplexing and routing of traffic (see Figure 8–1a). Conceptually, a virtual path switch is only required to examine the VPI part of the header for multiplexing/demultiplexing and routing. Virtual channel switches, on the other hand, must examine the whole routing field; that is to say, the VCI and the VPI.

The intent of this approach is to speed up switching and routing at intermediate points. Instead of requiring the switch (in some instances) to examine the entire virtual channel field, the VCI value can be transported transparently through a virtual path connection.

In Figure 8–1b, two virtual path connections exist from user A to ATM B and from ATM B to user B. A node can use the VCI identifier on any virtual path connection (with the use of the VCI). Also, as suggested in Figure 8–1b, virtual path connections exist between the user and the network and within the network itself. This approach could be applied to the North American carrier system in which the VCI 14 could exist in one operator and VCI 23 could exist in another operator. This approach is a “clean” interface and also lends itself to relatively small lookup tables.

The VP and VC switches actually treat virtual path and virtual channel connections the same in the sense that they can be set up on a demand basis or they can be set up on a permanent basis. The network also can use QOS features, such as traffic usage, to associate with either a virtual channel or a virtual path.

Figure 8–2 provides an example of how VPIs and VCIs are translated and mapped at the switch. Figure 8–2a is an example of VPI



OR:






Figure 8–1 ATM switches: virtual path and virtual channel switches.

switching. The VCIs are passed unaltered through the switch, and VCIs 14 and 15 remain bundled in a VPI. The incoming VPI 7 is translated to outgoing VPI 4, but the VCI values are not altered.

In contrast, Figure 8–2b shows both VPIs and VCIs being translated and mapped to different values. VCI 14 (in the incoming VPI 7 bundle) is translated and mapped to outgoing VPI 7 with the VCI value changed to 23. The incoming VCI 15 (also bundled into VPI 7) is translated and mapped into VCI 88 and bundled into outgoing VPI 10.



(a) VPI translation



(b) VPI/VCI translation

Figure 8-2 VPI/VCI switching.

SPACE AND TIME SWITCHING

Switching has been defined for many years to encompass two realms: space switching and time switching (see Figure 8-3). In either case, the objective of the operation is to transport traffic from an input interface (port) to an output interface (port). The ATM switch functions like classical switching systems in that traffic is physically moved from




Figure 8-3 Space and time switching.

input ports to output ports. This concept is called space switching in the sense that different input and output lines are utilized.

The other realm is called time switching (shown in Figure 8-3b), and it has long been used in time slot interchange (TSI) switches. With time switching, the information and time slots are switched to different slots from the input interfaces to the output interfaces. The difference with ATM is that the time slots do not exist for identification purposes such as in a straight TDM switch. The slot is identified by an ATM label (VCI/VPI). Because ATM switches do not use TSI operations, it is possible that the staging switches could contend with each other for the same time slot. Therefore, ATM uses queuing to ameliorate this problem.

DIGITAL CROSS CONNECTS

ATM uses several of the concepts found today in digital cross connect equipment. The switch uses line and trunk identifiers to map incoming connections to outgoing connections. Figure 8-4 shows a line side and




Figure 8-4 A digital cross connect.

a trunk side of a cross connect switch. Each 64 kbit/s DS0 channel is associated with a DS1 number and a physical port (line) number. A mapping table correlates this channel with an outgoing channel. An end-to-end digital circuit is made up of a set of DS1 and DS0 at each physical interface between two end users.

Figure 8–4 shows the mapping table at the bottom of the figure. The top part of the figure shows two examples of how the cross connect table is used. The two sets of arrows should aid you in these two examples. In the first example, DS0 #1 of DS1 #1 on line 1 is mapped to trunk 2, DS1 #1 and DS0 #5. In the second example, DS0 #2 of DS1 #1 on line 1 is mapped to trunk 1, DS1 #2 and DS0 #2. As we shall see in the next section, ATM uses the line and trunk switching concept for its VCI and VPI mapping and switching operations.

THE SWITCHING FABRIC


Before delving into the details of ATM switching, several definitions are needed. First, a switching fabric describes the components of the switch, which include its hardware and software architecture. Second, a network element is a part of the overall switching fabric, such as a software or hardware component.

Figure 8–5 shows a general depiction of ATM switch operation. Later discussions examine specific implementations. The ATM machine receives a cell on an incoming port and reads the VCI/VPI value. This value has been reserved to identify a specific end user for a virtual circuit. It also identifies the outgoing port for the next node that is to receive the traffic. The ATM switch then examines a routing table for the match of the incoming VPI number and corresponding incoming port to that of an outgoing VPI number and corresponding outgoing port. Some implementations do not make use of a routing table, but rely on the switching fabric to be self-routing; these switches are explained later.

The header in the outgoing cell is changed with the new VPI value placed in a field of the cell header. The new VPI value is used by the next ATM switch to perform subsequent routing operations.

POINT-TO-MULTIPOINT OPERATION

Figure 8–6 shows a point-to-multipoint operation, also called a multicast or broadcast. For this example, the switch translates both the VPI and VCI, and the input/output interfaces are not shown. The concept is quite simple. The cross-connect table contains multiple entries. In this example, an incoming cell is sent to two output ports. The table contains the mapping values to change the VPI and VCI values in the cell headers.



VPI Virtual path Identifier

Note: VPI and VCI may be examined at the VCI/VPI switch

Figure 8-5 ATM routing operations.

Although not shown in the figure, the cross-connect table could have bidirectional significance. That is, either of the cells labeled VPI = 8, VCI = 5 and VPI = 2 and VCI = 9 could be mapped to a cell labeled VPI = 6 and VCI = 4.

MULTIPLEXING, LABEL SWITCHING, AND LABEL MAPPING

An ATM switch also performs multiplexing functions. The information from the input ports are multiplexed onto the output ports. The operations also include switching—the traffic from the input ports is switched to different output ports. While the cells are switched through the switching fabric from input to output ports, their VPI/VCI values, which are called labels, are also translated from one incoming value to




Figure 8–6 Point-to-multipoint operation.

another outgoing value. The headers uniquely identify the users' traffic, although the virtual circuit identifiers can be reused even if they are on different physical ports.

Figure 8–7 shows that an ATM switch is performing space switching by moving traffic from input ports a, b, c, d to output ports w, x, y, z. The switch is also performing header translation, which is sometimes called header mapping. This figure also shows the use of buffers (in this example, these are called output buffers because they exist at the output ports). These buffers are used to assure (most of the time) that cells do not arrive at the same output port at the same time.

The figure also depicts how the VPIs in the cell headers are used to determine the output port for the incoming cell. For example, a cell with VPI 1 arrives at port a. Routing operations reveal that the cell is to be relayed to output x to and its VPI changed to 5. As another example, a cell with VPI 3 arrives at port c. It is relayed to output port z and its label is mapped to VPI 8.





Figure 8-7 Routing and header translation.

Figure 8-7 summarizes the major functions of the ATM switch. It performs space switching, routing, multiplexing, and queuing. It also provides header mapping (header translation) by mapping the VPI values (and usually the VCI values as well).

Figure 8-8 shows how the VPI/VCIs are used at the ATM switches (nodes) to set up a connection and to create the routing table entries. In this example, three connections are created by three originating end users, by sending Q.2931 messages to the network. These messages contain a wide variety of information and are examined in Chapter 10, Call and Connection Control.

Among the fields in the Q.2931 connection request message is a destination address. This address is used by the ATM node to determine the



Note: Destination address (Dest = n) is sent to each ATM node to determine the route & is also sent to end-user node at terminating UNI

(2): Two connections are mapped at this node for end points ABC & DEF

Figure 8-8 Setting up the path through the ATM switches.

route that is to be established for the connection. Each node accepts the Q.2931 message, examines the destination address, and then consults a routing table to determine the next node that should receive the message.

A routing table is stored at each node and reflects the state of the network, the available bandwidth at each node, and other such information. It is updated periodically as conditions in the network change. Consequently, when the Q.2931 message is received by an ATM node, it knows the “best route” for this connection—at least to the next neighbor node.

These operations are dependent on how a vendor chooses to perform route discovery and maintain routing tables. The network may search for available bandwidth, and then set up the VPI/VCIs, or it may set them up on a link-by-link basis.

As each node sets up the calls, it reserves a VPI/VCI for each connection. Given an input VPI/VCI value, it selects an unused VPI/VCI

value for the output port. The next node receives this value, selects the route, and chooses the VPI/VCI values for its output port, and so on, to the final terminating UNIs. It is the job of the network to select values that are not being used on the same physical interface. This approach allows the VPI/VCI values to be reused.

After the connection is established, only the VPI/VCI values are needed, not the destination address. Figure 8–8 also shows that the originating users do not give the network a VPI/VCI at the originating UNI. The network assigns these values to the user, usually after the connection is completed and verified through the network to the terminating endpoint. This practice varies, but the ATM Forum UNI specification requires that the network take responsibility for assigning VPI/VCIs. The originating end user can be informed of the values it is to use when the network returns a connection confirm message to the user.

PROTECTION SWITCHING

For large ATM networks that must support many users, the issue of network robustness is paramount, and the network provider must devise ways to ensure that the customer's traffic is delivered safely. One method for providing this service is called protection switching, which uses alternate links between ATM nodes to transmit the ATM cells, in case the primary link fails. This concept is depicted in Figure 8–9.

Two links (such as a T3 or SONET link) connect two ATM nodes, node A and node B. One link is the working route (the working link), and the other link is the backup route (the protection link).¹ Let us assume traffic is coming into node A through port a. The cells are switched to port w or x, depending on the conditions of the links attached to node B from node A. In this example, a selector at node B, denoted in the figure by the , determines which link is to be used for the relaying of traffic from port a onto ports y or z (the traffic on port b is not germane to our discussion at this time).

Two types of protection switching can be implemented in the ATM node: (a) 1+1 protection switching, and (b) 1:n protection switching. With 1+1 protection switching, both links (routes) are used simultaneously to send the same ATM cells; that is, two copies of each cell is sent to node B.

¹The terms working copy and protection copy are also used to identify these concepts. These terms are found in SONET and SDH literature.




Figure 8-9 Protection switching.

This node uses the selector mechanism to decide which cell copy is to be relayed to ports y or z.

The concepts of 1+1 protection switching are shown in Figure 8-10. The incoming cell on port a has a VPI/VCI value of 1. The cross-connect table at node A is set up to direct this cell to both ports w and x, with the VPI/VCI of 1 mapped to VPI/VCI 7 across port w and VPI/VCI 8 across port x.

In the event of failure of the working link (the link attached to node A's port w and node B's port c), there is another copy of the cell available on the protection link (the link attached to node A's port x and node B's port d). The cross-connect table at node A is configured as shown in Figure 8-10 to achieve the duplicate transfer of the cell to node B, and the




Figure 8-10 1+1 protection switching.


selector at node B selects the cell with VPI/VCI = 8 for further relay. Obviously, the cell with VPI/VCI = 7 did not arrive at node B, and the cross-connect entry of port out = w, header mapping = 7 is superfluous (depicted in the figure by shading).

In contrast to 1+1 protection switching, 1:n protection switching uses a backup (protection) link for 1 to n working links. A common practice is to use the protection link to send traffic, so there are no duplicate cells on this link. In the event of failure of the working route, traffic is diverted to the protection link, and any traffic that was occupying the backup link must be either discarded, or held in a buffer until the working link can be restored. If restoration occurs quickly, the buffered traffic can continue its journey through the network. Otherwise, the traffic will be discarded. Consequently, the traffic on the protection link should be of a lower priority, and not subject to tight timing requirements.

In Figure 8–11, the protection link is transporting asynchronous traffic, such as available bit rate (ABR) or unspecified bit rate (UBR) cells. As long as the working route is operating satisfactorily, the protection link can carry the ABR/UBR traffic.


Also, this example shows the 1:n ratio of 1:1, but installations are free to set the number of working links that are supported by the protection link.

The cross-connect table at node A is configured in a typical manner, as shown at the bottom part of the figure.



ABR Available bit rate
UBR Unspecified bit rate

Figure 8–11 1:n protection switching: Working link operational.



ABR Available bit rate
UBR Unspecified bit rate

Figure 8-12 1:n protection switching: Working link is down.

The 1:n protection scheme for recovery operations is shown in Figure 8-12. The working link between nodes A and B is down (the link between node A's port w and node B's port c). The cross-connect table is used to divert the cells to the protection link. The ABR/UBR traffic on the protection link must be held until bandwidth is available for it to be transmitted, or it must be discarded.

The tradeoffs between 1+1 and 1:n protection switching are straightforward. 1+1 provides instantaneous recovery, without loss of traffic, yet requires an expensive redundant link for each working link. On the other hand, 1:n necessitates reconfiguration operations, with the resulting delay and probable loss of traffic, yet the sharing of the backup link translates into a less expensive network topology. The choice of these techniques is a matter for the network operator to decide.

LABEL SWITCHING VERSUS IP ADDRESS ROUTING

With the advent of Internet subnet operations, the routing operations to support diverse topology and addressing needs are greatly enhanced. As an added bonus, the 32-bit IP address space is utilized more effectively.

However, these features translate into a more complex set of operations at the router. Moreover, as the Internet and internets continue to grow, the router may be required to maintain large routing tables. In a conventional routing operation, the processing load to handle many IP addresses in combination with subnet operations can lead to serious utilization problems for the router.

We just learned that label switching uses the VPI/VCI identifier to route the traffic. The VPI/VCI can be used in place of the IP address for routing by encapsulating the IP datagram inside the ATM cell (in the 48-octet payload). The label is examined on a complete value, and not on a bit-by-bit closest match, as done in conventional IP routing. Many of these operations can be performed in hardware, which further speeds-up the process. A companion book in this series, *Advanced Internet Features*, examines this issue in more detail.

SWITCHING TECHNOLOGIES

This section provides an overview of several prominent switching technologies. Vendors often design their switches with a combination of the variety of technologies described here.

Shared Memory Switch

The shared memory switch, as its name implies, provides a common memory for the storage of the cells and the switching fabric (see Figure 8–13). As we learned earlier in this chapter, these cells are organized into separate queues and are routed based on the VPI/VCI values in the




Figure 8–13 Shared memory switch.

header. This illustration shows the use of a multiplexer and demultiplexer at the input ports and output ports, respectively. The cells arriving at the receiving multiplexer are placed onto a single line to the switch, and the switching functions move the traffic from the input queues to the output queues and to the output port(s).

Shared Bus Switch

Figure 8–14 shows a functional view of a shared bus switch architecture. It also shows the switch supporting both frame relay PDUs and ATM PDUs. In most implementations such as shown here, the switch is implemented with two busses for fully redundant operations. The busses actually operate in the cell mode, and traffic is diced into 48-octet pieces with a 5-byte header attached (some implementations are using different size cells). The destination card reassembles the traffic.

Cell traffic is carried through the bus as cells, therefore, both cells or frames can be sent across the bus. As just stated, frames are converted to ATM cells and vice versa, which allows frame and cell relay services to interwork.

Crossbar Switch

Crossbar switching is arranged in one or a combination of three architectures: (1) concentration (more input lines than output lines), (2) expansion (more output lines than input lines), and (3) connection (an




Figure 8–14 Example of a shared bus switch.




Figure 8–15 Crossbar switching.

equal number of input and output lines). In its simplest form, a circuit switch is a $N \times M$ array of lines that connect to each other at crosspoints (see Figure 8–15). In a large switching office, the N lines are input from the subscriber (terminals, computers, etc.) and the M lines are output to other switching offices. The switching fabric of a crossbar switch consists of N^2 crosspoints. Therefore, the size of the switch is limited.

Crossbar Switch with Arbiter (Knockout Switch). A variation of the crossbar switch is shown in Figure 8–16 [DUBO94]. It is known by several names in the industry—perhaps the best known name is the knockout switch. The switch operates as a conventional crossbar switch in that the horizontal input is connected to a vertical output. Notice that each input line is connected to each output line. In addition, each output port has an arbiter. This component receives cells from the vertical output lines, and places them into a first-in, first-out (FIFO) queue. The queue allows the arbiter to decide which cells gain access to the output port.

Multistage Switching

Crossbar switching is usually performed in more than one array or stage in order to reduce the number of crosspoints. If N lines are to be connected, then N^2 crosspoints are required. Multistage switches are economical for networks when $N > 16^3$ [AMOS79]. Multistage networks are designed with fewer crosspoints, yielding a more economical arrangement.




Figure 8-16 Crossbar switching with output arbiter (knock-out switch).

The input and output lines are divided into subgroups of N inputs and N outputs. The input first stage consists of $n \times k$ matrices. The k output is connected to one of the k center matrices. The third stage consists of $k \times n$ matrices connected from the center stage to the n outlets. The center stage matrices are all N/n by N/n arrays that permit connections from any first stage to any third stage.

The total number of crosspoints for this three-stage switch is [BELL82]:

$$N_x = 2Nk + k(N/n)^2$$

where: N_x = total number of crosspoints; N = number of input and output lines; n = size of each input/output group; k = number of center stage arrays.

Banyan and Delta Switching Networks

A technique that has been employed in the past in high-speed switches is the Banyan network [GOKE73]. As illustrated in Figure 8-17, this switch is characterized by only one path existing between an input to the final output port. With this approach, routing operations are quite simple and straightforward, but cells may be blocked (collide) if more than one arrives at a switching element at the same time.




Figure 8-17 The Banyan-Delta switching fabric.

Banyan networks are classified further into subgroups. An (L)-level Banyan network is one in which only the adjacent elements are connected, so each path passes through the same number of L states. A special Banyan implementation is a delta network. It is characterized by having an equal number of inputs lines to outputs lines, with each output identified with a unique value (for example, a destination address). Each digit in an address identifies a next stage of a switching element in the switching fabric.

A number of switching technologies have been proposed for ATM. Some of these proposals and implementations are based on the concept of single-state networks where a single stage of switching element is connected to the inputs and outputs of the switch. Multistage networks (in which several stages of switching elements are connected by various links to form the switching fabric) have also seen use. As of this writing, it is certain that ATM switching will be a melding of several switching techniques, and these techniques will form a switching element that is reproduced and interconnected to form the switching fabric.

Figure 8-18 shows a multistage delta network, called a 3-stage folded network. With this approach, traffic that arrives through the input lines (inlets) are examined by (in this example) three decision stages (switching




Figure 8–18 A 3-stage folded delta network.

elements). The traffic is passed to subsequent stages on output lines (outlets) based on an analysis of the address (which could be a tag used at the switch). Each stage is responsible for examining one bit in the address, typically from the most significant bit (MSB) to the least significant bit (LSB). In this example, a 3-bit address is used for simplicity.

The resulting address resolution is made one stage at a time in each bit of the address. As Figure 8–18 shows, the resultant output is based on the 3-bit address analysis of the three stages.

Figure 8–19 provides an example of an address resolution performed by the Banyan multistage switch. The bold line indicates the decision processes performed at the various stages. The destination address (010) is examined by each of the three stages on a bit-by-bit basis from the MSB to the LSB. The traffic, with its associated header, and address of 010 is routed to the output line identified as 010.

Collisions may occur in a delta network if cells arrive at a switching element at the same time. These cells may collide at an intermediate element, even though they are destined for the same output line, or they may collide at the final element. The probability of collisions occurring can be




Figure 8-19 Example of address resolution.

reduced by placing buffers in each switching element and allowing the buffers to queue arriving cells, which is an approach that is implemented in a number of high-speed switches. Alternately, the switch can operate at a higher speed than the inlets, thus giving an arrival time at the switching elements with no collision probability. Other mechanisms may be employed, and systems today can operate at a cell loss rate of 10^{-10} .

Figure 8-20 shows three alternatives for the use of buffers and queues within the switching fabric. The first approach places the buffers on the inlet side. Each inlet has a dedicated buffer, which is serviced on a first-in, first-out (FIFO) operation. Logic in the buffer determines when the buffer is to be served and when the cell is to be out-queued to the bus. This approach encounters collisions if two cells at the head of the queue in different buffers are destined for the same switching element. Also, if a queued cell is awaiting the availability of a line to a switching element, subsequent cells in the queue are also blocked even though an inlet may be available.

Output buffers are yet another alternative. This approach requires that a single cell can be served only by an outlet, which also could create




Figure 8-20 Switch buffers and queues.

output collisions and contention. Collisions can occur only if the switching fabric operates at the same speed as the incoming lines to the switching elements. In other words, several cells are contending for the same outlet. As discussed earlier, using different speeds within the switching fabric vis-à-vis the communications lines can reduce this problem. So, if collision is not to occur, cell transfer must be performed at n times the speed of the inlets, where n = the number of possible lines.

The last illustration shows central queuing, which entails nondedicated buffers between the inlets and the outlets; that is to say, a central memory element is used by all inlets to the outlets. This process does not work on the FIFO discipline, because the inlet traffic is placed into a single queue. To service the queues efficiently and fairly, there must be some mechanism to access the entries in the queues and to index the traffic in either sequential or random access operation.

Network switches can be developed that use parallel switching techniques. Figure 8–21 shows a parallel Banyan switch, also called vertical stacking. While switching techniques vary with parallel switching, as a general rule, the ATM cells that belong to the same connection are passed on the same plane of the multidimensioned switch. How this is determined varies among switch designers. During call establishment, decisions are made as to which plane will be used for the connection. The switch contains logic that examines each incoming call, places it on the appropriate plane, and then switches it to the appropriate output line. This approach has proved to be highly efficient, and it provides very high throughput.

VSLI AND ASIC-BASED SWITCHES

In the 1980s, the rise of high-speed and low-cost very large scale integrated (VSLI) circuits, application-specific integrated circuits (ASICs), and random access memory (RAM) chips paved the way for the use of ATM. With ATM, and its use of fixed cells, it was now possible to build very high-speed switches, non-blocking in design (by using RAMs). Much of this recent work has now focused on shared memory switch fabrics. The concept is illustrated in this Figure 8–22.²

The idea is for the switch to support different link speeds and interface types—in this example, OC-3, OC-12, and OC-48 rates. Of course, it

²This idea is found in many ATM switches, and we are highlighting work performed at Lucent (See *Bell Labs Technical Journal*, 2(2), Spring, 1997, “Advances in Shared-Memory Designs for Gigabit ATM Switching,” by Kai Y. Eng, and Mark A Pashan.





Figure 8-21 A parallel Banyan switch.

is impractical to build one switch for each link type, so the approach is to multiplex the incoming signals to the rate of the highest rate input, a concept called hierarchical multiplexing.

The output demultiplexer may be buffered to handle cells that leave the switch fabric that are destined for the same output link. However, much work has been accomplished in defining buffering operations in the switch fabric that minimize or eliminate the need for output multiplexer buffers. Additionally, while output buffering for each output link is the best approach from the standpoint of performance, it is costly. In contrast, a shared-memory technique reduces costs, and with the use of VLSI chips and RAM, can provide the desired performance.

The ATM switch processes the cells first through a bit-slicing converter, a wide data bus, and a large RAM. During a specific period, the incoming cells are written to RAM and the outgoing cells are read from RAM. The bit-slicing converter accepts cells for the WRITE operation for the input lines and converts the serial stream to a parallel stream (the



Note: Muxes also perform table look-up for cell routing
 SF Switch fabric

Figure 8-22 ATM hierarchical multiplexing switch.

bit-slicing operation), then multiplexes each cell (one after the other) onto the data bus in such a way that N arriving cells are written into different RAM locations. For the outgoing operations, as many as N cells are read from the RAM through the data bus, and converted back to a serial stream before being sent to the output links.

The fabric interface at the mux is also responsible for a table look-up for cell routing. The VPI or VCI in the cell header is used to determine the output link on the switch, and this information is placed in the cell show so it can be self-routing. With the Lucent approach, the VPI/VCI are replaced in the cell header as part of the conventional mapping/translation operation.

SUMMARY

The VPI and VCI are used by an ATM switch to make routing decisions. Switching is not part of the ATM standards, so vendors use a wide variety of techniques to build their switches. Several approaches to cell switching are now implemented, but research continues in this area.



Index



- Analog and digital systems, nature of, 23–37
- analog systems, 23–24
- analog-to-digital conversion process, 28–32
 - other coding schemes, 31–32
 - adaptive predictive DPCM, 32
 - differential pulse code modulation (DPCM), 31–32
 - sampling, quantizing, and encoding, 30–31
- bandwidth, 25–27
- broadband and base signals, 27–28
 - other definitions of broadband, 28
- classification of terms, 34–35
- cycles, frequency, and period, 24–25
- plesiochronous networks, 33
- synchronous clock hierarchy, 33–34
- timing and synchronization in digital networks, 32–33
- timing variations, 35–37
 - bit or clock slips, 36–37
 - slips—controlled and uncontrolled, 35–36
- Anchorage Accord, 20, 21
- Asynchronous transfer mode (ATM) basics, 98–133
 - ATM layers and OSI layers, 107–109
 - cell relay bearer service (CRBS), 130–133
 - interfaces, 102–106
 - labels, 129–130
 - layers, 106–109
 - PDU(s) (cells), 114–117
 - rationale for cell size, 117–128
 - relationship of AAL, ATM, and network, 109–112
 - review, 98–99
 - topology, 99–102
 - typical protocol stacks, 112–114
- ATM adaptation layer (AAL), 134–180
 - AAL type 1 (AAL1), 141–142
 - AAL type 2 (AAL2), 146–152
 - AAL type 5 (AAL5), 175–176
 - AAL types 3, 4, 3/4, and 5 for data, 161–166
 - AAL/ATM primitives, 177–180
 - AAL1 PDU, 142–146
 - AAL3/4, 166–174
 - available bit rate (ABR), 176–177
 - dividing CS into further sublayers, 140–141
 - naming conventions, 141
- ATM adaptation layer (*cont.*)
 - original AAL types 3 and 4 (AAL3, AAL4), 166
 - principal tasks of, 134–136
 - rationale for AAL types, 140
 - sublayers, 136–138
 - traffic, classes of, 138–139
 - voice packetization, 152–161
- ATM and SONET, 85
- ATM market, 431–434
 - ATM and IP, 432–433
 - ATM in residential broadband, 433
 - future of ATM, 433
 - recent successes and projections, 431–432
- ATM switching operations, 181–206
 - ATM switching, 181–184
 - digital cross connects, 185–187
 - label switching *versus* IP address routing, 195–196
 - multiplexing, label switching, and label mapping, 188–192
 - point-to-multipoint operation, 187–188
 - protection switching, 192–195
 - space and time switching, 184–185
 - switching fabric, 187
 - switching technologies, 196–206
- ATM, principal specifications for, 18–20
- Bandwidth, 25–27
- Broadband integrated services digital network (B-ISDN) model, 87–97
 - ATM and B-ISDN model, 89–93
 - B-ISDN functions, 93–95
 - B-ISDN service aspects, 95–97
 - ISDN and B-ISDN, 87–89
- Broadband ISDN, 16–18
- Buffers, 35
- Call and connection control, 260–299
 - ATM address, 263–266
 - ATM connections on demand, 260–263
 - connection control examples, 270–278
 - connection control messages, 266–267
 - connection setups and clears, 267–269
 - functions of Q.2931 messages and information elements, 279–287
 - PNNI model, 292–294
 - Q.2931 messages, examples of, 287–292
 - Q.2931 timers and states, 269–270
 - UNI signaling, voice, and telephony over ATM to desktop, 295–299
- Cell, 12
- Change of frame alignment (COFA), 36
- Channel banks, 30
- Client operations, 51
- Condensations, 24
- Currents, 24
- Cycle, 24
- Data networks and protocols, typical, table, 4
- Datagram, 51
- Differential pulse code modulation (DPCM), 31–32
 - adaptive, 32
 - adaptive predictive, 32
- Emerged technologies, 58–86
 - ATM and SONET: reduction or enhancement, 85
 - comparison of switching systems, 58–60
 - integrated services digital network (ISDN), 70–79
 - signaling system number 7 (SS7), 79–85
 - T1/E1 systems, 60–64
 - X.25, 64–69
- Fast relay networks
 - and ATM, 11–14
 - and SONET, 15–16
- Fragmentation, 53
- Frame, 51
- Frames, 11
- Frequency (*f*), 24
- Gateway, 50
- Hertz (Hz), 24
- Hold-and-forward technology, 59
- Integrated services digital network (ISDN), 70–79
- Internet protocols (TCP/IP), 49–56

- Internetworking with ATM networks, 300–330
 - ATM B-ISDN intercarrier interface (B-ICI), 308–312
 - ATM network as backbone, 300–301
 - broadband low-layer information element, 304–306
 - frame UNI (FUNI), 327–328
 - multiprotocol over ATM (MPOA), 328–329
 - network-to-network interface (NNI), 307–308
 - RFC 1483 and RFC 1577, 322–327
 - specific internetworking services, 312–322
 - using Q.2931 to support protocol capability (tunneling), 301–304
- Jitter, 35
- Layered protocols: model for ATM and SONET networks, 38–57
 - ATM and OSI layers, 48–49
 - ATM and the Model, 47
 - concept of service provider, 41–46
 - encapsulation/tunneling, 45–46
 - internet protocols (TCP/IP), 49–56
 - internet layers, 51–52
 - IP functions, 52–54
 - TCP operations, 55–56
 - OSI network and transport layer, 56–57
 - protocol entities, 47
 - protocols and OSI Model, 38–41
 - service access points (SAPs), 47–48
- Open Systems Interconnection (OSI) Model, 39
- Operations, administration, and maintenance (OAM), 363–405
 - ATM functions at U- and M-planes, 367–372
 - ATM management information bases (MIBs), 388
 - ILMI MIB and ATM MIB, 403
 - integrated local management interface (ILMI), 389–403
 - layer management/ATM primitives, 403–405
 - network management model, 363–364
 - OAM at ATM layer, 381–388
 - OAM headers, 375–381
 - operation and maintenance (OAM) operations, 364–367
 - other ATM Forum specifications for network management, 405
 - SONET OAM functions, 372–375
- Packet switching, 2
- Period (T), 25
- Phase variation, 35
- Physical layer services for ATM, 406–430
 - 2.4 gbit/s physical layer specification, 425
 - ATM mapping into 100 mbit/s multimode fiber interface, 414–416
 - ATM mapping into 155.52 mbit/s private UNI, 416–418
 - ATM mapping into DS3, 410–412
 - ATM mapping into SONET STS-3C, 408–409
 - ATM/physical layer primitives, 407–408
 - circuit emulation service interoperability (CES-IS) specification, 412–413
 - internetworking ATM and SONET, 423–424
 - inverse multiplexing for ATM (IMA), 425–430
 - mapping DS1, DS3, CEPT payloads into SONET frames, 419–423
 - physical layer options for ATM, 406–407
 - private UNI for 51.84 mbit/s and subrates, 418–419
- Plesiochronous networks, 33
- Primary multiplexers, 30
- Protocols, 38
- Rarefactions, 24
- Router, 50
- Segment, 51
- Server operations, 51
- Service access points (SAPs), 47–48
- Service provider, concept of, 41–46
- Signaling system number 7 (SS7), 79–85
- Slips, 35–37
 - bit or clock, 36–37
 - controlled and uncontrolled, 35–36
- Sound wave, 24
- Store-and-forward technology, 59
- Subnetworks, 50
- Synchronous clock hierarchy, 33–34
- Synchronous Optical Network (SONET), 331–362
 - automatic protection switching (APS), 344–346
 - equipment, 356–362
 - foundations for, 333–335
 - layers, 343–344
 - mapping and multiplexing operations, 351–356
 - optical fiber, 336–337
 - payload pointers, 350–351
 - payloads and envelopes, 346–350
 - pertinent standards, 337–338
 - present transport systems and, 333
 - progress in penetration, 362
 - purpose of, 331–333
 - synchronous networks, 335–336
 - typical topology, 339–343
- T1/E1 systems, 60–64
- Telecommunications technologies, introduction to, 1–22
 - Anchorage Accord, 20
 - applications use of ATM, 14–15
 - broadband ISDN, 16–18
 - fast relay networks and ATM, 11–14
 - fast relay networks and SONET, 15–16
 - present and future requirements, 5–8
 - present infrastructure, 1–5
 - principal specifications for ATM, 18–20
 - virtual companies and virtual networks, 8–11
- Timing variations, 35–37
- Traffic management, 207–259
 - ABR operations, examples of, 251–258
 - allocation of bandwidth, 215–219
 - at UNI, basic concepts, 223–231
 - ATM bearer service attributes at UNI, 231
 - ATM cell transfer performance parameters, 234–236
 - ATM Forum and ITU-T traffic control and congestion control, 238–244
 - ATM layer provisions for quality of service (QOS), 236–238
 - ATM service architecture, 249–251
 - ATM services for LAN and internet traffic, 244–249
 - buildout delay procedures at receiving endpoint, 258–259
 - cell arrival rate and cell interval, 232–234
 - connection admission control (CAC) procedures, 222–223
 - control mechanisms, 212–214
 - guaranteed frame rate (GFR), work on, 259
 - in an ATM network, 207–212
 - queue management operations, example of, 219–220
 - traffic control and congestion control, 214–215, 231–232
 - usage parameter control (UPC), 223
 - variable delay, dealing with, 220–222
- Velocity (V), 25
- Virtual network/virtual private network (VPN), 8
- Wander, 35
- X.25, 64–69

SECOND EDITION
ATM VOLUME I
Foundation for Broadband Networks

UYLESS BLACK

All you need to know about ATM:

If you want to understand how ATM fits into today's state-of-the-art WANs and LANs, look no further. In this book, best-selling author and world-renowned communications consultant Uyless Black explains all you need to know: architecture, switching elements, traffic management, and much more. This brand-new second edition covers many important new ATM enhancements, including MPOA, LAN Emulation, Frame Based ATM, Layer 3 Switching, even Wireless ATM. Learn all you need to know to get results, including:

- ▶ The fundamental architecture of ATM networks, including the new Signaling ATM Adaptation Layer
- ▶ Using Frame User Network Interface (FUNI) to improve the utilization of access line bandwidth
- ▶ Using MPOA to integrate ATM with TCP/IP, Ethernet, and other key LAN protocols
- ▶ ATM network management—including the ATM Network Management Model and the Anchorage Accord specifications
- ▶ New protection switching techniques for better network backup

This new edition combines extensive, real-world implementation examples, up-to-the-minute coverage, and Uyless Black's unique insight into the issues that matter most to communications professionals. If you want to get up to speed with ATM technology, this is the place to start.

Also look for the *ATM Resource Library*, a boxed set which includes these great books by Uyless Black...

- ▶ *ATM, Volume I: Foundation for Broadband Networks, Second Edition*
- ▶ *ATM, Volume II: Signaling in Broadband Networks*
- ▶ *ATM, Volume III: Internetworking with ATM*

About the Author

UYLESS BLACK is a widely known and respected consultant and lecturer on computer networks and data communications. He is author of all the books in his Prentice Hall Series in Advanced Communications Technologies, including *ISDN and SS7*, *Second Generation Mobile and Wireless Networks*, and *Advanced Internet Technologies*.

Prentice Hall
Upper Saddle River, NJ 07458
<http://www.phptr.com>

ISBN 0-13-083218-9



9 780130 832184