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.

[1]T1 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	Propriety	LAN
Ethernet	Xerox, Digital, Intel	17 years	Standard***	LAN
Netware	Novell	12 years	Proprietary	LAN

[1]Approximate ages; all have evolved and have been enhanced
[2]WAN 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 lead 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-

[2]The 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 business 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

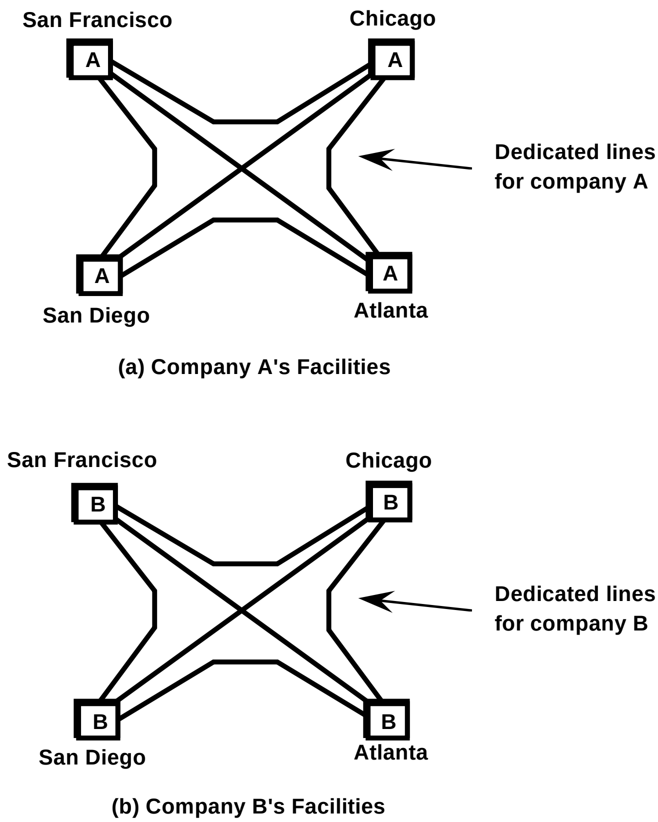

(a) Company A's Facilities

(b) Company B's Facilities

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]

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

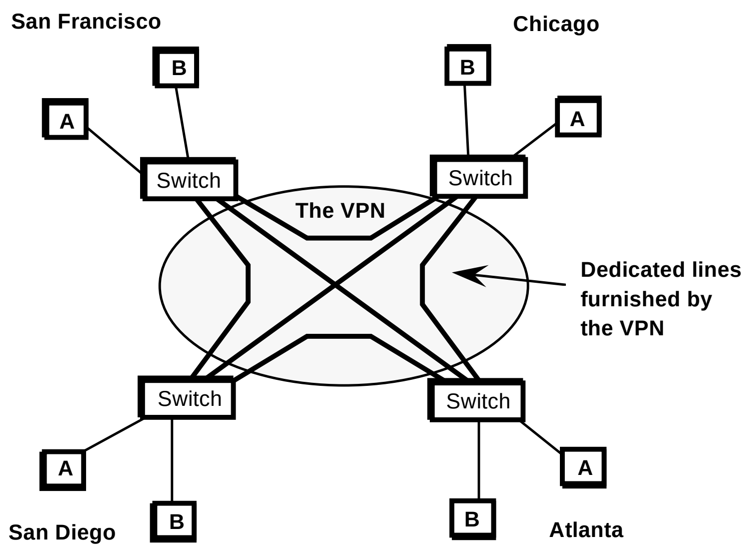

Figure 1–2 Virtual private networks.

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

[4]I 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.

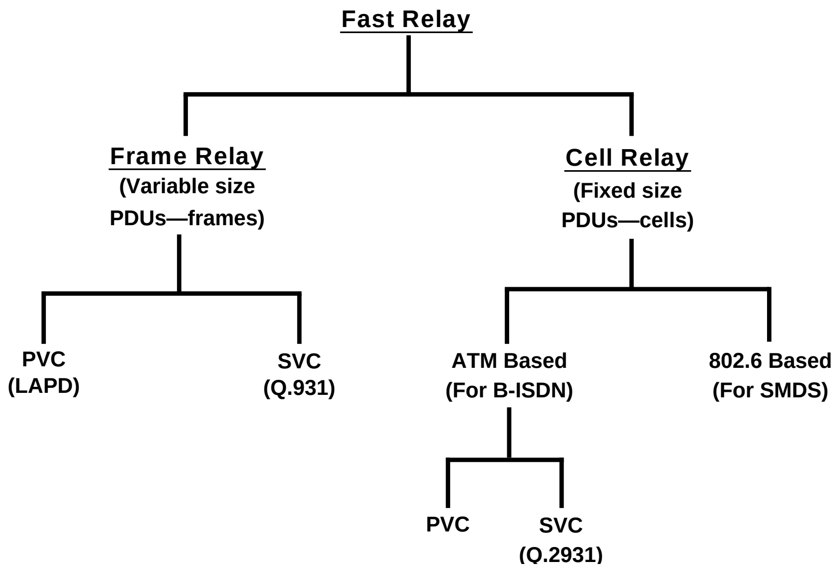

Figure 1–3 Types of relay systems.

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 remaining 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

[5]A 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 second/8000 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.

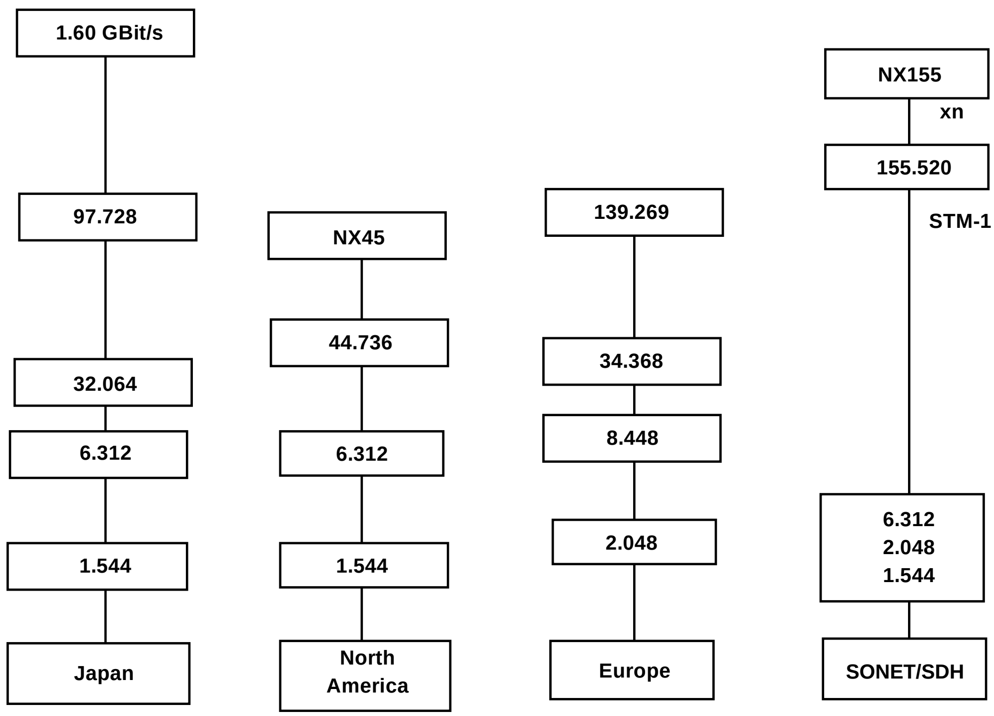

Note: Unless noted otherwise, speeds in Mbit/s

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

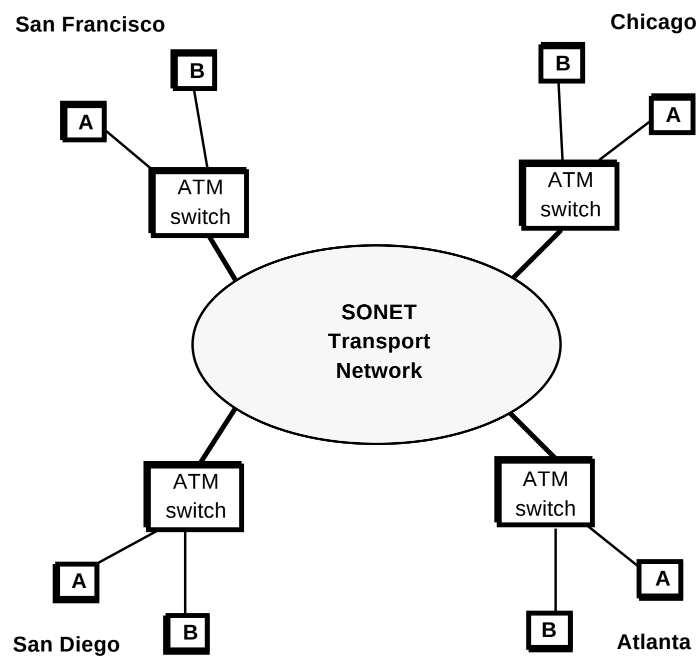

Figure 1–5 ATM and SONET.

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]

[1]Although 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-

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

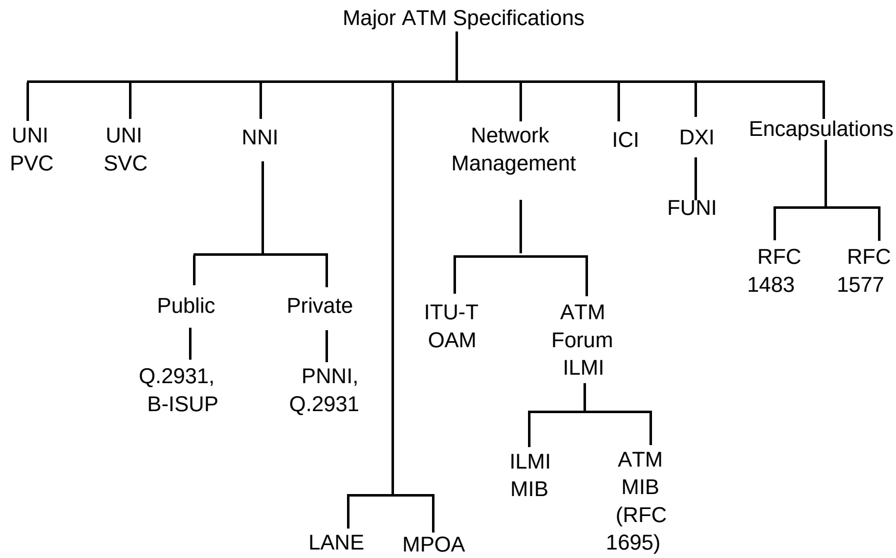

Where:

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 sparcity, 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.