NEAR EAST UNIVERSITY

Faculty of Engineering

Department of Electrical and Electronic Engineering

BROADBAND ISDN

Graduation Project
EE- 400

Student: Hani Agha (980910)

Supervisor: Professor Dr. Fakhreddin Mamedov

Lefkoşa - 2001

TABLE OF CONTENTS

ACKNOWLEDGMENT

Being at the end of a hard journey in my life; in which my future was being build. But this wouldn't have happen unless the help from almighty Allah who I always thank for creating reasons for this success in my life.

One of the major reasons was having a great parents who hasn't only this favor for me but a lot although my entire life.

Prof.Dr.Fakhreddin Mamedov was a special person in my study; I admire his way and his wide knowledge. Also admire learning from him, this I will always do. And I want to thank him also at this point

This term paper was also a chance for me to know some people closely, they were a very good help for me and they thought me what a real brother is. I thank them all and I hope to be able to help them or those who are like them some day.

Finally, But not final I hope that all my collage mates have success in all their life.

ABSTRACT

ISDN stands for Integrated Service Digital Network, and as the name suggests it allows digital communication. This is favorable as digital technology is a lot faster, and more accurate than the old analogue lines as they no longer require the process of modulation and demodulation. ISDN relies on already existing copper cable systems, causing its integration into our existing communications system to be smoother and less disruptive.

Narrowband ISDN has been designed to operate over the current communications infrastructure, which is heavily dependent on the copper cable. B-ISDN however, relies mainly on the evolution of fibre optics. According to CCITT B-ISDN is best described as 'a service requiring transmission channels capable of supporting rates greater than the primary rate.' Behind this statement lies the plan for a network and services that will have far more impact on the world we know today, than ISDN ever would.

When ISDN is referred to as a network it is to be considered a telephone network, not a computer network. Broadband ISDN allows its users to communicate over high speed, high quality digital channels. The media is supports include Telex, fax, voice telephone, video telephone, audio, high definition TV and computer networking.

In the past video, audio, voice and data services needed different types of communication channels. One of the main advantages of ISDN is the ability to integrate these features over the same network and cable plant. Not only is this possible using ISDN technology but the quality of the transmission is better also. In the past four networks were needed and video was distributed on coaxial lines, audio over balanced lines, voice used copper cable pairs and data services required coaxial or twisted pair cables. Using one network allows reductions in installation costs, as well as easier installation. Other features available include demand networking, automatic bandwidth and on the fly connectivity. Advances in the services available are due to ISDN being digital.

INTRODUCTION

As opposed to the B-ISDN, the present ISDN networks are referred to as Narrowband ISDN (N-ISDN). In the N-ISDN, the rates of information transfer speeds are based on the basic rate of 64kbit/s, ranging up to 1536/1920kbit/s. In the B-ISDN, the maximum transfer rates will be two orders of magnitude higher, at 155.52Mbit/s or 622.08Mbit/s.

The reason why these transfer rates are needed is that business communications and office automation through workstations and LANs is expected to make severe demands on the communications network. In particular, the B-ISDN with its ATM technology will be an efficient way to implement multimedia communications, with video information added to the conventional voice and data information.

Also in the background is recent progress in optical technology. It is now possible to construct high-quality, low-cost communication channels directly to the user's home or workplace, raising the possibility of B-ISDN services such as High Definition Television (HDTV) broadcasting offered at 155.52/622.08Mbit/s. These developments in very-high-speed switching are due to progress in optical-fiber, high-integration LSI and ATM switching technology. The above was a brief introduction of Recommendations for the B-ISDN, but many problems remain to be solved before the introduction of actual B-ISDN systems. In particular, many details have to be resolved for the Physical Layer, the ATM Layer and the protocols for the various kinds of Adaptation Layers. More work is also needed to determine the exact expression of variable and fixed-rate communications between terminals and networks, and the standardization of service classes and quality control.

The greatest need for multimedia communications is among business users, so the first applications of ATM technology for the B-ISDN will probably come in the form of dedicated networks for business communications. ATM technology is the perfect vehicle to transmit the mixture of voice, data and video information generated by applications of this kind.

In pace with the internationalization of the Japanese economy, more and more customers are choosing KDD's ISDN services as the base on which to construct international networks between Japan, Europe and North America. KDD will continue to strive to make the international ISDN more flexible, less expensive and easier to use. In order to meet the growing demand for multimedia communications, and high-speed, broadband services, we will continue research into ATM technology and the B-ISDN.

When ISDN is referred to as a network it is to be considered a telephone network, not a computer network. Broadband ISDN allows its users to communicate over high speed, high quality digital channels. The media is supports include Telex, fax, voice telephone, video telephone, audio, high definition TV and computer networking.

Most of the applications for ISDN have reached the extent of their development, and now the focus has shifted to services that an be provided across broadband ISDN cables. The ITU-T defines the services and associated standards of ISDN communications, have recommended the two service area for application with BISDN, Interactive Services, and Distribution Services.

1. WHAT IS BROADBAND NETWORKING? WHY DO WE NEED IT?

1.1 SUMMARY

Three "dimensions" may characterize all networks:

(1) Their bandwidth or data rate,

(2) Their switching characteristics;

(3) Their degree of "intelligence" (i.e., their level of automated functionality).

For the purposes of this project, broadband networks are defined, using these dimensions, as switched Intelligent networks with data rates at 45 Mbps or above.

There are four "megadrivers" that will lead to the need for such networks. These are:

(1) The transition to an information-based postindustrial economy,

(2) The growing strategic importance of information to business,

(3) The ever-growing power of computers and storage devices, and

(4) The coming video and imaging revolution.

Although broadband may lie in the future, we are already beginning to see the appropriate technologies and standards emerge. However, several factors still stand in the way of the development of broadband, including an immature standards environment, uncertain prices, and a lack of familiarity with broadband concepts on the part of the end user.

1.2 The Three Dimensions Of Broadband Networking

For our purposes here, a useful way of characterizing different kinds of networking is according to :

- Bandwidth or data rate
- Switching characteristics
- Network intelligence

The bandwidth (an analog measure) or data rate (a digital measure) of a network is really the amount of information that it can carry in a given period of time. Its switching characteristics define how information is carried from one point of the network to another.

Finally, there is "network intelligence," which is not a well-defined concept but has something to do with how automatic the features of a network actually are:

The essence of network intelligence is software control. (Figure 1-1) shows how the three dimensions of broadband networking are interrelated.

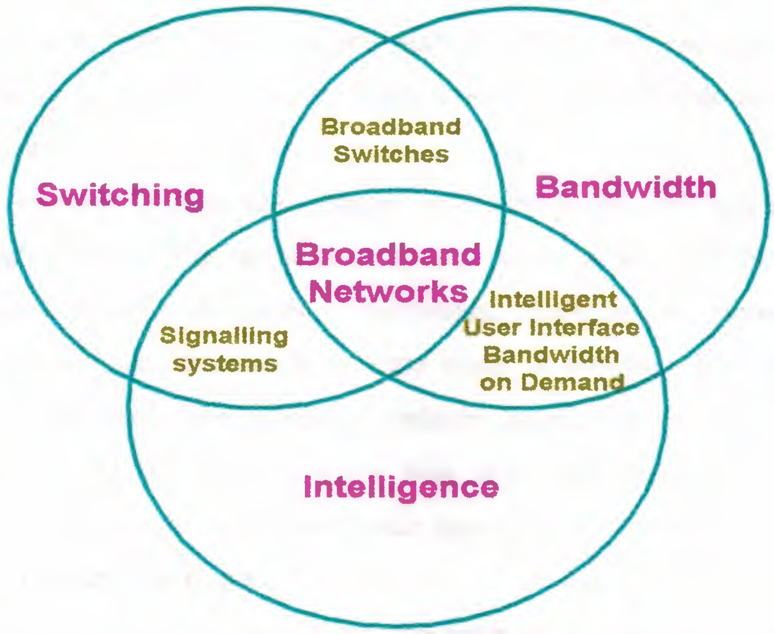

Figure 1.1: The Three dimensions of Broadband Networking

1.2.1 Bandwidth, Data Rates, and Broadband

The classic way to define information carrying capacity in analog networks is in terms of bandwidth-the difference between the higher and lower frequencies in the signal carrying the information. This is measured in Hertz (Hz), thousands of Hertz (kHz), millions of Hertz (MHz), or billions of Hertz (GHz). With the shift to digital networks and most of the networks talked about in this project are digital networks the references to frequency have been replaced by references to the number of digital "bits" the network can carry per second. This is measured in bits per second (bps), thousands of bits per second (kbps), millions of bits per second (Mbps), or billions of bits per second (Gbps).

For the most part we will be talking about bps, kbps, Mbps, or Gbps. Strictly speaking these are measures of "data rates" rather than "bandwidth" (which, as mentioned before, is analog terminology), but because both measure the information carrying capacity of a network, the terms "bandwidth" and "data rate" tend to be used interchangeably Technically speaking, this is an error.

By the time you have finished reading this project, you will have a pretty good idea of how many bits per second are required to support certain kinds of services, but to give you an initial idea, 9.6 kbps that is 9600 bps is the current standard for communications using a fairly ordinary personal computer. It is suitable for most text-oriented communications applications such as electronic mail, but to support some of the more exciting applications that are currently receiving publicity in the communications press, one needs to go into the realm of hundreds of kbps.

For example, a reasonable quality videoconferencing link needs about 300 to 400 kbps, although videoconferencing of lower quality can be achieved at 56 kbps. With 56 kbps, video telephony, showing the heads of the callers, would operate successfully assuming that the "talking heads" did not move around too much. A 400 kbps link would be adequate for carrying an entire conference involving multiple participants on both sides. However, the picture quality could still be well below what one might expect from a regular television broadcast, and such a video- conference once again assumes that there is a relatively low level of motion by conference participants.

For broadcast-quality digital video, we have to move into realms of tens of Mbps that is tens of millions of bits per second, and for certain applications, mostly associated with scientific visualization and other rather specialized supercomputer-related applications, a speed of Gbps billions of bits per second is required.

Implicit in the very word "broadband" is the idea that broadband networks have high bandwidths or data rates. But just how high? Some definitions of broadband actually leave this question moot. In its original meaning, the word "broadband" referred to the technology used in cable television (CATV). This analog technology allows multiple channels to be carried at different frequencies on the same coaxial cable. (This technique is called frequency-division multiplexing.) Implicit in this use of the word "broadband" is the idea that if a network is capable of carrying multiple video channels, it would have to have a high capacity because video (whether analog or digital) is "bandwidth hungry." and Electronic Engineers (IEEE) for broadband LANs (IEEE 802.3 1OBROAD36 and IEEE 802.4 Broadband) did not specify data rates higher than the "baseband" alternative. However, broadband LANs, unlike base-band LANs, were capable of the practical transmission of reasonably high-quality (analog) video, which might at a pinch be said to indicate that they had a higher information carrying capacity in an intuitive sense.

Be that as it may, the use of "broadband" to mean something to do with cable television is a little anachronistic and broadband LANs are not widely used these days. Unfortunately, the old meaning of "broadband" has left us with no real answer to the question we are trying to answer in this section: Just how big a bandwidth makes a network broadband?

In his project Broadband Coding, Modulation and Transmission Engineering Bernard Keiser asks the question, "What is a broadband communications system?" and answers this question with the following: "[a broadband communications system] is any telecommunications system capable of conveying information bandwidths in excess of voice bandwidth, which is often taken to be 3 kHz. In terms of digital transmission, a precise definition cannot be given, but the term might be taken to include digital transmission at rates in excess of 9.6 kbps, since the 9.6 kbps rate can be handled readily by a voice band modem."

However, Mr. Keiser's definition is surely too weak, because it would tend to include a huge range of networks such as many X.25 networks, ISDN networks, and so on, which are not normally included these days when broadband networks are discussed. But the Keiser definition does offer us a clue about where we should be going with our definition. The usual data rate associated with the transmission of voice over digital channels is 64 kbps, so, taking Keiser's suggestion that broadband networks are those that operate "in excess of voice bandwidths," we end up with a definition of broadband that says that broadband networks are those that operate in excess of 64 kbps.

It turns out that we are still being too inclusive, but we are getting nearer to a definition of broadband. Simply calling a broadband network any network that is capable of carrying data at rates above 64 kbps includes networks based on higher bandwidths built up from 64 kbps (or, what is much the same thing, 56 kbps) channels. In particular, it includes ISDN. ISDN the Integrated Services Digital Network is a set of international standards designed to allow the telephone companies to migrate their mainly analog, mainly voice-oriented networks to networks that are capable of supplying advanced voice and data services over a digital fabric while retaining the basic physical infrastructure of the existing telephone network. ISDN has been deployed in a moderate way in the United States and supplies subscribers with bandwidths of up to 1.5 Mbps by offering them two interface "packages" that consist of 64 kbps "B" channels bundled together.

An ISDN end user who simply wants to plug in a terminal subscribes to the ISDN "Basic Rate" package, which offers two B channels a total of 128 kbps. She may use this entire bandwidth for one application or may use each of the B channels for a separate function. For example, in a videoconferencing application one B channel may carry voice and the other B channel may carry data. The other ISDN package is the "Primary Rate" package, which offers 23 B channels and is intended for connecting nodal equipment (such as PBXs, LAN bridges, nodal processors, etc.) to the network. In both packages, in addition to the B channels, there is a D channel, which is used for controlling the network rather than transmitting information for the end user. The structure and applications for both the Basic Rate interface (BRI) and the Primary Rate interface (PRI) are profiled in (Table 1-1).

TABLE 1-1: Basic Rate And Primary Rate Isdn Connections: Structure And Applications

Interconnection	Structure	Applications
Basic Rate interface (BRI)	Two circuit-switched 64 kbps B channels and one 16kbps packet-switched D channel. (Occasionally a package of one B channel and one D channel is avail-able)	Terminal interconnection on the B Channels Signalling on the D channels. D Channel may also be used for telemetry or data communications
	BRI operates at a data rate of 144 kbps (plus 8 kbps for Framing, synchronization, And other overhead bits)	BRI is designed to meet much the same needs as the analog service that is provided through today's telephone jack. BRI service is what most residential and small business ISDN users are intended to subscribe to
Primary Rate interface (PRI)	Twenty-three 64 kbps B channels and one 64 kbps D channel Total data rate 1.544 Mbps	Interconnection of PBXs, multi-plexers, routers, and other CPE products that concentrate or switch telecommunications traffic to the public segment of the ISDN over B channels Signalling over D channels

ISDN plays a major role in the story that we will be telling in this project, but we will not classify ISDN as a broadband network. The reason for this is that it fudges over the distinction between a networks that has a high bandwidth made up by simply concatenating channels of lower bandwidth and a network that handles a single large chunk of undivided bandwidth. It is this latter kind of network that we have in mind when we use the term "broadband." With broadband networks defined in this way, it is typically the case that one can take this large chunk and split it up in smaller chunks, but it is the big chunk that is the fundamental "unit of transmission.

In taking this approach, we will be following in the footsteps of the major U.S. and international standards bodies, who are currently defining a successor to ISDN, which will be known as Broadband ISDN (B-ISDN).

There are a lot of other differences between regular (or "narrowband" ISDN) and B-ISDN apart from the bandwidths. B-ISDN standards makers have borrowed as much as they can from the earlier ISDN standards, but B-ISDN requires an entirely different communications infrastructure. Narrowband ISDN is specifically intended to operate over today's communications infrastructure, which is still heavily dependent on copper cable, but B-ISDN uses for the most part fiber optics. And where narrowband ISDN uses a "circuit-switched" technology, like the public telephone network, B-ISDN will use a "packet-switched" technology; called asynchronous transfer mode, like a computer network. In a circuit-switched network, a permanent physical circuit is maintained between the users of a network during a conversation. In packet-switched technology, information is broken down into packets and transmitted over either a "logical" or a "virtual" circuit established by the network or by sending packets over entirely different routes.

We shall return to the differences between circuit switching and packet switching in, however, we suggest, on the basis of the discussion that has gone before, that a broadband network can be loosely defined as a network that supports "unchannelized" data rates considerably in excess of those possible over narrowband ISDN. The term "unchannelized" here refers to the idea that the high bandwidths are not achieved just by joining together channels of smaller bandwidth. In particular, we shall discuss as broadband any network that operates at DS3 rates (45 Mbps) or above. The reason for taking this approach is that there is a sense, that broadband is leading edge technology and in terms of delivering bandwidths within public or semipublic networks, 45 Mbps (DS3) is currently the leading edge.

Treat these comments with a small pinch of salt. We will not be sticking to these definitions in any hard way. Definitions, after all, are only conventions and are good only as long as they remain useful. In particular, many trials of networks now going on at relatively low data rates still have important implications for broadband. In order to prove certain points about broadband networks we will mention such trials from time to time.

Also, the reader should note the important distinction between "network speed" and "access speed." Network speed is aggregate speed of the network itself. Access speed is the speed that the user gets to "see" at his terminal. Network and access speed may be equal. This is the case with most LANs, but it is certainly not always the case. When you dial up the public telephone network over your 9.6 kbps modem (i.e., at an access speed of 9.6 kbps), your message may be transmitted over certain backbone routes in the public network operating at 2.4 Gbps or even higher rates.

1.2.2 Switching and Broadband

Switching is a concept fundamental to networking, but much like the concept of broadband, it is hard to pin down precisely. For our purposes, if a network only allows information to be sent between two fixed points (a communications environment known as point-to-point link), it is said to be "unswitched." We will also say that a network is unswitched if information is broadcast to multiple points making no discrimination among those points that is, all points receive the same information. Such a communications environment is characteristic of satellite broadcasting.

If the information can be routed between selected points in a network on an as needed basis we are, dealing with a switched network. Many of the devices intended to do switching are mailer called switches, which makes things pretty clear. But other devices, called by their vendors multiplexer, routers, hubs, or cross connects, also perform the switching function in the sense switching is meant here. In fact, switched networks do not necessarily have to have an identifiable of switching device. This is most notably true of LANs.

There are, in fact, many approaches to switching. We have already mentioned circuit switching and packet switching, and, as promised earlier, we will be explaining these two fundamental categories of switching e only in more depth later. For the time being, however, we are going to try to define broadband networking in terms of its switching as still characteristics. As with broadband's bandwidth characteristics, there are no hard and fast answers.

Can we have broadband networks that are totally unswitched? For the purposes we are going to answer that question in the negative. The reason, again, is that we are trying to examine leading edge speed technology, and point-to-point (i.e., unswitched) links operating at high d that bandwidths are not particularly leading edge.

Point-to-point trunk lines operating at a few gigabits per second are now quite common in the public telephone network. High-bandwidth satellite broadcasts the other 3 kbps kind of unswitched communications environment we have defined-are also nothing special these days.

But if broadband networks are inherently switched, what kind of switching are we talking about? We will leave a detailed answer to this question to a later chapter. Suffice it to say here that more than just the traditional categories of switching are involved. As we shall see, the broadband field is heavily reliant on new switching technology; especially a stripped-down version of packet switching called fast packet. At the physical level, broadband networking may also derive some benefits from poses, the emerging area of optical switching, a field that promises very fast points switching (and possibly also very fast computing) at some time in the relatively near future.

1.2.3 Broadband Intelligence

Finally we come to "intelligence," the third and final networking dimension we are considering here. In some ways this is the hardest of the three dimensions to pin down. By way of a reminder, "intelligence," in the sense that we are using the term here, has little to do with the way that we apply the term to a human being or even a dog. Instead, it has to do with the degree to which the functions of the network are automated to meet end user needs. For example, can you:

- Automatically change the bandwidth coming out of your wall plug?
- Automatically order a service from your terminal and then cancel it when you no longer need this service?
- Automatically reroute (all or some) calls to another number?

- Identify incoming callers and automatically call up information on those callers from a database?
- Automatically call up statistics on network usage and performance?

If the answer to at least some of these questions is yes, then you are dealing with an intelligent network, in the sense that we are using the term here, and in this sense broadband networks must be intelligent. Broadband networks are supposed to serve the needs of many different kinds of end user: consumers in their homes; large, sophisticated businesses and small retailers; and manufacturers and service firms.

This covers a lot of ground, so a broadband network must be designed so that it is highly flexible and can respond to user needs rapidly and efficiently and most important of all, at the user's request. This means intelligent networks involving sophisticated software.

1.3 Four Broadband Megadrivers

So far, we have not really said that much. We have defined what broadband networks are but have said nothing about why anyone would want or need such networks. This is obviously a critical point.

There are four general trends that will lead increasingly to the deployment of the high-bandwidth intelligent switched networks that are broadband networks.

These broadband megadrivers, which are somewhat dependent on each other, are as follows:

- The transition to an information-based postindustrial economy
- The growing strategic importance of information to business
- The ever-growing power of computers and storage devices
- The coming video and imaging revolution

The ways in which these four broadband megadrivers interrelate are shown in (figure 1.2)

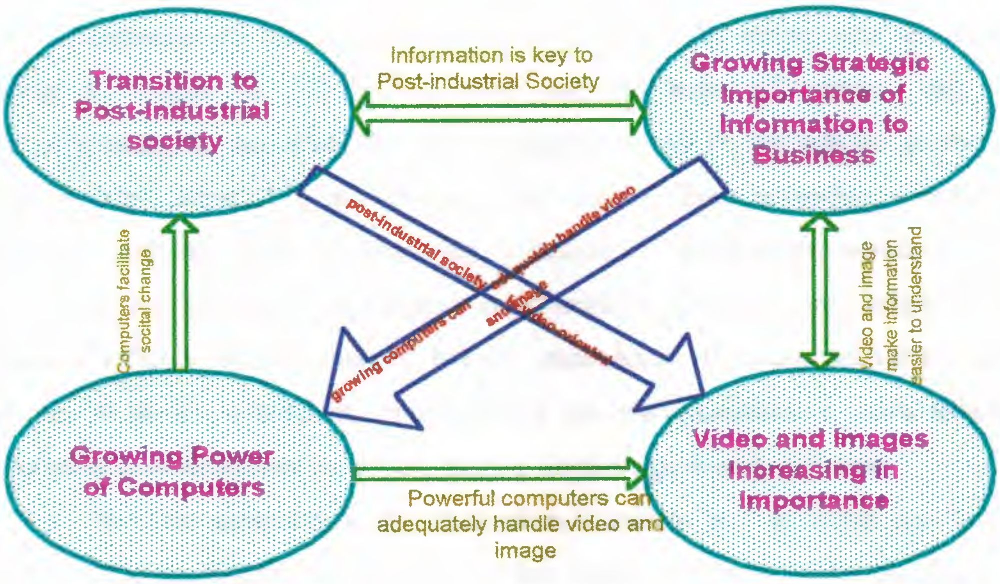

Figure 1.2 The Four Broadband Mega drivers

1.3.1 The Transition to an Information-Based Postindustrial Economy

Since the 1970s, writers such as Alvin Toffier, Daniel Bell, John Naisbitt, and others have pointed to an important transition through which our society is passing. Each of these writers has his own particular point of view, but they all share the position that our Western economies are shifting from being dominated by manufacturing to being dominated by services.

The historical model used by these writers is the Industrial Revolution, during which our economy shifted away from being based mostly on agriculture to being based mostly on manufacturing industry. This transformation turned out to have a profound impact on our entire society, changing the nature of work in very important ways. For example, with the coming of the Industrial Revolution, work life started to be organized around the time cycles associated with the factory rather than the agricultural cycle Most writers see the shift to a postindustrial service-oriented economy in a positive light. Others are not so sure. They worry that we are moving from a nation in which a substantial proportion of the population have high-paying factory jobs to one in which many of us will be working at low-paying service jobs someone working behind the counter at a McDonalds is the example usually cited.

Proponents of the postindustrial society counter by saying that this is not what is going to happen at all? Instead, they say, the types of services that will replace manufacturing, as the driver for the U.S. economy will be information services, defined in the broadest sense. We are moving, they believe, from becoming a nation of factory workers and managers to becoming a nation of information workers (or knowledge workers) and information managers. Such workers and managers include virtually every kind of office worker as well as lawyers, accountants, computer scientists, educators, and so on. With few exceptions, such workers are paid considerably more than the average worker at McDonalds and the content of their work is significantly different.

Both the advocates of the postindustrial society and its critics may turn out to be correct. The social critic Charles Murray has suggested that we are moving into a future in which we will become two nations, with a large class of knowledge workers at the top of the heap and an even larger class of low-skilled service workers at the bottom.

But this is not really the place to discuss the sociological dimensions of the coming information society. What is important here is the fact that information will become the stuff that makes up our economy. It will be bought and sold as fruits and vegetables have been bought and sold. Charles Murray has suggested that access to information resources will help the knowledge worker class dominate society. Murray sees this already happening with the growing class of knowledge workers making use of fax, while the proletariat uses the supremely inefficient U.S. Postal Service. There is certainly some truth in Murray's apocalyptic vision, but it should also be remembered that even McDonalds the archetype of the low-wage service job provider makes growing use of information technology McDonalds was, in fact, one of first companies to use ISDN in the United States.

However, an information economy needs an information infrastructure, and this means communications networks. Such networks will serve the information age much as roads and railroads served previous ages. Here is how John Naisbitt and Patricia Aburdene put it in their book Megat rends 2000. Associate in a Tokyo office from a mountain perch in Colorado, as if you were across a table. We are laying the foundations for an international information highway system. In telecommunications we are moving to a single worldwide information network. We are moving toward the capability to communicate anything to anyone, anywhere, by any form voice, data, text or image.

But what kind of network will support such activity? The passage just quoted contains some important clues. First, it suggests a network that carries a lot of information of different kinds. This means a network of high bandwidth, especially because, as we shall discuss in detail later in this chapter, the image communications specifically mentioned by Naisbitt and Aburdene can be particularly bandwidth hungry. The old public telephone and computer networks, which can typically deliver no more than 2 Mbps or so to the customer premises, are likely to be completely inadequate for the public network capacity requirements of the information age. As the result of the sheer bulk of information that will need to be transmitted across networks in the coming information age, networks will have to deliver tens, hundreds, or thousands of megabits per second to residential and business locations. Such networks will therefore clearly have the bandwidth characteristic of broadband networks. But the passage also suggests that the networks the coming information age will demand will have another characteristic of broadband networks they will be switched. Only a switched network, as we have defined it, could provide the ubiquity of communications that Naisbitt and Aburdene foresee.

Naisbitt and Aburdene do not have much to say about network intelligence, although it is unlikely that any network could do all the things that they say future networks will be able to do without some smarts. However, another well-known futurist Alvin Toffler has stressed intelligence as a key characteristic of future networks.

According to Toffler, in the past networks have tended merely to exhibit "intra-intelligence"; that is, the intelligence of the network was directed primarily at keeping the network up and running. In the future, Toffler believes that we will see the increasing dominance of "extra-intelligent" networks. In Toffler 's words, extra-intelligent networks will "analyze, combine, repackage or otherwise alter messages, sometimes creating new information along the way."

If Toffler, Naisbitt, and Aburdene are correct and I believe their forecasts will prove to be accurate the networks appropriate to the coming information age will be switched, will have high bandwidth, and will be highly intelligent. That is, they will be broadband networks.

1.3.2 The Growing Strategic Importance of Information to Business

As just outlined, the transition to an information age will mean more information sent to more locations. But it is not merely quantitative changes in networking that will be involved. As this transition proceeds, a qualitative change in the way that businesses view information will come about. As part of the fabric of the information age, businesses will succeed or fail based on how they deploy their networking infrastructure.

This trend is manifested by the gradual shift from a situation in which corporate telephone and data networks were viewed as cost centers to an environment in which the network becomes a means of directly enhancing profits. The classic case of this happening is Merrill Lynch's Cash Management Account (CMA), which, when it was launched in 1977, was the first financial product that offered customers a combined checking account, savings account, brokerage account, and credit card. Not only did customers get a single statement of account at the end of each month, but they could readily move funds from one type of account to another.

CMA effectively propelled Merrill Lynch into the banking industry and took traditional banks by surprise. In its first year of operation, CMA brought Merrill Lynch $5 billion in funds. By 1984 there were one million CMA account holders with total deposits of $70 billion. CMA was a major strategic coup. Response from the banks and from other brokerage houses either fell far short of Merrill Lynch's success with CMA or fizzled altogether. Although transparent to the customer, CMA was underpinned by Mertill Lynch's communications network and computer installations, without which such an innovation would have been impossible. CMA is one of the relatively few examples of a corporate communications infrastructure being used to launch an entirely new product. More often the strategic use of telecommunications is concerned with speeding up the pace at which business is conducted. Toffier cites two interesting examples of this trend:

- Shiseido, Japan's leading cosmetics firm, which sidesteps the normal distribution chain by connecting its network directly to retail stores. Shiseido can receive orders directly from retailers and broadcast new product information to retailers. The wholesaler is eliminated. Shiseido's profits are increased.
- American Hospital Supply (AHS), which places terminals inside hospitals so that

these hospitals can order supplies at the push of the button. Not only does this make it easier for hospitals to order from AHS than from its competitors, but it allows the hospitals to cut back on their own inventories, saving them significant amounts of money. It is not surprising that AHS's market share has increased.

Merrill Lynch's CMA and the distribution innovations of Shiseido and AHS are examples of major firms using electronic networking to change and enhance their relationships with a constituency to which their activities are addressed. Merrill Lynch and AHS were using networking to enhance their relationship with the financial services consumer. Shiseido was directing its marketing efforts toward the needs of its distributors. Other constituencies also may be addressed by what has become known as strategic telecommunications using telecommunications for competitive advantage. Telecommuting using networks to allow workers to stay at home and work from terminals is another example of strategic telecommunications, this time addressing the needs of the employee community.

There are many other examples of strategic telecommunications, and the topic is widely discussed, especially in journals, such as the Harvard Business Review, aimed at corporate strategists. This discussion, however, focuses mostly on how telecommunications can be deployed for competitive advantage. There is very little discussion of what kind of telecommunications, or what kind of networks, are (or will be) required to optimize this advantage.

This omission has largely been because there has not been much to say. Strategic telecommunications guru Peter Keen says of CMA, "The technology Merrill Lynch used was fairly standard. There was no state-of-the-art software or hardware." There are, however, some good reasons both technological and economic to suppose that in the future, strategic telecommunications will tend to demand advanced technology in general and broadband technology in particular.

Although Merrill Lynch found that none of its competitors could respond adequately to CMA, as we have already noted, other practitioners of strategic telecommunications are unlikely to be so lucky. Norman Weizer, a consultant with Arthur D. Little, has pointed out "onetime strategic moves simply do not confer sufficient long-term superiority. Rivals can quickly copy such systems. Thus, what once was strategic advantage typically weakens within 6 months into simple financial advantage." Weizer has a few suggestions of how to get round this problem. For one thing, Weizer claims that "To achieve success in the 1990s, a company needs a business and technology-wise staff, both inside and outside of the information systems department. The only sustainable competitive advantage in the 1990s is the ability of a company's people to identify and seize new opportunities more rapidly and effectively than competitors."15 Weizer clearly believes that advantages accrue to those companies that can recognize technological opportunities when they see them. In the future, many of those opportunities will stem from broadband networking.

In fact this is already beginning to happen. For example, at a recent presentation by Pacific Bell, the example was given of a California apparel maker who was using frame relay to give itself a one-week advantage over its competitor in product marketing. Prior to its use of frame relay the company was forced to send couriers to different showrooms with product samples. Using frame relay they were able to transmit very-high-quality color images of the samples and take orders for new product lines about a week before customers even got to see competitors' products.

However, as the first Weizer quote suggests, one can confidently predict that other apparel makers will soon copy any such innovations. Weizer, for example, cites the example of United Airlines' Covia reservation system soon being matched by American Airlines' SABRE. Weizer sees the solution in building a corporate information infrastructure that is as central to the functioning of the organization as any other major corporate asset. As Weizer puts it, "fundamental benefits derive from an infrastructure that forges all the of a company's information and communications systems into a coherent unit."

As it is built, such an infrastructure will call out for more and more bandwidth for a variety of reasons. For one thing, as information gains in strategic importance, more information will be carried on corporate networks. Second, as voice, data, image, and video communications are merged into a single communications fabric, more room will be needed for multimedia services that combine two or more of these modes of communications. Third, there will be growing need for bandwidth to connect up formerly separate networks. This trend, which is known usually as "enterprise networking," is exemplified by the need at many large corporations to connect up local computer networks over a region or across the nation (or indeed across the world) in order to make for a more efficient information flow. There is little point in automating an office with local area computer networks if the information is then going to sit in hard copy form waiting to be sent out by mail to another location.

The interconnection of local area networks (LANs) at different locations of large organizations is one of the major short-term factors driving broadband communications. Many of the early trials of broadband communications that are being carried out by the telephone companies and many of the pre-broadband and broadband services that are being brought to market by these companies at the present time are emphasizing LAN internetworking. Few of these services now operate at especially high data rates, although in some cases a clear migration path of higher data rates has been defined.

For now, however, note that the deployment and use of corporate networks implies not just more bandwidth, but also a need for switched bandwidth. Strategic networks must be flexible. They must be able to handle peaks and troughs in the demand for bandwidth, and they must be able to serve many locations, adding and disconnecting these locations from the corporate net as business requirements demand. This can only happen in a switched network and can only happen easily in an intelligent switched network.

Thus, although strategic telecommunications in the past has been based on creative new ways of deploying older networking technologies, competitive forces are likely to give telecommunications strategists more reasons to think in terms of new networking technology in the future. This technology will be high bandwidth, intelligent, and switched. It will, in other words, be broadband.

1.3.3 The Growing Power of Computers and Storage Devices

The third of the megadrivers pushing the need for broadband networking is the ever-growing power of computers and storage devices. This point will hardly need explanation to anyone who has watched the growth of the microcomputer industry over the past decade. Exactly how much more powerful today's machines are than the ones that preceded them somewhat depends on how one defines computer power. This is something of a contentious issue, because every vendor wants to be able to say that his machine is the most powerful.

One way of measuring computer performance is in terms of the number of operations of a particular kind that can be performed in a given amount of time. Thus a popular way of measuring the power of large computers is in terms of "megaflops" the number (in millions) of floating point operations that can be performed in a second. A floating point operation, somewhat simplified, is a calculation using real numbers. Using this measure, if we were to go back to the early 1970s, we would discover that a general-purpose computer had a power rating of one megaflop or under, although supercomputers would exceed 100 megaflops. Currently, a general-purpose computer would rate in the tens of megaflops, with some supercomputers moving towards 10 gigaflops (1 gigaflop = 1000 megaflops).

There have also been major improvements in the power of storage devices. Disk drives are not actually getting that much faster, because they are mechanical devices. However, they are becoming increasingly cost effective even floppy disks hold a lot more than they once did. In particular, optical storage, once a rather awkward technology, has matured. Even personal computer owners can now buy read-only optical storage, in the form of CD-ROM players, for a few hundred dollars, and erasable optical disks, once affordable only by the largest corporate users, have now come down to the several thousand-dollar mark. Again, this represents an increase in power of several orders of magnitude.

This is a huge jump in power in a relatively short period of time, and it has important consequences for the bandwidths required by computer networks. The most obvious consequence is that because computers are now capable of faster processing, they are also capable of throwing data onto networks much faster than before, so networks need extra bandwidth to support this increase. The combination of increased computer power and less-expensive memory also leads to a requirement for high-bandwidth pipes to link computer central processing units (CPUs) to large memory banks. The local networking standard mentioned earlier, FDDI, was originally developed precisely to serve this need, although computer rooms now represent only a small part of the FDDI market as a whole. HIPPI, another local networking standard that we will discuss in some depth later, was also originally intended as a computer room standard.

The increase in computing power that we have described so far is due to improvements in hardware and software technology. For example, some high-speed computing devices now use gallium arsenide (GaAs) rather than silicon for some parts of their circuitry. The physics of GaAs provides more electron mobility than silicon, and this translates into faster processing speeds. Meanwhile, the price/performance ratio of conventional silicon chip technology has continued to improve and can be expected to continue to do so over the next decade.

Another important development that has enhanced computer power has been the development and commercialization of “reduced instruction set computing” (RISC) microprocessors. RISC microprocessors are contrasted with the complex instruction set computing (CISC) microprocessors that are found in personal computers.CISC microprocessors were developed in the 1960s and 1970s when random access computer memory (RAM) was very expensive. In order to conserve RAM, each instruction to the microprocessor was relatively complex and was decoded into machine language by the permanent read-only memory (ROM) of the microprocessor. In RISC processors, instructions are short and decoding can be left to RAM, which is now far less expensive. Queues of instructions can be managed better, and more efficient use is made of the computer’s time, which makes for faster computing.

The computing power enhancements just described derive largely from developments in microelectronics. Further improvements of this kind are to be expected, but another entirely different kind of computer power enhancement is emerging, one that will put an even greater strain on bandwidth requirements of networks. This is the trend toward the “network itself becoming the computer,” a phrase that has become the slogan of distributed computing. In distributed computing, a single application runs on multiple processors, with the processors being connected over a network. Distributed computing networks can do more than a single computer alone. Thus distributed computing makes computers more powerful.

This distributed computing concept can be deployed in several ways. In massively parallel computers, processors connected across a network provide not just for more rapid computing but for the solution of “fuzzy” problems. This is because massively parallel architectures essentially mirror the structure of the brain.

Since the mid-1980s, “parallel” supercomputers have been constructed that substitute up to 1000 low-powered processors acting simultaneously for the single CPU.

The latest twist on the parallelism theme is the distribution of computer memory to each of the parallel processors so that the supercomputer in effect becomes a small network of computers, with relatively little in the way of a centralized architecture. "Massively parallel" super computers -supercomputers with more than just a few processors are the main growth area for supercomputing. The growth of "traditional" supercomputers is largely stagnant, at around 60 machines annually. The main constraint on the growth of parallel systems is that they are difficult to program. Breakthroughs in this area are likely, however.

Massive parallelism is mainly of interest in the supercomputer community. Although the installed base of supercomputers is likely to accelerate throughout the 1990s as the result of the development of the minisuper computer," of more importance to the world of general corporate computing is the trend toward client/server applications. Client/server is a form of distributed computing in which a considerable amount of processing occurs on a user workstation, but the application at the same time has access to a relational database on a server, with access to the server provided across a network. A considerable number of client/server software packages are already available, especially for financial applications.

The idea behind distributed computing is that certain tasks are off-loaded to one computer while another computer deals with other tasks, speeding up the computing process and, in some cases, allowing certain specialized types of computers to deal with tasks for which they are most suited. Distributed computing is significantly different from the distribution of computing intelligence that resulted initially from the introduction of the minicomputer and accelerated when IBM made the personal computer a respectable tool of corporate computing in 1981. With the older trend of distributed intelligence, different applications continued to reside on different computers but processing power became more localized. Mainframe hosts gave way to personal computers, for example. In the new distributed computing, several computers act together as if they were one computer.

As I have already noted, distributed computing by definition requires some kind of communications link or network running between multiple processors on which a shared application is being run. Hence the slogan, 'the network is becoming the computer." But if this is the case, then the network will have to behave in much the same way as the internal network of a regular computer. This may not sound too challenging, until one starts doing a few calculations.

Let us consider an engineering workstation operating at 30 million instructions per second (30 MIPS). These days this is nothing special. Let us suppose, also, that the word size, the size of the smallest possible instruction, is 32 bits. Any computing operation is going to consist of two words: one, the instruction itself, the other, the "operand," the piece of information on which the instruction acts. Putting this all together then, the internal bandwidth the data rate at which the network inside the computer must operate in order to bring information out of storage and into the central processing unit (CPU) to achieve 30 MIPS is calculated by multiplying the number of words that must be brought out of storage to perform an instruction by the number of instructions that must be performed in a second (30 million), and then multiplying again by the number of bits in each word (32). This works out to 1.92 Gbps. Again, this is way beyond today's commercial technology. As a matter of fact, except for point-to-point communications, it is way beyond any technology we would expect to see in the next five to ten years.

As workstations and personal computers get ever more powerful, the demands on bandwidth I have just described will become more and more acute. Similar applications for supercomputers will require bandwidths and network technologies that we have barely dreamed of. Put another way, our networking capabilities lag our computing capabilities by several orders of magnitude. For example, a computer operating at 1 billion instructions per second (1 BIPS) and with a 64-bit word size has an internal bandwidth of 128 Gbps. No existing transmission technology even in the best research labs could support this kind of bandwidth over any distance.

Today's networks are simply not up to the bandwidth challenge presented by all this increase in the power of computing and storage devices. And this increase in power is not just demanding of bandwidth. It also requires enhanced network intelligence and switched capabilities so that powerful computing and storage devices can find each other across a net- Work for distributed computing applications or merely to exchange large files. In other words, the logic of increased computing power leads us to construct broadband networks, but more powerful computing is not just a driver for broadband networking. It is also an enabling technology. Enhanced computer power has been incorporated into customer premises equipment such as routers, hubs, and switches, and has enabled important network functions, especially error checking and flow control, to be devolved to the network termination and eliminated from network nodes. This development is at the core of fast packet switching, the key transport technology for broadband networks.

A second way in which enhanced computing power is a broadband enabler is its ability to add intelligence to the network, especially in the form of advanced signalling systems and network management systems. An advanced signalling system is required to provide the full range of services of which broadband networks are capable. In particular, it is required to support bandwidth-on-demand services, which are closely identified with broadband technology. The future broadband signalling system will presumably be a superset of Signalling System Number 7, the system currently being developed for ISDN. Sophisticated network management systems will also be required for broadband networks in order to deal with the multiservice aspect of these types of networks. However, network management and signalling remain missing pieces of the broadband jigsaw puzzle at the present time.

1.4 Today's Broadband Environment

We have now taken a brief look at the four broadband megadrivers. These megadrivers are what will move broadband out of the R&D lab and telephone company trial phases and turn it into a commercialized reality. To summarize what went before: We will need broadband networks to handle the information glut that is coming to us in the future information society. We will need broadband networks to handle the strategic information needs of business in the late 20th century. We will need broadband networks to connect up the next generation of computing devices and peripherals. And finally, we will need broadband networks to provide high-quality digital video and image communications.
But broadband lies in the future. So where are we today in terms of bandwidth, switching technology, and intelligence?

1.4.1 Bandwidth Today and Tomorrow

Let us begin, as before, with bandwidth. We need to consider two measures here. One is the data rate at which terminals, personal computers, workstations, and similar devices access networks. The other is the speed of the network itself.

Network access rates begin at 9.6 kbps. If you walk into a computer store and want to buy a modem for your personal computer, you will probably end up with a modem operating at this speed. A decade ago the common standard for personal computer modems was a mere 0.3 kbps and we are now at a point where modems operating at 14.4 kbps are widely available.

Modem rates define the lower end of data rates for access to wide-area networks (WANs) that is, networks, such as the public telephone network, that extend well beyond the customer's premises. Many terminals now access WANs over digital channels operating at 9.6 kbps and upwards. ISDN promises higher data rates for network access. Access using ISDN begins at 16 kbps using one D channel, and it is possible to use 1SDN's Primary Rate Interface to hook into the network at up to 1.544 Mbps. That is currently about the limit for most commercial terminals, although there are now some terminal interfaces operating at 100 Mbps and 150 Mbps.

Many terminals, personal computers, and multi-user hosts do not access wide-area networks directly, but rather do so through a LAN. Access speeds to LANs are much higher than those just cited for access to WANs. For example, the lowest-speed Ethernet LANs require interfaces operating at 1 Mbps. At the other end of the scale LAN interfaces for the new Fiber Distributed Data Interface (FDDI) LAN function at 100 Mbps.

Two new and related standards - the Higher Performance Parallel Interface (HIPPI) and Fibre Channel define local networking at speeds up to 1.6 Gbps. HIPPI and Fibre Channel are today primarily of interest to the supercomputing community and users of large mainframe computers, but this may change. In particular, IBM, Hewlett-Packard, and Sun the three largest workstation vendors have recently decided to promote Fibre Channel as a technology for high-speed workstation interfaces. Table 1-4 compares the established local networking standards.

LAN interfaces are at least a couple of orders of magnitude faster than WAN interfaces. However, there is a catch. LANs are shared bandwidth solutions. LAN interfaces provide full access for a given terminal to the full bandwidth of the LAN-1 Mbps in the case of a low-end Ethernet or 100 Mbps in the case of FDDI but only for a short period of time. In part, different kinds of LANs are distinguished by the manner in which they grant access to a network's total bandwidth.

LANs therefore operate at the same bandwidths as the access interfaces for the workstations and personal computers attached to them. WANs on the other hand may have segments operating at rates entirely

TABLE 1.4 Local area-networking standards

	Standards	Speeds	Media	Vender support	Key applications
Ethernet	IEEE 802.3	1Mbps, 10Mbps	Twisted pair , coax,fiber	Ubiquitous	Work groups
Token-Ring	IEEE 8020.5	4Mbps, 16Mbps	Twisted-pair, fiber	Widespread, Led by IBM	Workgroups and small backbones
FDDI	ANSI X3T9.5	100 Mbps	Twisted-pair, fiber	Widespread	Workgroups and small backbones, and metropolitan area networks
HIPPI	ANSI X3T9.3	800Mbps, 1.6Gbps	Twisted-pair	Limited	Computer room networks
Fibre Channel	ANSI X3T9.3	100Mbps, 200Mbps, 400Mbps, 800Mbps	Twisted pair, coax,fiber	Limited	Computer room networks and workgroups

Different from those at which the terminal devices attached to them gain access. Most corporate networks today are made up of communications channels operating at Ti rates-1.544 Mbps. A few include segments operating at T3 rates-45 Mbps. There are hardly any terminals or workstations attached directly to WANs at 1.544 Mbps and virtually none attached at 45 Mbps.

Ti and T3 are old Bell System standards, reflecting the rates at which the old AT&T ran its digital trunks. Today's telephone companies are still using T3 trunking but increasingly need rates above T3 for busy routes between major cities or in major downtown areas. An entirely new digital hierarchy called the Synchronous Optical

Network (SONET) has been defined to meet this need. The SONET hierarchy begins at approximately 50 Mbps and extends to approximately 10 Gbps, with extensions beyond that point relatively easy to define. Because of difficulties with defining and implementing network management in SONET networks, there is still very little SONET equipment installed in today's networks. However, there is a considerable amount of nonstandard transmission equipment already installed in the public networks, operating at rates similar to those defined in SONET.

1.4.2 Switching: The Evolution to ATM

Today's dominant switching technology is still circuit switching, which is used throughout the public telephone network and is the switching "philosophy" that underpins the millions of PBXs on customer premises throughout the world. There are, however, various generations of circuit switches. The first generation was entirely manually operated. The second generation used electromechanical relays to perform some of the switching functions that needed little intelligence to perform. The third generation had greatly increased intelligence because it moved most of the control of the switch over to a computer. This type of switch was called a stored program control (SPC) switch.

The fourth generation of switch was fully digital, with the information being switched itself in digital form. Previous generations of switch dealt with information in analog form. With the fourth generation, the switch effectively became a rather specialized form of computer.

Today, most of the public network is switched by switches that are either third or fourth generation. PBXs are also mostly of these two generations. There are a few switches around especially on customer premises from earlier generations, but their number is declining.

These circuit switches were originally intended for voice communications but increasingly are asked to handle data and video communications, too. This may happen in either a passive or active manner. Passively, the public network finds itself the transport medium for huge volumes of data transmitted from terminals equipped with modems. Actively, public carriers are offering a growing number of circuit-switched services aimed at the data community and for videoconferencing services.

Although these new circuit-switched services are reengineered some-what to meet the needs of data and video, the truth is that circuit switching was invented and optimized for voice. Packet switching in which data is packaged into packets was invented to handle data communications requirements, including the long bursty sessions that tend to be associated with computer-to-computer communications. A computer can handle many more connections on a packet switched fabric than it can on a circuit-switched fabric.

Packet switching is the fabric chosen by most networks, which are mostly dedicated to data. These include private networks and a few specialized public networks. Although voice and video can be carried on packet networks, there are some inherent technical problems with doing so, and to date, pocketsized voice and video has mainly been experimental. Packet-switched service providers, which are also known as public data networks or value-added networks (VANs), have tended to be institutionally separate from voice-oriented earners.

However, this is beginning to change. The two biggest public packet switchers are now owned by companies whose business is mainly on the voice side. Sprint bought the old Telenet Company and renamed it Sprint-Net. British Telecom, the dominant carrier in the United Kingdom, bought Tymnet and renamed it BT Tymnet and then BT North America. However, most of the other packet-switched carriers have nothing to do with voice. For example, the third biggest packet switcher is CompuServe, which is owned by H&R Block, the tax prepares. And in addition to commercial services, there is a huge sprawling network of networks, loosely called the Internet, serving mainly the educational and research communities. All of the networks that make up-the Internet are packet-switched data networks.

But whether commercial carrier or educational network, these packet-switched networks all offer similar services raw bandwidth for interconnecting computers, electronic mail (including an E-mail directory), access to databases, and electronic data interchange (EDI, the electronic exchange of standard forms and business documents). There is still hardly a hint of packet-switched networks trying to offer anything along the lines of a packet-switched real-time voice or video service.

However, as we have already noted, the broadband networks of the future will be based on a special version of packet switching called fast packet switching. We will discuss the concepts behind this is the next chapter. In this chapter, we will merely note the current state of fast packet and note a rather interesting fact about how it is being deployed.

Today, fast packet comes in two versions. One is a relatively primitive version that is still only suitable for data transmission. This version is called frame relay. Frame relay is now commercially available and in late 1993 was being used by perhaps 500 end user companies all large corporations in the United States. The other version of packet switching is cell relay. Cell relay should eventually be able to be used with equal ease to supply voice, video, and data services at least in theory. To date, however, cell relay products and services have mostly been data oriented.

B-ISDN will use a specific kind of cell relay, called asynchronous transfer mode (ATM), to provide a vast range of services, from the most mundane voice or low-speed data service to the most exotic multimedia service.

ATM is currently being hailed as the network panacea the switching/multiplexing technology that can do it all. Its influence is being felt both in the world of public network switching and in the LAN community, where it is said to pose a threat to FDDI the 100Mbps LAN. What is interesting, though, is that rather than replacing existing switching fabrics, ATM actually looks as though it will supplement existing switching technologies.

This pattern is already emerging. For example, here is how the Bell companies' Switched Multimegabit Data Service (SMDS) works: A customer buys a Ti access line, which takes traffic to the local central office. The carrier then backhauls the SMDS traffic to a specialized switch (which will increasingly be an ATM switch) shared by a central office. AT&T's evolutionary philosophy of introducing ATM into the public carrier switching products also has the ATM switch as an adjunct processor. Siemens is already most of the way down the same strategic road. Its vision of the medium-term future is one in which conventional traffic is handled by a Siemens Stromberg-Carlson circuit switch while the Siemens MAN cell switch does the honors with the broadband traffic.

2. PRINCIPLES AND BUILDING BLOCKS OF B-ISDN

2.1 B-ISDN Principles

The motivation to incorporate broadband & features into ISDN is neatly documented in CCITT Recommendation.

The B-ISDN recommendations were written taking into account the following:

- The emerging demand for broadband services.
- The availability of high speed transmission switching and signal processing technologies (bit rates of hundreds of Mbit/s are being offered)
- The improved data and image Processing capabilities available to the User.
- The advances in software application processing in computer and telecommunication industries.
- The need to integrate interactive and distribution services and circuit and packet transfer modes into one universal broadband network. In comparison to several dedicated networks, service and network integration has major advantages in economic planning, development implementation Operation and maintenance. While dedicated networks require several distinct and costly customer access lines, the B-ISDN access can be based on a single optical fibre for each customer. The large-scale production of highly integrated system components of a unique B-ISDN will lead to coat effective solutions.
- The need to provide flexibility in satisfying the requirements of both user and operator (in terms of bit rate, quality of service etc.).

ISDN is conceived to support a wide-range of audio, video and data applications in the same network' [47]. B-ISDN thus follows the same Principle5 M 64 kbit/s based ISDN (cL CCITT Recommendation 1.120 [46]) and is a natural extension of the latter [47]:

A key element of service integration is the provision of a wide range of services to a broad variety of users utilizing a limited set of connection types and multipurpose user-network interfaces.

Whereas telecommunication networks of the pre-ISDN era have usually been specialized networks (e.g. for telephony or data) with rather limited bandwidth or throughput and processing capabilities, the future B-ISDN is conceived to become a universal (standardized) network supporting different kinds of applications and customer categories. CCITT Recommendation 1.121 [47] presents an overview of B-ISDN capabilities:

B-ISDN supports switched, semi-permanent and permanent, point-to-point and point-to-multipoint connections and provides on demand, reserved and permanent services. Connections in B-ISDN support both circuit mode and packet mode services of a mono- and/or multi-media type and of a connectionless or connection-oriented nature and in a bidirectional or unidirectional configuration.

A B-ISDN will contain intelligent capabilities for the purpose of providing advanced service characteristics, supporting powerful operation and maintenance tools, network control and management.

We believe the reader of this list of intended B-ISDN capabilities must be deeply impressed; B-ISDN is tailored to become the universal future network!

B-ISDN implementations will, according to the CCITT, be based on the asynchronous transfer mode (ATM).

2.2 Asynchronous Transfer Mode

The asynchronous transfer mode (ATM) is considered the ground on which B-ISDN is to be built [47]: asynchronous transfer mode (ATM) is the transfer mode for implementing B-ISDN

The term transfer comprises both transmission and switching aspects, so a transfer mode is a specific way of transmitting and switching information ma network.

In ATM, all information to be transferred is packed into fixed-size slots called cells.

These cells have a 48-octet information field and a 5-octet header. Whereas the information field is available for the user, the header field carries information that pertains to the ATM layer functionality itself, mainly the identification of cells by means of a label (see Figure 2.1).

figure 2.1 : ATM Cell Structure

ATM allows the definition and recognition of individual communications by virtue of the label field inside each ATM cell header; in this respect, ATM resembles conventional packet transfer modes. Like packet switching techniques, ATM can provide a communication with a bit rate that is individually tailored to the actual need, including time-variant bit rates.

The term asynchronous in the name of the new transfer mode refers to the fact that, in the context of multiplexed transmission, cells allocated to the same connection may exhibit an irregular recurrence pattern as cells are filled according to the actual demand. This is shown in (Figure 2 (b)).

In the synchronous transfer mode (STM) (see Figure 2.2(a)), a data unit associated with a given channel is identified by its position in the transmission frame, while in ATM (Figure 2.2(b)) a data unit or cell associated with a specific 'virtual channel may occur at essentially any position. The flexibility of bit rate allocation to a connection in STM is restricted due to predefined channel bit rates and the rigid structure of conventional transmission frames. These normally will not permit individual structuring of the payload or will only permit a quite limited selection of channel mixes at the corresponding interface at subscription time. Otherwise the network provider would have to manage a host of different interface types, a situation that the designer would try to avoid, for obvious reasons (for example, STM switching of varying B and H channel mixes per interface requires switching equipment that can simultaneously handle all sorts of channels potentially used by customers at any time).

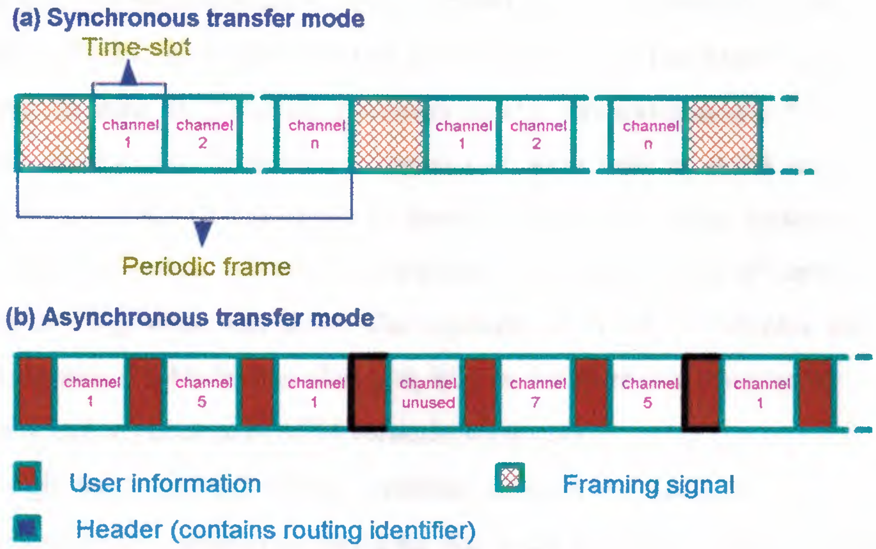

Figure 2.2 STM and ATM principles

The sum up, whereas today's In ATM-based networks the multiplexing and switching of cells is independent of the actual application. So the same piece of equipment in principle can handle a low bit rate connection as well as a high bit rate connection, be it of stream or burst nature. Dynamic bandwidth allocation on demand with a fine degree of granularity is provided. So the definition of high-speed channel bit rates is now, in contrast to the situation in a STM environment, a second-rank task.

The flexibility of the ATM-based B-ISDN network access due to the cell transport concept strongly supports the idea of a unique interface, which can be employed by a variety of customers with quite different service needs.

However, the ATM concept requires many new problems to be solved. For example, the impact of possible cell loss, cell transmission delay and cell delay variation on service quality needs to be determined.

Networks are characterized by the coexistence of circuit switching and packet switching; B-ISDN will rely on a single, new method called ATM, which combines adva7ntageous features of both the circuit- and packet-oriented techniques. The former requires only low overhead and processing, and, once a circuit-switched connection is established, the transfer delay of the information being carried over it is low. The latter is much more flexible in terms of bit rate assigned to individual (virtual) connections. ATM is a circuit-oriented, hardware-controlled, low overhead concept of virtual channels which (by contrast with X.25 access no flow control or error recovery. The implementation of these virtual channels is done by fixed-size (relatively short) cells and provides the basis for both switching and multiplexed transmission. The use of short cells in ATM and high transfer rates involved (e.g. 150 Mbit/s) result in transfer delays and delay variations which are sufficiently small to enable universal applicability to a wide range of services including real-time services, e.g. voice and video. The capability of ATM to multiplex and switch on the cell level supports flexible bit rate allocation, as is known from packet networks.

The overall protocol architecture of ATM networks comprises:

- A single link-by-link cell transfer capability common to all services
- Service-specific adaptation functions for mapping higher layer information into ATM cells on an end-to-end basis, e.g. packetization/depacketization of continuous bit streams into/from ATM cells or segmentation/reassembly of larger blocks of user information into/from ATM cells (core-and-edge concept).

Another important feature of ATM networks is the possibility of grouping several virtual channels into one so-called virtual path. The impact of this technique on the B-ISDN structure will be addressed in the following chapter.

2.3 Optical Transmission

The development of powerful and economic optical transmission equipment was the other big driving force for B-ISDN. Optical transmission is characterized by:

- Low fibre attenuation (allowing for large repeater distances)
- High transmission bandwidths (up to several hundred Mbit/s)
- Comparably small diameter (low weight/volume)
- High mechanical flexibility of the fibre
- Resistance against electromagnetic fields
- Low transmission error probability
- No cross-talk between fibres

- Tapping much more difficult.

The high bandwidth of optical transmission systems- currently up to Gbit/s can be transported via one optical link - has led to early implementations in public networks to support existing services like telephony. Fibre-based local area networks are also widely in use nowadays, providing a bit rate in the order of magnitude of a hundred Mbit/s to the users.

So for B-ISDN the use of optical fibre-based transmission systems is straightforward from a technical viewpoint, at least in the trunk network and in the local access part of the network where considerable distances have to be bridged.

In B-ISDN at least about 150 Mbit/s will be offered to the user across the broadband user-network interface. Though much higher bit rates could comfortably be transmitted on optical fibre links, the costs of the electronics in-volved in the transmission equipment (e.g. sender/receiver in network terminations, terminals etc.) together with considerations on expected service needs i.e. bit rates simultaneously required at the interface led to the conclusion that a B-ISDN 'basic' interface at about 150 Mbit/s would be sufficient and adequate in many cases.

In addition, a second interface type with at least 600 Mbit/s in the direction from the network to the user is also foreseen. Handling of 600 Mbit/s ATM signals is still a challenge, the economic implementation of which is currently not so assured.

The deployment of highly reliable optical transmission systems with rather low bit error probabilities benefits a simplified network concept with, for example, potentially reducible data link layer functionality.

3. B-ISDN NETWORK CONCEPT

3.1 General Architecture of the B-ISDN

The architectural model of the B-ISDN is described in CCITT Recommendation 1.327 [53]. According to this recommendation, the information transfer and signalling capabilities of the B-ISDN comprise:

- broadband capabilities
- 64 kbit/s based ISDN capabilities
- user-to-network signalling
- inter-exchange signalling
- user-to-user signalling.

This is depicted in Figure 3.1:

Broadband information transfer is provided by ATM. The ATM data unit is the cell, a fixed-size block of 53 octets. The 5-octet cell header carries the necessary information to identify cells belonging to the same virtual channel. Cells are assigned on demand, depending on the source activity and the available resources.

ATM guarantees (under normal, i.e. fault-free, conditions) cell sequence integrity. This means that a cell belonging to a specific virtual channel connection can nowhere in the network overtake another cell of the same virtual channel connection that has been sent out earlier.

ATM is a connection-orienied technique. A connection within the ATM layer consists of one or more links, each of which is assigned an identifier. These identifiers remain unchanged for the duration of the connection.

Signalling information for a given connection is conveyed using a separate identifier (out of band signalling)

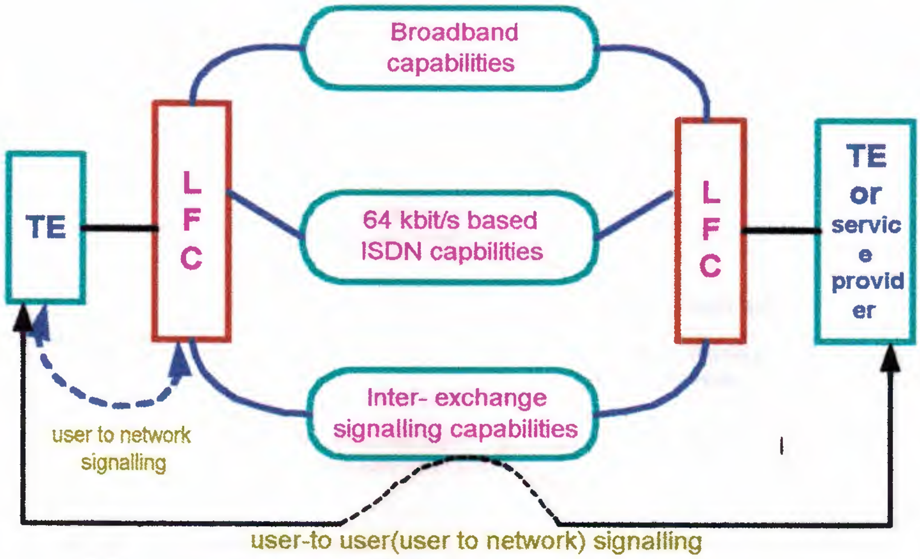

LFC Local function capabilities

TE Terminal equipment

Figure 3.1 Information transfer and signalling capabilities

Though ATM is a connection-oriented technique, it offers a flexible transfer capability common to all services including connectionless data services.

3.2 Networking Techniques

3.2.1 Network Layering

CCITT Recommendation 1.311 [50] presents the layered structure of the B-ISDN depicted in (Figure 3.2)

In this section, we only address the ATM transport network whose functions are split into two parts, namely physical layer transport functions and ATM layer transport functions.

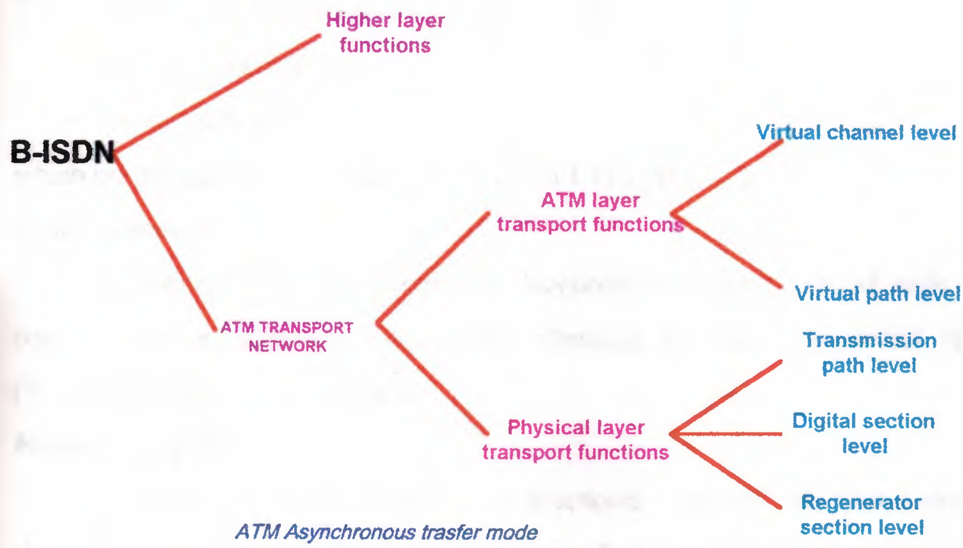

Figure 3.2: B-ISDN layered structure

Both the physical layer and the ATM layer are hierarchically structured. The physical layer
consists of:

- Transmission path level
- Digital section level
- Regenerator section level

Which are defined in the following way:

Transmission path:

The transmission path extends between network elements that assemble and disassemble the payload of a transmission system (the payload will be used to carry User information; together with the necessary transmission overhead it forms the complete signal).

Digital section:

The digital section extends between network elements, which assemble and disassemble continues bit or byte streams.

Regenerator section:

The regenerator section is a portion of a digital section extending between two adjacent regenerators.

The ATM layer has two hierarchical levels, namely:

- virtual channel level
- virtual path level

which are defined in CCITT Recommendation 1.113 [451:

Virtual channel (VC):

A concept Used to describe unidirectional transport of ATM cells associated by a common unique identifier value.' This identifier is called the virtual channel identifier (VCI) and is part of the cell header.

Virtual path (VP):

A concept used to describe unidirectional transport of cells belonging to virtual channels that are associated by a common identifier value.' This identifier is called the virtual path identifier (VPI) and is also part -of the cell header.

(Figure 3.3): demonstrates the relationship between virtual channel, virtual path and transmission path: a transmission path may comprise several virtual paths and each virtual path may carry several virtual channels. The virtual path concept allows grouping of several virtual channels.

Figure 3.3: Relationship between virtual channel, virtual path and transmission path

Concerning the levels of the ATM layer (virtual channel and virtual path), it proved helpful to distinguish between links and connections [45]:

- *Virtual channel link:* A means of unidirectional transport of ATM cells between a point where a VCI value is assigned and the point where that value is translated or removed.'
- *Similarly, a virtual path link:* is terminated by the points where a VPI value is assigned and translated or removed.

A concatenation of VC links is called a virtual channel connection (VCC), and likewise, a concatenation of VP links is called a virtual path connection (VPC). The relationship between different levels of the ATM transport network is demonstrated in (Figure 3.4)

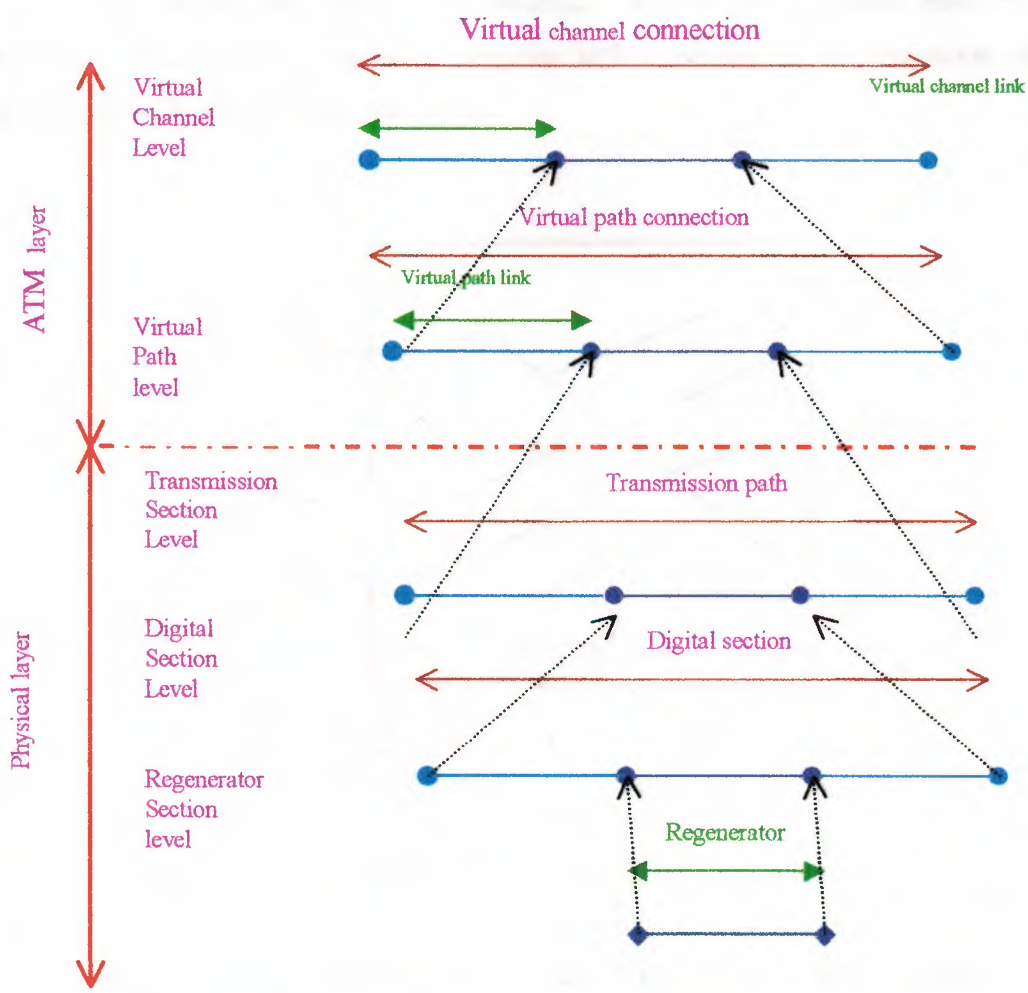

ATM Asynchronous transfer mode

Endpoint of the corresponding levels

Connecting point of the corresponding levels

Figure 3.4 Hierarchical layer-to-layer relationship

A VCC may consist of several concatenated VC links, each of which is embedded in a VPC. The VPCs usually consist of several concatenated VP links. Each VP link is implemented on a transmission path, which hierarchically comprises digital sections and regenerator sections.

3.2.2 Switching of Virtual Channels and Virtual Paths

VCIs and VHs in general only have significance for one link. In a VCC/VPC the VCI/VPI value will be translated at VC/VP switching entities.

VP switches (see Figure 3.5) terminate VP links and therefore have to translate incoming VPIs to the corresponding outgoing VPIs according to the destination of the VP connection. VCI values remain unchanged.

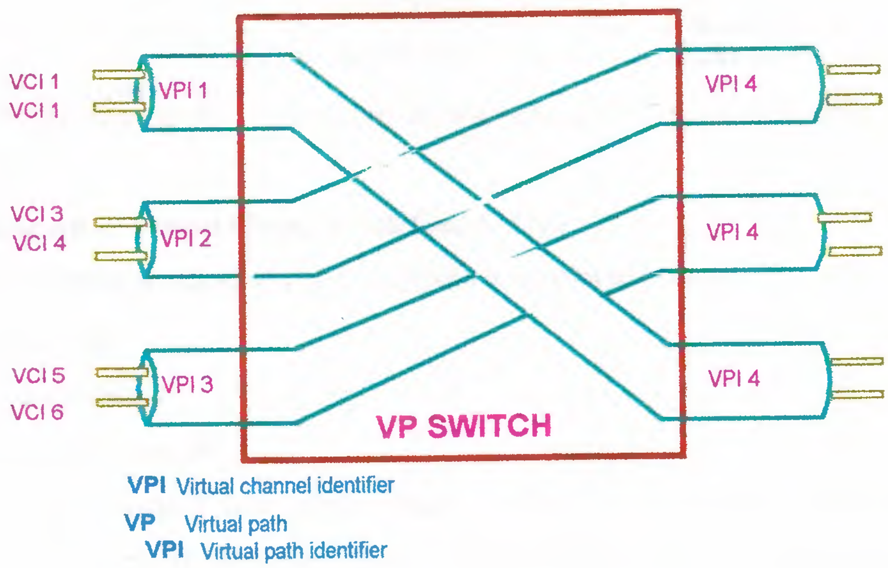

Figure 3.5: Virtual path switching

VC switches (see Figure 3.6) terminate both VC links and necessarily VP links. VPI and VCI translation is performed. As VC switching implies VP switching, in principle a VC switch can also handle mere VP switching.

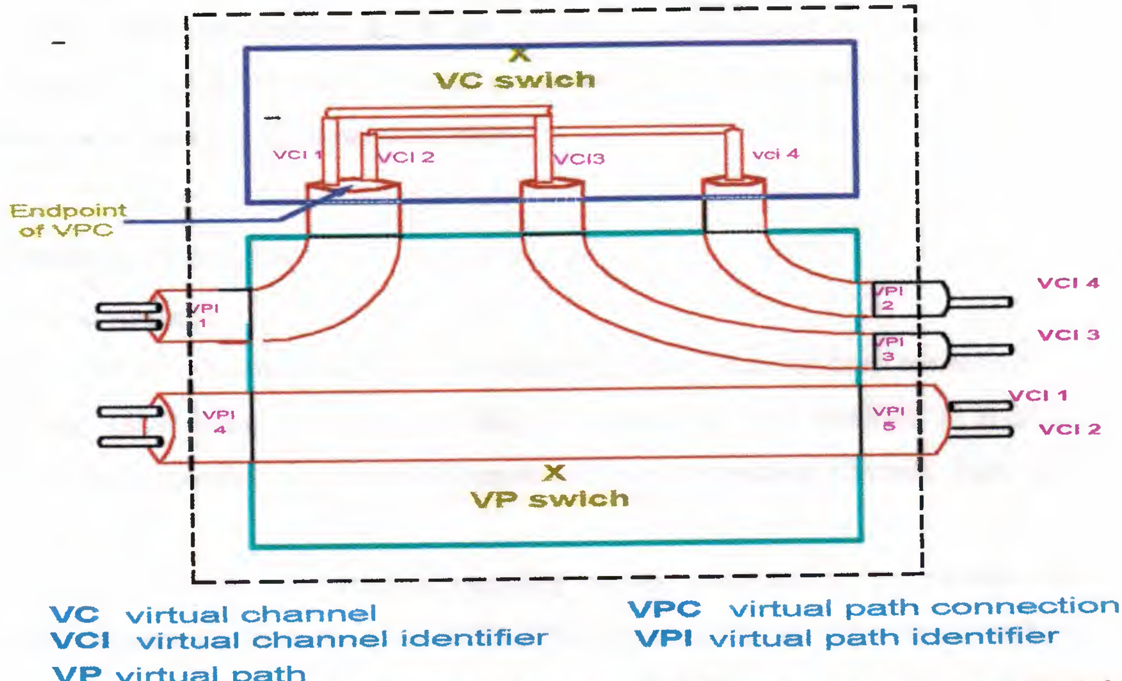

Figure 3.6 Virtual channel / Virtual path switching

3.2.3 Applications of Virtual Channel/Path Connections

Virtual channel/path connections (VCCs/VPC5) can be employed between:

- User and user
- User and network
- Network and network.

All cells associated with an individual VCC/VPC are transported along the same route through the network. Cell sequence is preserved (first sent - first received) for all VCCs.

User-to-user VCCs are able to carry user data and signalling information, user-to-network VCCs may for instance be used to access local connection related functions (user-network signalling), and network-to-network VCC applications include network traffic management and routing.

A VPC between users provides them with a transmission 'pipe' the VC organization of which is up to them. This concept may, for example, be applied to LAN-LAN coupling.

A user -to-network VPC can be used to aggregate traffic from a customer to a network element such as a local exchange or a specific server.

Finally, network-to-network VPCs can be used to organize user traffic according to a predefined routing scheme or to define a common path for the exchange of routing information or network management information.

3.3 Signalling Principles

3.3.1 General Aspects:

B-ISDN follows the principle of out-of-band signalling that has been established for the 64 kbit/s ISDN where a physical signalling D channel has been specified. In B-ISDN the VC concept provides the means to separate logically signalling channels from user channels.

A user may now have multiple signalling entities connected to the network call control management via separate ATM VCCs. The actual number of signalling connections and the bit rate allocated to them can be chosen in B-ISDN in a way That satisfies a customer's need optimally.

3.3.2 Capabilities Required for B-ISDN Signalling

B-ISDN signalling must be able to support:

- 64 kbit/s ISDN applications
- New broadband services.

This implies that existing signalling functions according to CCITT Recommendation Q.931 [74] must be included in B-ISDN signalling capabilities; on the other hand, the nature of. B-ISDN the ATM transport network and the increasing desire for advanced communication forms like multi-media services require specific new signalling elements. In the following, an overview of necessary B-ISDN signalling capabilities is given.

ATM network-specific signalling capabilities have to be realized in order to:

- Establish, maintain and release ATM VCCs and VP Cs for information transfer
- Negotiate (and perhaps renegotiate) the traffic characteristics of a connection.

Other signalling requirements are basically not ATM related but reflect the fact that more powerful service concepts appear. Examples are the support of multi-connection calls and multi-party calls.

For a multi-connection call, several connections have to be established to build up a 'composite' call comprising, for example, voice, image and data. It must also be possible to remove one or more connections of a call or to add new connections to existing ones. A certain capability in the network to correlate the connections to a s required. In any case, release of a call as a whole must be possible.

These correlation functions should be performed in the origination and destination B-ISDN switch only, since the transit nodes should not be burdened with such tasks.

A multi-party call consists of several connections between more than two endpoints conferencing). Signalling to indicate establishment/release of a multi-party call and adding/removing one party is required. (A communication that is part of multi-party call may be of multi-connection nature itself.)

In a broadband environment, asymmetric connections (i.e. low or zero bandwidth one direction and high bandwidth in the other) will gain relevance; signalling elements support such connection types have to be established.

Another broadband issue impacting signalling is interworking, e.g. B-ISDN with non B-ISDN services, or between video services with different coding schemes.

3.3.3 Signalling Virtual Channels

In B-ISDN, signalling messages will be conveyed out-of-band in dedicated signalling virtual channels (SVCs). Different types of SVC are provided [50]; they shown in (Table 3.1).

There is one meta-signalling virtual channel (MSVC) per interface. This channel is bi-directional and permanent. It is a sort of interface management channel used to establish, check and release the point-to-point and selective broadcast SVCs.

Table 3.1 Signalling virtual channels at B-ISDN UNI

SVC Type	Directionality	Number of SVCs
Meta-signalling channel	Bidirectional	1
General broadband SVC	Unidirectional	1
Selective broadband SVC	Unidirectional	Several possible
Point-to-point SVC	Bidirectional	One Per Signalling endpoint

Whereas the meta-signalling virtual channel is permanent, a point-to-point signalling channel is allocated to a signalling endpoint only while it is active.

A signalling end point at the user side may be located in a terminal or in the B-NT2 (e.g. private branch exchange). In a multi-functional terminal, multiple signalling endpoints may occur.

The point-to-point signalling channels are □i-directional. They are used to establish control and release VCCs or VPCs to carry user data (VPCs as well as VCCs, may also be established without using signalling procedures, e.g. by subscription).

The broadcast SVCs are unidirectional (network-to-user direction only). They are used to send signalling messages either to all signalling endpoints in a customer's network or to a selected category of signalling endpoints. The general broadcast SVC reaches all signalling endpoints; it is implemented in any case. Selective broadcast SVCs may be provided in addition as a-network option td- be able to address all terminals belonging to the same service profile category (a B-ISDN service profile contains information which is maintained by the network to characterize the services offered by the network to the user; the service profile may be specified as in CCITT Recommendation Q.932, Annex A [75] or otherwise).

The provision of SVCs in the network is Currently under discussion; so far no firm decisions have been taken but the principles outlined here are expected to apply:

To illustrate the SVC concept of B-ISDN, an example (based on CCITT Recommendation I. 311) is given in (Figure 3.7), which highlights different possibilities for carrying signalling information from the customer to the network and vice versa.

Four different VP links/connections are depicted in the figure. The first (a) is a signalling VP link which transports all the signalling information to be exchanged between the customer and the local exchange, including. Meta-signalling. When a signalling capability to a point in the network other than the local exchange is required (e.g. in order to communicate with .a special service provider located elsewhere), such signalling can be done on an extra VPC. (c) Which may carry signalling and user data. This VPC goes through the local exchange and is terminated at the appropriate place. (The other two VPs (b) and (d) shown for completeness in the figure carry user data only. In case (b) the corresponding VCs are switched in the local exchange and in case (d) the VP as a whole goes transparently through the local exchange.)

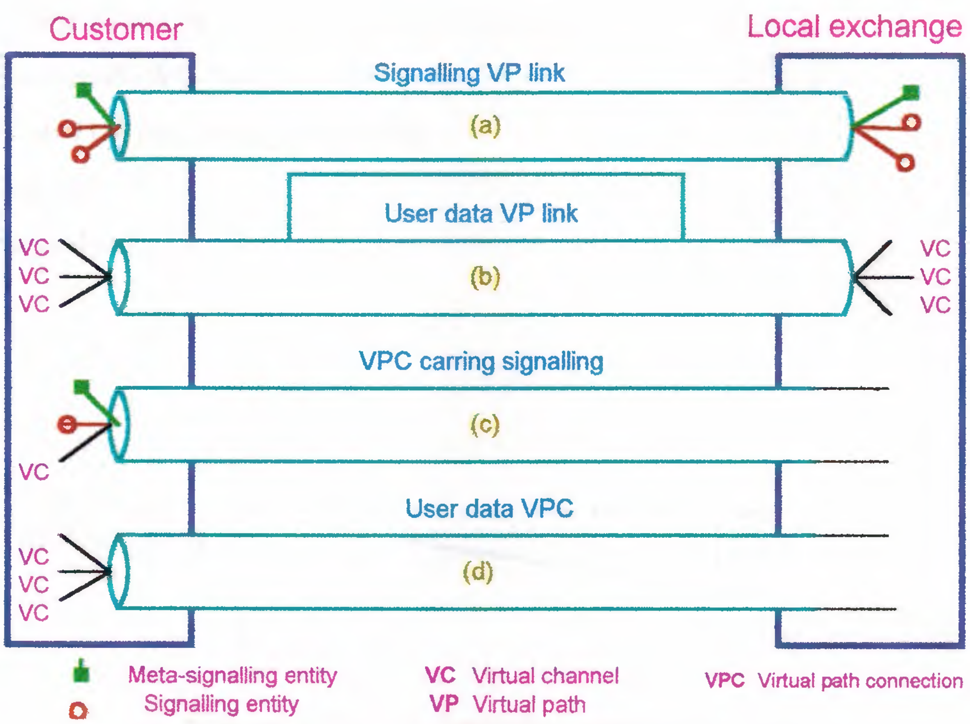

Figure 3.7: allocation of signalling virtual channels to a customer

3.4 Broadband Network Performance

Broadband networks based on ATM cell customer must meet certain performance in requirement in order to be accepted the users and network providers transport network performance but also by higher layer mechanisms. In some cases, these will be able to compensate for shortcomings in the ATM transport network performance. Cells belonging to a specified virtual connection are delivered from one point in the network to another, e.g. from A to B. A and B may denote the very endpoints of a virtual connection,. Or they may delimit a certain portion of the cell transport route, e.g. A and B may indicate national network boundaries of an international ATM connection. Due to a certain transfer delay, cells sent from A arrive at B within $\Delta t > 0$ (see Figure 3.8).

Note that the cell exit event occurs according to CCITT Recommendation L35B 44] when the first bit of the ATM cell has completed transmission across **A,** and the cell entry event occurs when the last bit of the ATM cell has completed transmission across B.

In order to adequately describe the quality of ATM cell transfer, CCITT Recommendation L35B [44] first defines the following outcome categories:

- Successfully delivered cell
- Errored cell and severely errored cell
- Lost cell
- Inserted cell.

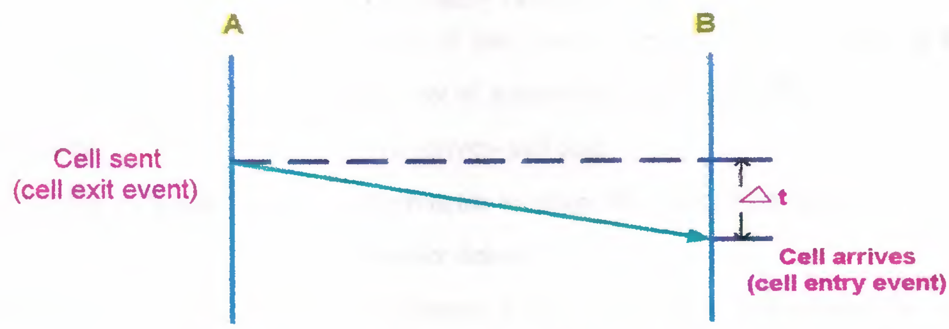

figure 3.8 cell transfer (schematically)

If Δt less than a maximum allowed time T (the exact value is not yet specified), then the cell has been successfully delivered, otherwise a lost cell outcome occurs (:e either the -cell arrives after T or it never reaches B). Errors in the VC/VP label field of the ATM cell header that cannot be corrected or cell buffer overflows in the network -(e.g. in an ATM switch) lead to lost cells.

If a cell arrives at B that has not been sent from A on this virtual connection, then such a misdelivered cell produces an inserted cell outcome. Label errors that are not detected or that are erroneously corrected may produce such inserted cells.

Successfully delivered cells arrive at their destination point in time but nevertheless may be errored when one or more bit errors exist in the cell information field. A cell is called severely errored when at least N > 1 bit errors occur in the cell information field with N yet to be specified (e.g. N = 2).

By making use of the above considerations, one can define the performance parameters. The parameters and their definitions are listed in (Table3.2)

Table 3.2 ATM performance parameters

Parameter	**Definition**
Cell loss ratio	Ratio of the number of lost cells to the sum of the number of lost and successfully delivered cells
Cell insertion rate	Number of inserted cells within a specified time interval (or per connection second)
Cell error ratio	Ratio of errored cells to the number of successfully delivered cells
Severely errored ratio	Ratio of the number severely errored cells to the number of successfully delivered cells
Cell transfer delay	Cell arrives- cell sent
Mean Cell delay variation	Arithmetic average of a specified number of cell transfer delays
Cell delay variation	Difference between a single observation of cell transfer delay and the mean cell transfer delay on the same connection
Cell transfer capacity	The maximum possible number of successfully delivered cell outcomes occurring over a specified ATM connection during a unit of time

Note that in (Table 3.2) cell loss and cell error that ratios are given whereas the cell insertion outcomes are measured by a rate (i.e. events per time unit). As the mechanism by which inserted cells are produced has nothing to do with the number of cells on the observed connection, this performance parameter cannot be expressed as a ratio, only as a rate.

Bit errors in errored cells can be corrected to a certain extent by error protection methods applied to the cell information field contents.

Lost and inserted cells can cause severe trouble in the event that they are not detected; e.g. for constant bit rate, the real-time services synchronism between sending and receiving terminals may be disturbed. Lost and inserted cell events become detectable (in many cases) by the supervision of a sequence number in the cell information field or an equivalent mechanism.

Both- cell transfer delay and cell delay variation must be kept within a limited range in order to meet service requirements. Cell delay and delay variation are induced, for example, by ATM multiplexers or ATM switches.

3.5 Traffic Control And Resource Management

3.5.1 Over view of Traffic Control Functions

To ensure the desired broadband network performance outlined in the previous section, an ATM-based network will have to provide a set of traffic control capabilities. CCITT Recommendation 1.311 [50] identifies the following:

- Connection admission control
- Usage parameter control
- Priority control
- Congestion control.

Their definition and description are given in the subsequent sections.

Connection Admission Control

Connection admission control is defined as the set of actions taken by the network at the call set-up phase (or during call renegotiation phase) in order to establish whether a (VC/VP) connection can be accepted' (50].

A connection can only be accepted if sufficient -network resources are available to establish the connection end to end at its required quality of service. The agreed quality of service of already existing connections in the network must not be influenced by the new connection.

Two classes of parameters are foreseen to support connection admission control:

- A set of parameters describing the source traffic characteristics.

- Another set of parameters to identify the required quality of service class.

Source traffic can be characterized by its:

- Average bit rate
- Peak bit rate
- Burstiness
- Peak duration.

The exact definition of bur mess and peak duration is still a pending issue in CCITT (burstiness might be defined as the ratio of peak bit rate to average bit rate). The description of source traffic parameters and quality of service class parameters cannot be completed unless more profound knowledge about ATM service requirements is compiled.

Usage Parameter Control

Usage parameter control is defined as 'the set 'of actions taken by the network monitor and control user traffic in terms of traffic volume and cell routing validity. Its main purpose is to protect network resources from malicious as well as unintentional misbehavior, which can affect the quality of service of other already established connections by detecting violations of negotiated parameters.

Usage parameter control will apply only during the information transfer phase of a connection. Connection monitoring encompasses all connections crossing the user network interface, including signalling [50].

Usage parameter monitoring includes the following functions:

- Checking of the validity of VPI/ VOl values
- Monitoring the traffic volume entering the network from active VP and VC connections in order to ensure that parameters agreed upon are not violated
- Monitoring the total volume of the accepted traffic on the access link.

What is actually performed will depend on the access network configuration.

The parameters subject to monitoring and control may be the same as those used for source traffic characterization to support connection admission control, namely average and peak bit rate, burstiness and peak duration. However, further studies are required.

Usage parameter control can simply discard those cells, which violate the negotiated traffic parameters. In addition, a 'guilty' connection may be released. Another, less rigorous option would he tagging of violating cells. These cells can be transferred as long as they do not cause any serious harm to the network. Thus the overall throughput of ATM cells might possibly be raised.

To illustrate the concept of usage parameter control, (Figure 3.9) shows different access network arrangements with the appropriate usage parameter control measures applied to VCs or VPs at the access point where they are terminated within the network.

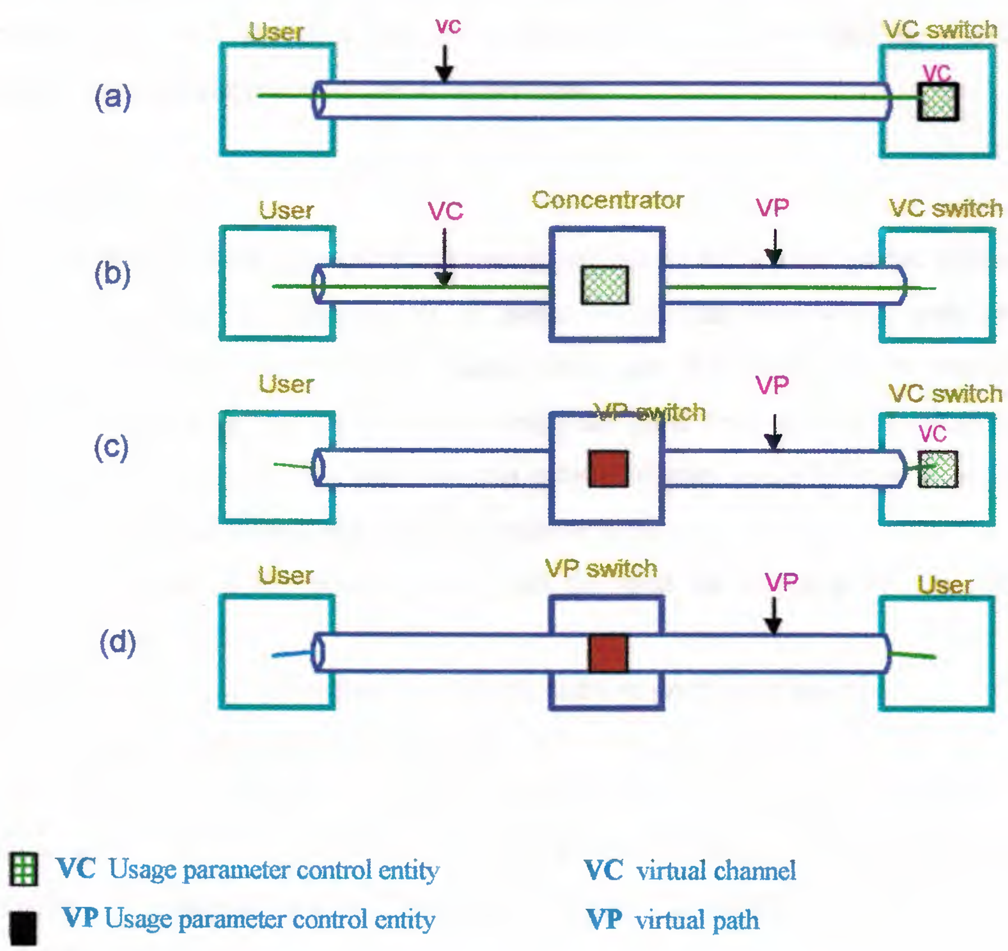

Figurer 3.9: Illustration of usage parameter control

In case (a), a user is connected directly to a VC switch. Usage parameter control is performed within the VC switch on VCs before switching them.

In case (b), a user is connected to a VC switch via a concentrator. Usage parameter control is performed within the concentrator on VCs only.

In case (c), a user is connected to a VC switch via a VP switch. Here usage parameter control is performed within the VP switch on VPs only and within the VC switch on VCs only.

In case (d), a user is connected to another user via a VP switch. Usage parameter control is now performed within the VP switch on VPs only.

Usage parameter control principles may also be applied to control the volume of traffic coming from other networks at the entry of an ATM-based network (from the ATM network's viewpoint, the other network is considered a big user). This so – called inter-network usage parameter control merits further study.

Priority Control

ATM cells have an explicit cell loss priority bit in the header; so two different ATM priority classes can be distinguished. A single ATM connection virtual path or channel level) can comprise both priority classes when the information to be transmitted is classified by the user into more and less important parts (one possible application may be layered video-coding). In this case the two priority classes could be treated separately by connection admission control and usage parameter control.

The details of this priority control and the need for additional priority mechanisms are still an open issue.

The different buffering mechanisms for systems with two priorities is described:

***Common buffer with pushout mechanism*:**

Cells of both priorities share a common buffer. If the buffer is full and a high priority cell arrives, a cell with low priority (if any is available) will be pushed out and lost. In order to guarantee the cell sequence integrity, a complicated buffer management mechanism is necessary.

***Partial buffer sharing*:**

Low priority cells can only access the buffer if the total buffer filling is less than a given threshold SL (SL < total buffer capacity). High priority cells can access the whole buffer. By adjusting the threshold SL it is possible to adapt the system to various load situations.

Buffer separation:

For the two priorities, different buffers are used. This mechanism is simple to implement but cell sequence integrity can only be maintained if a single priority is assigned to each connection.

The results showed that the system performance can be improved using priorities and that the partial buffer sharing mechanism is a very good strategy (best compromise between performance and implementation).

Congestion Control

Congestion in the context of B-ISDN is defined as 'a state of network elements (e.g. switches, concentrators, transmission links) in which, due to traffic overload and/or control resource overload, the network is not able to guarantee the negotiated quality of service to the already established connections and to the new connection requests. Congestion can be caused by unpredictable statistical fluctuations of traffic flows or fault conditions within the network.

Congestion control is a network means of minimizing congestion effects and avoiding the congestion state spreading. Congestion control can employ connection admission and/or usage parameter control procedures to avoid overload situations.

For example, congestion control could reduce the peak bit rate available to a user and monitor this (and react accordingly if the user exceeds the new parameter value).

Another congestion control procedure is called fast reservation protocol. It is claimed to allow an intelligent multiplexing of sporadic and variable bit rate sources both without real-time requirements. When using such a procedure, cell losses caused by the statistical behavior of the individual traffic streams can be reduced.

3.5.2 Traffic Control Procedures and Their Impact on Resource management

Traffic control procedures for ATM networks are currently not standardized within CCITT. In fact, it has not been decided yet whether, and to what extent, they should be standardized at all. On the one side, network providers may have a desire for flexible network tools to be able to react adequately to customers' needs (which 4 speaks against standardization), while on the other side, for the benefit of users and especially terminal manufacturers, settled network standards are indispensable. For instance, a terminal basically needs to know how the network handles its ATM cells under normal and also fault conditions in order to shape its cell flow according to the rules of the ATM network and thus to be able to use the transport facilities of the ATM network optimally.

The choice of traffic control algorithms directly impacts on a network's resource allocation strategy. For example, if only the peak bit rate of a connection were considered for admission and usage parameter control, then this peak bit rate would have to be allocated to the connection. If this connection had a low average bit rate, then most of the time the network would exhibit a quite poor efficiency. Nevertheless, such simple strategies can assist in quickly introducing ATM-based networks; as long as knowledge about the management of ATM traffic flows is rather limited, due to the fact that neither the source traffic characteristics nor the actual traffic mixes on a link are sufficiently clear, it might be wise to stay on safe ground even if a considerable amount of network capacity is wasted. Besides peak bit rate reservation, a restricted utilization (e.g. 70 %) of the cell transport capability of a link helps to avoid congestion.

The goal is simultaneously to:

- Achieve good ATM network efficiency
- Meet the users' quality of service requirements

with a method that is generally applicable. Therefore, more sophisticated traffic control measures and resource management actions are being taken into account.

In this approach, both average and peak bit rate are considered and, in addition, the upper bound for the bit rate variance that is dependent on the behavior of the source (a good representation of the bit rate variance is especially important when-the source is a video codec). The results are promising: this method has a considerable advantage compared with mere peak bit rate reservation if the peak bit rates of most of the connections on a link are small in relation to the total bit rate of the link and if the peak-to-average bit rate ratio (burstiness) of these connections is high.

Different usage parameter control mechanisms have been proposed, namely leaky bucket, sliding window, jumping window and exponentially weighted moving average. The comparison of these mechanisms showed that the leaky bucket and the exponentially weighted moving average are most promising with respect to flexibility and implementation complexity.

The basic problem of ATM networks is the statistical behavior of the cell arrival process, e.g. at a buffer where cells generated by several sources are multiplexed together.

The ATM traffic can be described by a three-level hierarchical model as depicted in (figure3.10).

The call level has a typical time scale of seconds up to hours, the, burst level is related with the millisecond range up to seconds, and the cell level with the microsecond range. These levels have different impacts on network implementation, whereas cell level analysis can provide information for buffer dimensioning (e.g. in ATM switches or mtiltiplexers), the burst level characteristics influence mainly call admission strategies, and the call level considerations relate to ATM link dimensioning.

Based on the traffic model shown in (Figure 3.10), mathematical models have been defined with different levels of abstraction including:

- Burstiness (peak bit rate/average bit rate).
- geometrically distributed burst lengths.

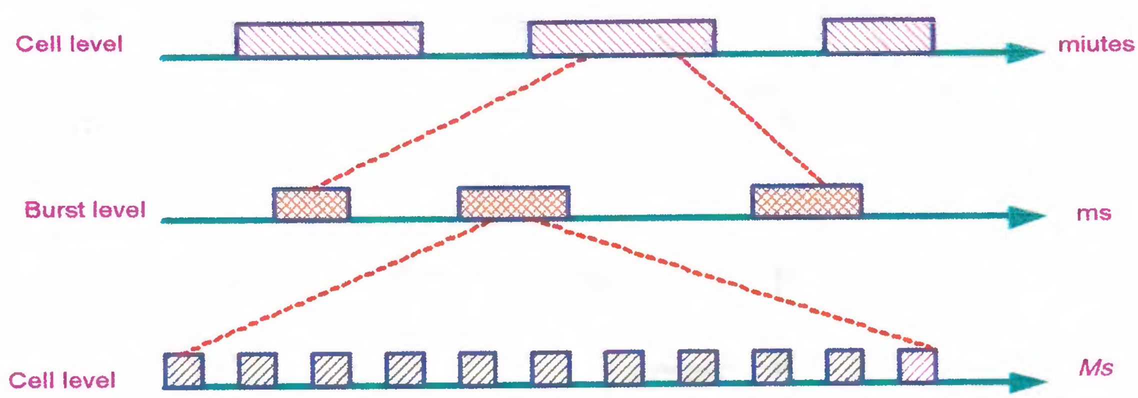

Figure 3.10: Hierarchical modelling of ATM traffic

- Switched Poisson process
- Markov-modulated Poisson process.
- Generally modulated deterministic process.

It has been found that the quality of service parameters (e.g. jitter, loss probability) are very sensitive to the assumed source characteristic. Therefore, it is necessary that, for performance evaluation, these detailed source traffic models are used.

3.6.2 OAM Levels in B-ISDN

The OAM levels of the ATM transport network coincide with the levels introduced before on network layering (see Figures 3.2 and 3.4):

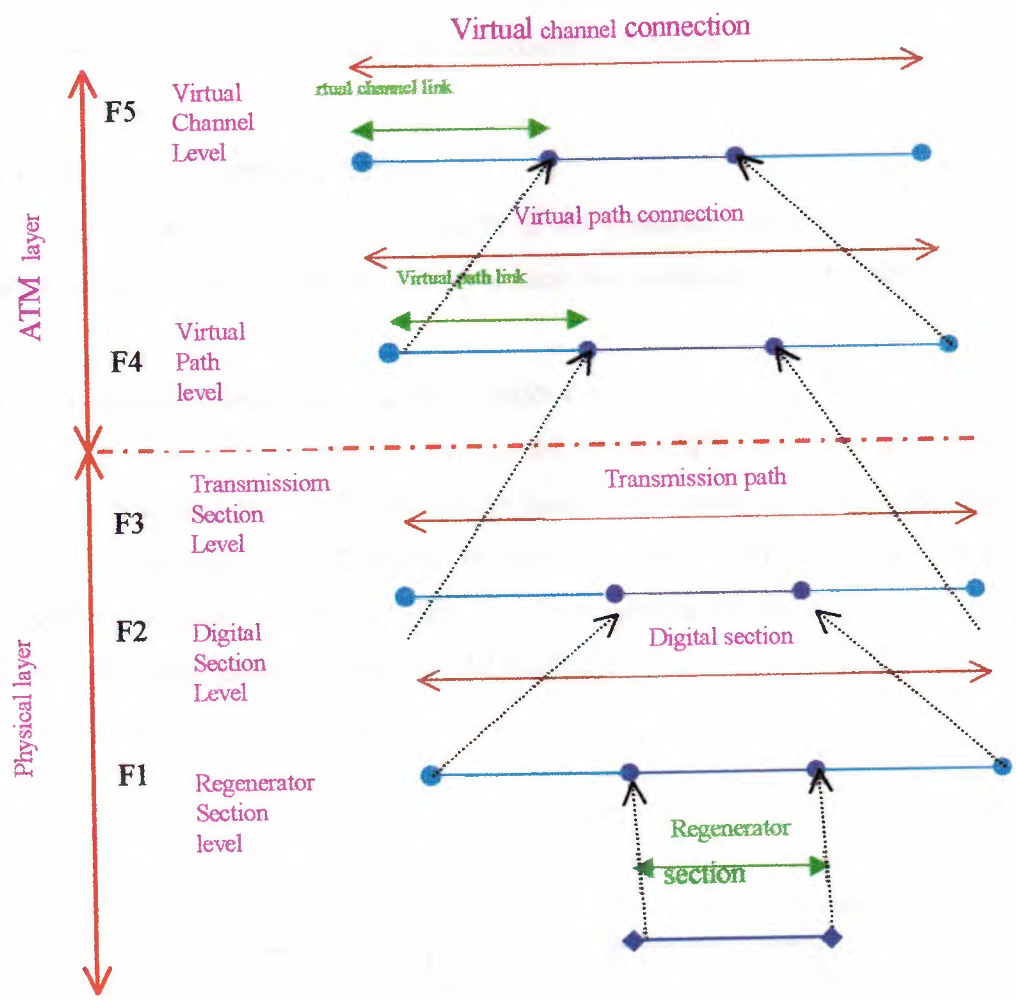

ATM *Asynchronous transfer mode*

OAM *Operation and maintenance*

● *Endpoint of the corresponding levels*

◆ *Connecting point of the corresponding levels*

Figure 3.11: OAM hierarchical levels,

The ATM transport network comprises the physical layer and the ATM layer, and these layers are subdivided into regenerator section, digital section and transmission path level, and virtual path and virtual channel level (of Figure 3.11).

The corresponding CAM information flows of each level- denominated Fl to F5 are also depicted in (Figure 3.11).

The OAM Bows are bidirectional. As an example of an OAM flow, consider the monitoring of a VPC by means of monitoring cells sent out at one endpoint of the VPC and mirrored at the other endpoint, to be evaluated at the sending side.

3.7 Customer Network Aspects

One essential block of the B-ISDN is the customer network (CN). Sometimes it is called customer premises network (CPN) or subscriber premises network (SPN).

3.7.1 Reference Configuration of the B-ISDN UNI

The reference configuration for the 64 kbit/s ISDN user-network interface (UNI), which is described in CCITT Recommendation 1.411- was accepted to be general enough also for application in the B-ISDN environment.(Figure 3.12) shows the principles of the reference configuration for the B-1SDN UNI. It contains the following:

- Functional groups: B-NT1, B-NT2 and B-Tel
- Reference points: T_B and S_B

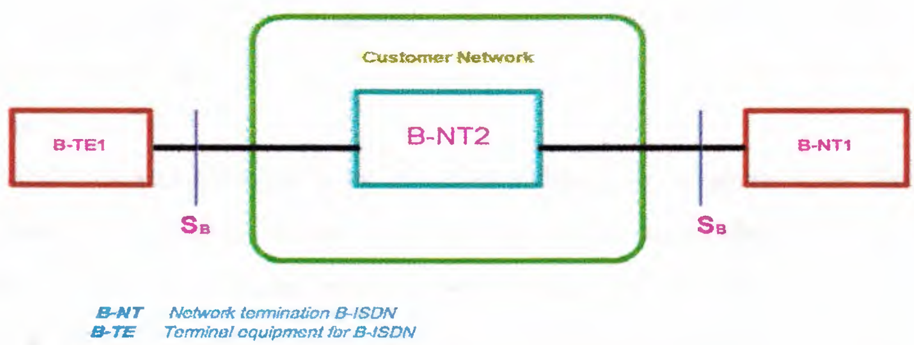

Figure 3.12: Reference configuration of the B-ISDN UNI

B-NT1 and B-NT2 are generic terms and denote broadband network terminations. B-TEl is the acronym for a broadband terminal with standard interface. At the reference points T_B and S_B, physical interfaces may or may not occur. If they are realized, they must comply with the specified standard. While the 3.NT1 performs only line transmission termination and related OAM functions, the B-NT2 may be, for example, a private branch exchange or LAN, which performs multiplexing and switching of ATM cells.

The CN covers the area where users have access to the public network via their terminals. This is the part of the telecommunication network located at the user side of the B-NT1.

For the functional description of the CN, the reference configuration shown in (Figure 3.12) can be used. The interface between the CN and the public network is usually at the reference point TB. So the CN coincides with the functional group B-NT2 (see Figure 3.12).

3.7.2 Customer Categories

Different aspects for the classification of CNs are possible, e.g. environment (residential, business), number of users or topology.

The residential environment is characterized by a small number of people (e.g. a family) using broadband services mainly for entertainment purposes. This category is considered to be more or less homogeneous and will be restricted to one flat or house. In many cases, no internal switching capabilities are necessary within a residential CN.

The counterpart of the residential CNs are the business, CNs which are subdivided according to their size:

Small business CNs:

Small business CNs has a lot of commonality with residential CNs. Such a CN will be installed in small areas like an office-or a shop with -up to about 710 employees. Often the office or shop is combined with private ~housing and therefore the requirements of the residential CN (e.g. -entertainment services) must- also be fulfilled. One major difference to the residential CN is that there may be a need for internal switching.

Medium business CNs:

In medium and large business CNs, as well as factory CNs, no distribution service for entertainment purposes need be supported. Normally, only interactive services, e.g. telephony, videoconference or high-speed data transmission, will be used. One of the main characteristics is the internal switching capability. A medium business CN has typically between 10 and 100 users.

Large business CNs:

Large business CNs has more than 100 users. They spread over several buildings or floors of a building. Sometimes distances of up to 10 km must be spanned.

Factory CNs:

Factory CNs can have the size of a medium or large business CN. They often have to satisfy exceptional physical requirements such as being robust against extraordinary heat, dust or electromagnetic interference.

The common characteristic feature of-the medium, large and factory CNs is the provision of internal communications This facility-is not so often required in the residential and small business environment.

3.7.3 General Requirements

Two types of requirements are expected to be met by a CN:

- Service requirements
- Structural requirements.

Service requirements deal with the service mix as well as the consequences of having to support these services. The service mix is dependent on the customer categories. This type of requirement includes the bit rate to be supported, the information transfer characteristics like mean and maximum delay, delay variation, error performance and throughput.

It is possible to characterize realistic service mixes for each of the customer categories, but it is rather difficult to estimate the evolution of coding techniques (e.g. variable bit rate video-coding), which influence required bit rate of a specific service.

The second type of requirements is called structural requirements, which include aspects of flexibility, modularity, reliability; physical performance and cost.

Flexibility of the CN is the ability to cope with system changes. This can be subdivided into four parts

***1. Adaptability*:**

Is the requirement, which measures how the - CN deals with changes that do not alter the global scale of the CN (e.g. new wiring). This is very important in the terminal area for the residential as well as the business environment.

***2.Expandability*:**

Indicates how the CN can grow (e.g. introduction of new services, increasing the bit rate to be supported, or installation of new terminals, expanding the scale of the CN).

3.Mobiltiy:

is the ability of interchanging terminals. This requires a universal terminal interface.

4.Interworlcing:

describes how the CN can interface to other networks. This is very important in areas which are already covered by existing networks (e.g.LAN, private MAN).

Modularity is the provision of a flexible structure. The network should not be limited to a few applications. Therefore, it is necessary that a modular system is used for the provision of CN capabilities.

Reliability deals with the sensitivity of the CN to errors (e.g. bit errors, terminal failures and user-induced errors). This requirement is extremely important when a large number of people or terminals are afflicted by the error or in all cases where a. very good operation of the CN must be guaranteed (e.g. hospitals or fire departments). Reliability requires redundancy and therefore within large CNs even the terminal connection is duplicated.

Physical performance concerns the optimum use of the physical medium. It includes aspects of coding efficiency and cable length and influences the hardware cost. Installation and maintenance are covered by the operating performance. This must be very simple in the terminal environment so that changes can be carried out rapidly and cheaply.

Cost is one of the most important requirements influencing the acceptance of the CN.1n the residential area, low costs are essential whereas in the business area it is necessary to have reasonable costs during the introduction phase. Fast system growth can only be achieved if the incremental costs can be kept low.

3.8 B-ISDN Local Network Topology And Technology

3.8.1 Local Network Structure

The conceptually simplest realization of the B-ISDN local network is the star topology where there is access per customer. The remaining problem to be solved could then be to define the characteristics of the digital section extending between T_B and V_B reference points (see Figure 3.13 (a)).

For the B-ISDN local network, however, other structures are also being investigated as Candidates for implementation. These are, for example, multiple star, ring, Ims and tree configurations. Some examples are depicted in (Figure 3.13)

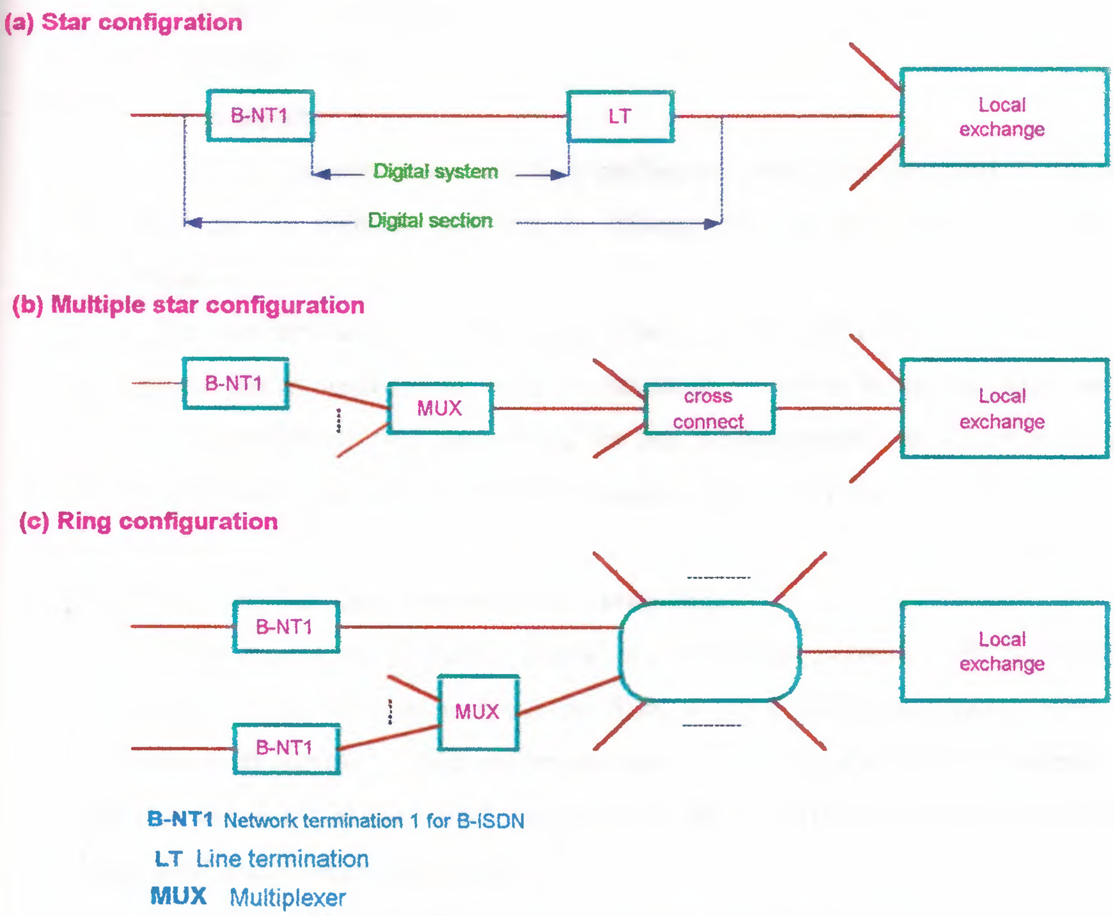

Figure 3.13: : Examples of local network strucures

Case (b) shows the multiple star configuration where several customers' signals are multiplexed on to one access line. Optionally, there might be an additional cross-connect to provide different paths through the network. Its functionality may comprise:

- Connection of customers with big traffic volumes directly to the local exchange.
- Separating traffic that is to be switched in the local exchange from the traffic that is to be routed on a fixed path (permanent connection) through the network.

Configuration (c) is a ring structure where many customers share a common transmission medium (e.g. MAN). Customers with low traffic may share the access to the ring as shown in the lower part of (c). The ring may be a single or double ring depending on:

- Number and location of customers
- Throughput requirements
- Availability criteria
- Cost aspects.

When a customer requests a high performance and reliability level the network may be connected via multiple interfaces to different ring access nodes or even to different exchanges.

The ring structure saves transmission lines in terms of meters of cable to be laid in the ground but is restricted in terms of bandwidth available to an individual user and in terms of upgradability (the same holds for bus configurations). In a star or multiple star configuration it is easier to connect additional customers to the B-ISDN.

3.8.2 Transmission Characteristics and Technology

Transmission of B-ISDN signals in the access network will be predominantly performed via optical systems based on SDH (CCITT Recommendations 0.707-709 [34, 35, 36]) which provide bit rates of approximately 155 Mbit/s (STM-1), 622 Mbit/s (STM-4) and 2.5 Gbit/s (STM-16). PDH systems and other transmission media (e.g. radio links) could also be used where appropriate.

Use of single-mode fibres according to CCITT Recommendation 0.652 [31] is favoured in the access network as this allows longer distances to be covered. Optical signals can be generated by laser diodes. The electronic parts will be based on CMOS technology for bit rates up to 155 Mbit/s and on bipolar technology for higher bit rates.

Two optical fibres (one, for each direction, of transmission) may be employed, or alternatively optical wavelength division multiplexing (WDM) on one single fibre to separate both transmission directions by different wavelengths (e.g. 1530 rim and 1300 run). Later, a more advanced technology like coherent multi-channel transmission will possibly be employed.

The bit error probability of the optical transmission link should be less than 10^{-9} An ATM switch will guarantee a cell loss probability of about 10^{-9}. Cell loss in an ATM switch is primarily caused by buffer overflow in the case of traffic congestion or detection of uncorrectable errors in the cell header. If the transmission bit error probability is less than10^{-9}, the latter effect is negligibly small-as -it requires the occurrence of specific multiple-bit errors in one cell header.

3.8.3 Maintenance Aspects of Optical Transmission

The following failures may occur in optical transmission systems:

- Laser light emission failure
- Laser-sends continuous 1s
- Receiver indicates continuous 0s or 1s.

In the case of such a failure the corresponding message will be delivered to a line maintenance entity (usually located in the local exchange) which is responsible for setting the appropriate alarms, i.e. urgent alarm in the case of a unit out of order and warning in the case of a unit exhibiting deteriorated performance.

3.9 Trunk Network Structure

A possible trunk network implementation is shown n Figure 3.14(Realization of internodal signalling connections is not discussed in this chapter. Signalling may be established via the exiting SS7 routes or, in the long run, over the network itself.).

(In Figure3.14) the following network elements are visualized

- B-ISDN exchange (VC/VP switch)
- ATM cross-connect (VP switch) -
- STM multiplexer/cross-connect.

The ATM cross-connects act as VP switches and can flexibly provide VP connections though the network. STM cross-connects may be deployed in addition to facilitate rearrangement of physical paths, e.g. in the case of transmission failures (protection switching). Finally, STM multiplexers merge, for example, 155.520 Mbit/s STM-1 signals into the higher bit rate signals STM-4 or STM-16 of about 622 Mbit/s or 2.5 Gbit/s, respectively.

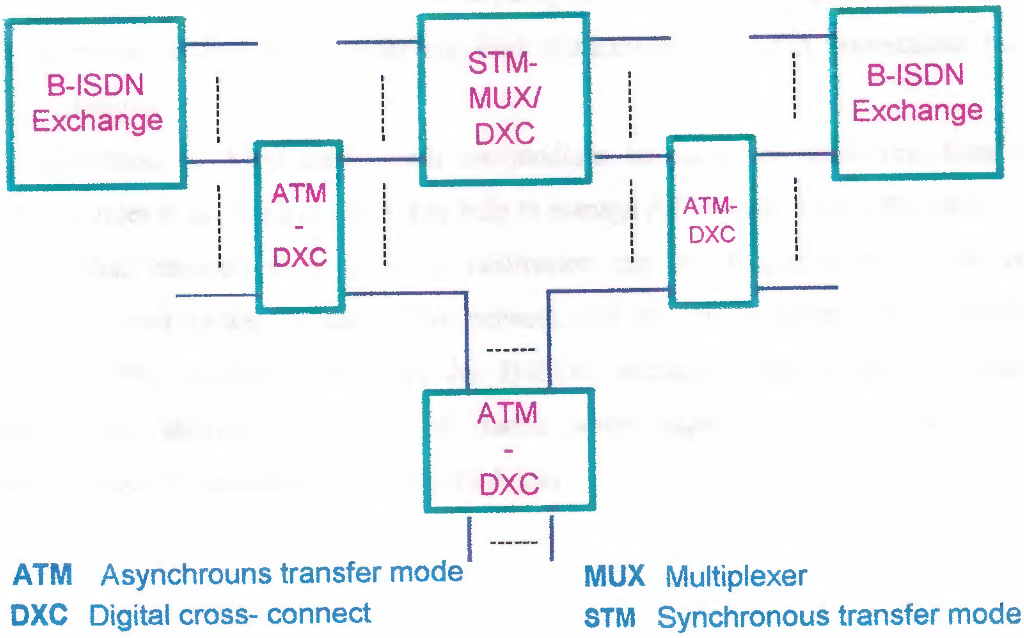

Figure 3.14: Example of trunk network structure

ATM cross-connects processes each arriving cell; according to its VPI value, a cell is routed into the direction that is defined by the associated VPC. Therefore, VPCs with arbitrary bit rates can be established and switched through an ATM cross-connect.

The user or the network provider by means of, for example, ATM layer management or probably in a later stage signalling procedure can in principle initiate Establishment and release of VPCs.

The time required for abolishment or release of VPCs will be much shorter than the realization of the request for a reserved, permanent channel in today's networks, i.e. VPC establishment can be performed within a couple of seconds. The same time scale applies to the automatic rearrangement of SDH paths in the case of transmission failures (protection switching).

Multiplexers and cross-connects can be used to decouple the logical point-to-point configuration of the ATM switching network from the actual topology of the fibre-based transmission network and to achieve:

- Flexibility concerning realization of hierarchically structured networks or net-i works with any other structure (e.g. meshed rings)
- Easy provision of additional network capacity in the case of growing traffic
- A means to be able to offer required redundancy for ATM connections between exchanges.

Segregation of ATM traffic with low/medium bit rates per connection from ATM with high bit rates in the trunk network may help to manage ATM traffic more efficiently.

Individual channels with the same destination can be grouped to be carried on the same VP. Transit nodes of the ATM network act as VP switches with reestablished (redundant) VPC capabilities between the B-ISDN exchanges and need not handle VC switching. This alleviates operation of transit nodes especially in the case of many simultaneous low bit rate connections on a trunk line.

4. ENABLING TECHNOLOGIES FOR BROADBAND

4.1 Fiber Optic Technology

4.1.1 An Overview

Fiber optics is fundamental to the whole broadband endeavor. Although some segments of future broadband networks will be operated over microwave radio (because cabling is impractical) and some segments over copper cable (because this is less expensive than fiber and can carry the same data rates over short distances), it is fiber optics that makes broadband networking possible.

TABLE 4.1 Broadband Networks: Aspects And Technologies

Network Aspect	Technology
Bandwidth	Fiber optics Compression
Switching	Fast packet/ATM
Intelligence	Software

The theory behind fiber optics is quite simple. To simplify somewhat, the data rate of a transmission system is a function of the bandwidth of the electromagnetic radiation used in the system. The infrared light used in modern fiber optic systems has much higher frequencies than the electrical waves used in conventional transmission systems. Hence, fiber optic systems can carry much higher bandwidths. In theory, fiber optics can easily handle transmission running into the terabit-per-second range. In practice, however, fiber optic transmission systems are highly inefficient.

One key reason for this is that fiber optic systems are seldom wholly optical. Instead, fiber optic networks have historically consisted of point-to-point fiber optic links with electrical switches, multiplexers, and amplifiers serving at the nodes of the network. The inclusion of electrical components has tended to reduce the speeds at which such optical networks can operate.

Much of today's research in the fiber optics field has been concerned with making fiber optic networks more fully optical. In this way it is hoped that the full potential of fiber optic networking to deliver very high bandwidths for broadband networking can be achieved. As we discuss later, this research direction has led to the creation of optical amplifiers, multiplexers, and switches. Much of the rest of the research currently being carried out in the area of fiber optics is directed at reducing the cost of implementing fiber optic systems. Although this cost has declined considerably in the past decade, cost is still a limiting factor in using fiber optics for many applications.

4.1.2 How Fiber Optics Works

Flashing lights on and off to convey information is hardly a new idea. Paul Revere used light signalling at the Old North Church during the American Revolution. Railroad flagmen have long used lanterns moving in a special pattern to communicate commands from one end of the train to another. And in 1880, four years after the invention of the telephone, one crank invented a "photo phone" that he hoped would eventually replace the telephone. This crank was a man called Alexander Graham Bell.

The problem with all these different forms of optical communications is that objects got in the way of the light beam. This has been a constant problem with all over-the-air optical communications. Clouds, fog, flocks of birds, buildings, and anything else that blocks light are obstacles to light communications sent through the air. Another difficulty with over-the-air light communications is the problem of dispersion. As a light beam shines outward from its source it becomes more diffuse and weakens. Try shining a flashlight into the night sky and see what happens.

One obvious way of dealing with such problems is to confine light used for communications in some kind of pipe. For a short while "light pipes" were constructed from rigid straight-line segments and a series of mirrors. But this approach to optical communications never went beyond the experimental stage. It turned out to create more problems than it solved. Not only was the civil engineering needed for light pipes prohibitively expensive, but light pipes were also easily disturbed by movements and vibrations in the earth.

Light pipes turned out to be a blind alley in the evolution of optical communications and were replaced by hair-thin strands of optical fiber.

The basic concept of fiber optics is not hard to understand. At one end of a fiber optic circuit a light source flashes on and off millions of times each second. This signal is then transmitted through a glass fiber to a light detector; which picks up the signal at the other end of the link. The idea goes back to the 1950s when the American Optical Corporation developed fiber capable of carrying light in this way. However, what really got the idea of optical communications moving was a 1966 paper by Charles Kao and George Hockman, two researchers at Standard Telecommunications Laboratories in the United Kingdomfi This paper suggested that glass fiber could transmit information at optical frequencies over several hundred feet if only the glass from which it was manufactured could be made pure enough.

By removing impurities and using glass with a high silica content, it became possible to manufacture a glass with super-high transparency If a light is shone through a block of ordinary window glass 1 meter thick, it will have lost more than 65 percent of its power by the time the light gets out the other side. Before such a power loss occurs in an optical fiber the light will have traveled several thousand feet!

But how is the light kept in the glass? The answer to this obvious question is that the glass core of the fiber through which the light travels is surrounded with a transparent cladding having special properties. This cladding is usually also made of glass, although occasionally it is made of plastic. Glass may not seem like a very promising material for cladding, because it is transparent. However, the core and the cladding of the fiber are of different transparencies, so most of the light traveling from the core to the cladding is reflected back into the core, due to a physical phenomenon called total internal reflection. This phenomenon will be familiar to anyone who has ever snorkeled and looked up at the surface a few feet above her because of total internal reflection she cannot see what lies above the water.

Total internal reflection is also what keeps light in an optical fiber. Because of total internal reflection, the light hitting the cladding of the fiber does not get through it. Instead the light "bounces" against the cladding and back into the fiber core.

But different rays of light travel down different paths within the fiber core and arrive at the other end of the fiber at different times. This phenomenon is called pulse dispersion, and it limits the amount of information that can be carried over the fiber. As the light source is pulsed more and more rapidly, we eventually reach a point where the light detector at the other end of the fiber cannot distinguish between rays that arrive at the end of one pulse and those that arrive at the beginning of the next. Obviously the further the signal has to travel, the longer the time between the arrival of the first and the last rays emanating from a single pulse. In other words, the pulse dispersion problem gets worse the further the optical signal must travel down the fiber.

Fiber in which pulse dispersion is a significant problem is called multimode fiber (because the rays take different paths). However, since the mid1980s a new type of fiber, called single-mode fiber, has become available. The core of single-mode fiber is so fine that light waves travel down it as if they were a single ray Fiber of this kind has a core with a diameter of less than 10 micrometers. Using it, pulse dispersion is all but eliminated; so huge amounts of information can be carried over very long distances.

Both single-mode fiber and multimode fiber are now used in fiber optics. At first multimode was priced considerably below single-mode, but now single-mode is actually less costly than multimode. Multimode fiber continues to be used because, with its larger core, it is easier to connect to optical transmitters and receivers, so the overall cost of a multimode system may be less than an equivalent system using single-mode technology. For fairly short hauls, multimode may offer perfectly adequate bandwidths. Thus multimode fiber has been the fiber of choice for fiber-based local area networks. However, long-haul systems are almost always single-mode. Since the mid-1980s, when the long-distance telephone companies started putting in optical transmission systems, they have used single-mode fiber almost exclusively, and even for local data communications single-mode fiber is increasingly used.

Fiber is not used "as is." As part of the manufacturing process fiber is usually coated with plastic to protect it from abrasion and interaction with the environment. Later multiple strands of fiber are combined in a single cable, providing further protection, In addition to the fiber itself, the cable contains strength members extra strands that help the cable stand up to all the stresses and strains that it will have to undergo during its lifetime. There are numerous different kinds of cable, each suited to a different purpose and each containing different numbers of fibers.

The fiber and cable is just half the fiber optics story. There are also the light sources and light detectors.

The light sources may be either light emitting diodes (LEDs) or diode lasers, which are similar in structure to LEDs. Both are tiny chips of semiconductor material about the size of a grain of salt with metallic layers at the top and bottom that act as electrical contacts for connecting the laser or LED to a power source. But there are important differences in what goes on inside the two types of light source. We will not discuss these differences here, however, but will note only that diode lasers are more complex and harder to manufacture than LEDs. They are therefore more expensive. But they have remarkable properties that are particularly useful in fiber optics.

Perhaps the most important of these properties is that fiber optic systems based on diode lasers can carry more information than those based on LEDs. This is because lasers can be pulsed on and off at a much faster rate than LEDs. The usual measure of this is called the "response time" (also known as the "rise time") of the light source. If the steady-state power of the light source is its "normal" power when not being flashed on and off, then its response time is defined as the time that it takes for the light source to increase from 10 percent of its steady-state power to 90 percent of its steady-state power. The response times of LEDs vary from a few to a few hundred nanoseconds (a nanosecond is one billionth of a second). Lasers have response times of a fraction of a nanosecond. Another reason for using lasers for communications over fiber is that the light from a laser beam is restricted to a narrow frequency range and this results in less signal distortion. In the jargon of optics, a laser beam is said to be highly "coherent."

The relative economics of lasers and LEDs have changed over the years in a manner similar to that of multimode and single-mode fiber. In the early days of fiber optics, LED-based transmission equipment predominated. This was not only because diode lasers were very expensive. It was also because they were considered too expensive, too short-lived, and too unreliable to be much use in a practical communications link. But as laser technology improved, they became the light source normally used in public networks, and the less expensive LEDs were found in local area networks and links between computers in a data center. However the performance and cost of the two types of fiber optic light source are growing ever closer. High-end LEDs are gradually beginning to match the performance of some diode lasers, and diode lasers are beginning to become available at the same cost as these high-end LEDs.

The current generation of optical transmission systems uses an electrically modulated laser with no amplification between repeaters. Such systems are more than adequate for most practical purposes. Commercial systems using this architecture are available supporting SONET OC-48 (2.5 Gbps) transmission over 80 km. However, long-haul systems operating at above 10 Gbps could probably not be supported using conventional modulation schemes.

Two approaches to "external modulation" are now being explored: absorptive and interferometric. Absorption modulators appear to offer the best chance for short-term commercialization because they use existing semiconductor materials. This means that their bulk manufacture will be easy to achieve, thereby lowering costs. Using common semiconductor materials also offers the possibility for easy monolithic integration with the laser. Interferometric approaches to external modulation, using phase modulators, offer the potential for higher performance, but such modulators have usually been constructed out of lithium niobate ($LiNbO_3$), an expensive material whose long-term stability in photonic devices is still in doubt. Last, but certainly not least, we come to the light detectors used in fiber optic communications. Like diode lasers and LEDs, the receivers that convert the optical signal transmitted through the fiber back into electricity are also photo electronic semiconductor devices, but whereas LEDs and diode lasers turn electricity into light, the semiconductors at the other end of the link turn light into electricity.

Photo detectors can be manufactured from a wide variety of semiconductor materials, including most of those that can be used for semiconductor light sources. The semiconductor material that is chosen determines the frequency of light to which the detector is sensitive. Until the late 1970s, detectors for fiber optic systems were usually made from germanium, but since then indium gallium arsenide (InGaAs) has become the material used in high-performance fiber optic systems.

One of the most significant aspects of fiber optics has been that it has achieved effective, high-bandwidth transmission over very long distances with no amplification or regeneration. For next-generation systems used in gigabit networks, however, a major challenge will he to increase output power at the transmitter and obtain high sensitivity at the receiver. An optical post amplifier at the transmitter and an optical preamplifier at the receiver can achieve this. Optical amplifiers using erbium in the core of silica fiber are already available and are increasingly used in commercial fiber optic systems, especially in the cable television industry. Optical amplification has (at least) two significant advantages: (1) Optical amplifiers are passive devices and therefore could be relatively inexpensive to manufacture and maintain, and (2) their operation is based on a stimulated emission phenomenon, which occurs very rapidly and can therefore handle very high modulation rates. The latter is obviously of particular importance in the context of gigabit networking.

Important developments boosting the sensitivity of optical detectors are also occurring, especially in the area of coherent detection (see below). Combining some of the new technology available for fiber optic systems, data rates of 10 Gbps over distances as long as 4500 km without electrical signal regeneration have been achieved under experimental conditions. However, for higher rates, wavelength-division (also called wave-division) multiplexing (WDM) of several high-speed channels is a more likely choice than boosting the capacity of single channels. WDM technology is rapidly maturing as both a multiplexing and switching technology. Most of the next-generation transmitting, receiving and amplification devices currently being built can be redesigned to operate at multiple wavelengths.

4.1.3 Coherent Fiber Optics

Fiber optic detection as we have described it is properly characterized as "direct detection," and is the norm in most commercial fiber optic systems installed today. Another form of detection is known as coherent fiber optics. Although not yet fully commercialized, it has a special relevance for broadband networks.

Coherent systems are much more complex and expensive to build than conventional fiber optic systems, but they offer the possibility of achieving significantly higher bandwidths than most systems installed today. This is due to higher rates of receiver sensitivity. One useful (although admittedly simplistic) way of thinking about this is that coherent receivers need to receive fewer photons for a given piece of information than a regular receiver, and this translates into higher data rates i.e., more pieces of information successfully received per interval of time.

Here is how a coherent system works. The light from the transmitter laser is amplitude, frequency, or phase modulated and then transmitted through the fiber. At the receiving end of the link the modulated carrier signal is mixed with light from a local oscillating laser signal, and the combination is then detected. This may be done in several ways, but in all cases the point is that the information in the carrier signal can be detected by detecting the deviations from the local oscillator caused by the mixing of signals. These deviations can be usually be found in the radio frequency/microwave range, where they can be detected using electrical equipment that is of a much higher degree of sensitivity than the optical equipment used in conventional fiber optics.

Although bandwidth improvements may be the main advantage of coherent optics from the point of view of the kind of networks that we are considering in this book, another advantage is that because of the higher sensitivity of coherent detection systems, in cases where different channels are carried over the same systems, channels can be spaced closer together. According to one source, an improvement of two orders of magnitude compared with conventional fiber optic systems can be expected. This does not represent an improvement in bandwidth per Se, but it does represent an improvement in usable bandwidth, which from the point of view of practical engineering is much the same thing.

In particular, coherent detection could potentially be used to put tens of thousands of video channels on a fiber optic network, where only hundreds could be carried before. This is clearly in line with the objectives of broadband networking as discussed inChapter1.However, widespread practical implementation of coherent optics is limited to some extent by the availability, reliability, and cost of suitable lasers. For example, because small differences are being detected in coherent fiber optics, the laser must exhibit stable frequencies, and tenability of the lasers used in coherent schemes is also to be desired.

4.2 Broadband Switching Technology

In fact, all the applications for which broadband networks are intended place heavy demands on the switches that are installed in such networks. First, almost by definition, these switches must be capable of switching circuits with very high bandwidths into the gigabit-per-see-cir and range in some instances. The conventional circuit switches and packet switches used in today's networks, which are limited to switching channels of a few megabits per second, cannot accomplish this. There are cross connects that can handle channels of much larger bandwidth, and these could be regarded as switches, but they are used not to rapidly establish and disconnect calls between people thousands of times per day in the manner of a regular telephone exchange, but rather to reconfigure networks a few times a day or few times a week.

The function of broadband switches is to provide the end user with what is called bandwidth on demand (BOD) in response to the fact that many broadband applications require high bandwidths, but require them for just a short space of time. Such applications are said to be "bursty" There are many examples of such applications, including videoconferencing, LAN internetworking, computer backup and disaster recovery image file transfer, client/server applications, messaging, and multimedia networking. Not all of this application requires broadband bandwidths, but many of them do. Such applications may run for just part of the day this is typical for computer backup, for example. Or traffic patterns may be so variable that the bandwidth required at one second (or microsecond) may be quite different from the next second or microsecond. This latter pattern is that associated with multimedia networking or videoconferencing. Bursty applications can be supported over a network made up of fixed-bandwidth private lines in fact they often are. But such an approach is inefficient, because bandwidth gets wasted. Much more than elegant engineering is at stake here. Ultimately, wasted bandwidth translates into wasted user dollars. Even the low-cost bandwidth associated with fiber optic transmission is not low-cost enough to allow bandwidth to go to waste.

Thus network managers are looking for a broadband switching technology that is finely tuned to the traffic patterns of their network and applications. The ideal technology would:

(1) Be able to switch very-high-bandwidth channels to the desktop, providing this bandwidth when and only when it is required.

(2) Provide a means by which bandwidth would be paid for only when it is used.

Today no widely available switching technology quite matches these two criteria, but both circuit-switched and packet-switched technology can make approximations to BOD.

4.2.1 Circuit-Switched ROD

Both ISDN and Switched 56 kbps are circuit-switched services with significant capabilities to provide BOD. These services offer moderate to high bandwidths, and by using inverse multiplexers, 56 kbps or 64 kbps channels can be built up to bandwidths of around 8 Mbps. Also facilitating the use of these services for BOD is the fact that they are available on a dial-up basis. Inverse miiltiplexers allow responses to changes in-the demand for bandwidth to occur within 1 second.

These services are available at the wide-area level, but circuit switching may also be an appropriate technology for providing BOD on the customer premises, with the company PBX as the platform for this application. Back in the 1980s, the prevailing philosophy was that the PBX would become the central office controller, switching voice, data, and perhaps even video. But the PBX-as-office-controller vision turned out to be incorrect, mainly because LANs first Ethernet and Arc Net, then token-ring, then FDDI came along and handled data a lot better than the PBX could do.

Today, many PBXs have data terminals attached to them and do some data switching, but data is a secondary consideration in most PBX networks. PBX vendors have not generally targeted BOD applications with their products, nor have the corporate telecommunications managers who traditionally deal with PBXs shown any stron~ interest in such applications. However, PBXs clearly do offer some kind of reasonable support for BOD.

4.2.2 Packet-Switched BOD

As we noted in the previous chapter, the alternative to circuit switching is packet switching. Until the 1960s, communications switches were generally divided into two types. The most familiar type (at least to the general public) were the circuit switches used for public and private telephone networks, but there were also "message switches" that used a "store-and-forward" technology in which terminals send messages to a central switch or computer which stores them and eventually transfers them through a network of other nodes until they reach their final and intended destinations.

Thus no permanent circuit is established between users during the communication. Message switching was originally used in telex networks and is now also associated with fax, electronic mail, and voice mail networks.

Packet switching is an extension of the message-switching concept. In conventional message switching, the whole message is kept intact while it is being passed through the network. In packet switching, the message is broken down into information "packets," each of which follows its own path through the network. Packet switching has its origins in the defense community of the early to mid-1960s but has now become the primary means of interconnecting computers over wide areas.

Packet switching allows terminals to share high-speed trunks more efficiently than would be possible using circuit-switched channels or point-to-point private lines. This efficiency is reflected in the rarifying policy of public packet-switching networks, which typically charge by usage and not by distance.

The development of packet switching in the 1960s was intended to provide a fabric more in tune with the traffic patterns associated with computers rather than with those of telephones for which circuit-switched networks were engineered. As such packet-based solutions which now include LANs; the newer fast-packet technologies, such as SMDS, frame relay, and ATM; and old-time X.25, SNA, and IP networks' are obvious options for networking bursty BOD applications. Although some of the older packet services, notably X.25, are too slow to support many BOD applications, SMDS and frame relay are well suited to certain BOD applications particularly, LAN internetworking. However, the degree to which these services which are primarily intended for LAN internet-working can support real-time voice and video is still unclear at the time of writing. Clearly, voice and video support will be required for these services to give full support to BOD applications.

Similar problems also arise with LANs. Although some of the earliest conceptions of LANs were as shared bandwidth networks capable of handling voice, data, and video, and Datapoint marketed a video LAN several years ago, the overwhelming bulk of traffic that travels through LANs at the present time is data oriented. The existing LAN infrastructure is therefore capable of supporting data-oriented BOD applications, and as bandwidth requirements for such applications increase (notably because of the growth in image traffic), demand will largely be satisfied by upgrading from today's mainly Ethernet and token-ring networks to FDDL This transition will be aided by the advent of low-cost FDDI interfaces that will make FDDI-to-the-desk economically viable. Segmenting LANs with bridges and routers, thereby organizing traffic are also facilitating support of data-oriented BOD applications more efficiently.

There has also been a rapidly growing interest in running video over LANs, and commercial products with this capability are appearing from both major vendors and start-ups. However, all the evidence points to the fact that no matter what the bandwidth of the LAN, video support begins to run into trouble when many users are on the network, because at some point video quality begins to degenerate. This problem is inherent in the shared bandwidth technology common to all LANs and is not ultimately solved by shifting from Ethernet/token-ring to FDDI. For example, the best that commercial products seem to be able to offer is support for 40 users on FDDI. This may well be enough for many of today's implementations of video and multimedia applications, but it clearly could not support a ubiquitously distributed broadband application such as video/multimedia mail.

One way of handling real-time isochronously (voice and/or video) services over LANs is to graft a circuit-switched/TDM technology onto an existing LAN. This has been done in the case of the ANSI X3T9.5 FDDI-II standard and in the IEEE's 802.9 standard. However, FDDI-II has only a few supporters in the vendor community and 802.9, which is a LAN that incorporates ISDN channels, has as yet produced very little excitement.

4.2.3 The ATM Advantage

As we discussed briefly in the previous chapter, the core switching technology for broadband networks will be asynchronous transfer mode (ATM). As we also noted in this chapter,

ATM is an advanced standardized version of a switching/multiplexing concept often referred to as fast packet switching. This technology represents the adaptation of older packet-switching technology to modern network conditions. Fast packet -switching is not a new technology but rather a new class of packet-switching techniques that get over some important limitations (notably low data rates and lengthy packet delays) imposed in the older packet-switching standards. A typical conventional packet switching network using the CCITT's X.25 protocol can offer no more than 64 kbps, although 2 Mbps X.25 networks have been built. Speeds beyond this level would require massive processing power to deal with all the overhead that is part of the X.25 protocol.

The reason for the existence of this overhead is that at the time when packet switching was being developed, a considerable amount of error checking and flow control had to be performed within the network. A high level of error checking was required because the bit error rates (BERs) associated with the physical network were much higher than today. In addition, a high level of overhead processing had to be done within the network because, for the most part, the equipment deployed at the network termination was quite dumb.

Essentially, the fast packet concept proposes streamlining the error checking and flow control functions of conventional packet switching. This is turn reduces the protocol processing that needs to be done on the network. Improvements in network hardware mean that error recovery can be off-loaded from switches to terminals, because errors are assumed to be sufficiently infrequent that they can be left to terminals and high-level software and because terminals are now much more intelligent than when X.25 first evolved. In conventional packet-switching networks, error checking protects every node-to-node hop.

As the result of the adoption of the fast packet switching technology, the speeds of packet networks have been increased considerably and packet delays have been improved. The first of these trends is clearly important from the perspective of broadband. Improvements in handling delay are also critical, because such improvements make it possible to use packet-switching technology for real-time voice and video commendations. Less than 80 ms end-to-end delay is what is required for a comfortable voice conversation. Greater delay tends to result in participants in such a conversation cutting each other off as anyone who has talked on a satellite link knows. Conventional packet-switching networks offer typical end-to-end delay times of about 125 ms and so are not suitable for voice transmission.

In the past, running voice and video over a packet network has been carried out mostly in a research environment, in order to prove that it can be done. Typically, the quality of transmission associated with such demonstrations has been very poor With fast packet technologies and ATM specifically now rapidly being established in public and private networks, packet-based video and voice is of much higher quality.

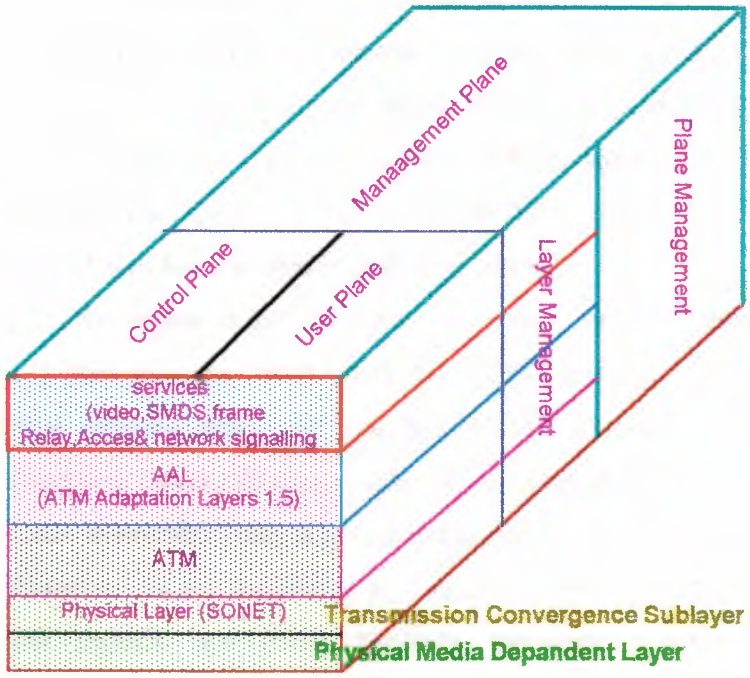

Figure 4.1:Broadband ISDN Protocol Model

Although, in practice, most of the fast packet technology currently being deployed is mainly being used for data applications, mostly internet-working LANs and host computers. But the capability of fast packet switching to handle real-time voice and video is a major factor making ATM the chosen switching/multiplexing fabric for Broadband ISDN (BISDN).

4.2.4 ATM and Broadband ISDN

Although ATM is now being used as a switching/multiplexing technology within enterprise networks, it was originally developed with B-ISDN in mind and forms a vital role in B-ISDN as the second layer in the B-ISDN protocol model. This model is shown in (Figure 4.1). At the bottom of the model is the B-ISDN PH (physical) layer, which performs the traditional physical layer functions such as mapping cells into transmission systems in the transmit direction or extracting the cell in the receive direction. In B-ISDN, this activity is dealt with by the SONET protocols. The ATM layer of the B-ISDN model, which sits immediately above the PH layer, manages the formatting of the cell for transmission and interpreting the contents of a valid cell on reception. The ATM adoption layer (AAL) which sits above the ATM layer in B-ISDN, is effectively a protocol conversion layer that adapts higher-layer protocols to conform to the ATM layer protocol.

For current purposes, however, the reader should bear in mind mostly where ATM functions in? The communications process, relative to other layers in the model. This is not significantly different, even where ATM is used outside of a formally defined B-ISDN, as is the case when ATM switches are used in enterprise networks.

In addition, the model helps to explain the rather peculiar terminology associated with ATM. In its standardization of B-ISDN, the CCITT has tried to use terminology that does not override existing usage for the International Standards Organization's Open Systems Interconnection (OSI) reference model. Thus, although ATM is a packet-switching technique, the basic information unit is not described as a "packet," but as a "cell." This is because in conventional packet switching, a "packet" is a network layer entity, whereas a "cell" is an OSI physical layer entity. Similarly, the term "transfer mode" is used to designate that the process takes place on the physical layer, rather than the more usual term "transport." In conventional telecommunications systems, the 0S1 Transport Layer carries out the switching/multiplexing functions.

The exact relationship between the 0S1-reference model and the B-ISDN protocol model is ill defined and is an issue that is not being widely discussed at the present time. It is, for the most part, an academic issue, and there is currently just too much to do of practical importance to develop broadband networks.

The CCITT chose ATM as the basis for B-ISDN because of its ability to handle the very wide variety of data streams that are envisioned as traveling over B-ISDN. This includes both variable and fixed-rate bit streams, low- and high-speed traffic, and traffic with differing "quality of service" (QOS) characteristics, such as delay, bit error rates, and so on. As the result of processes on the ATM and AAL layers, customer data of many different kinds are split up into fixed-length cells. ATM allows multiplexing of these cells from one of more customers.

The ATM cell size is 53 bytes, of which 48 bytes is given over to customer information (including higher-level protocol information). The remaining 5 bytes are given over to cell identification, message type identification, and error detection. Two types of ATM cell headers are specified in the B-ISDN standards, one for the network-to-network interface (NNI), the other for the user-to-network interface CUNI). These headers differ in how they handle flow control.

The choice of a 53-byte cell was essentially a compromise between the North American and European telecommunications communities, because this cell size can readily support the multiplexing of data streams operating at data rates in either the North American or the European digital hierarchies. Of more importance than the actual sue of the cell, however, is the fact that the cell is relatively small. This facilitates the delivery of delay-sensitive traffic such as voice and video traffic, which is a critical issue for B-ISDN, which must continue to support existing telephony services and for which the ability to transport high-quality video traffic is a raison d'être. Unfortunately, the small cell size associated with ATM also means that the customer data must be fragmented for transporting over a B-ISDN fabric. This requirement led to the need for some kind of group identity scheme that allows cells to be identified as belonging to a particular data stream. The scheme utilized in ATMIB-ISDN is based on a system of virtual channel identifiers (VCIs) and virtual path identifiers (VPJs).

Here is an example of how the system of identifiers works: a B-ISDN data stream may contain multiplexed data from (say) two 64 kbps voice channels, a 10 Mbps Ethernet LAN, and a 45 Mbps video channel. At the ATM layer, these three data streams will have been broken down into cells, each with their own unique set of identifiers. Thus, all the cells associated with the Ethernet LAN have the same set of identifiers. Cell identification is the dominant function associated with the cell header. In the NNI cell header all but 12 of the 40 bits are associated with identification. This is shown in (Figure 4.2).

As noted earlier, identifiers may relate to either "virtual channels or to virtual paths. These concepts can be illustrated by examples from the North American digital hierarchy, and this discussion may also serve to indicate the main differences between ATM and the synchronous time division multiplexed technology that accounts for most public networking today.

In TDM technology a DS1 signal has 24 DSO signals, distributed across 8000 frames. Each octet in each frame represents a segment of a DSO and each octet is identified by its position with respect to a frame bit.

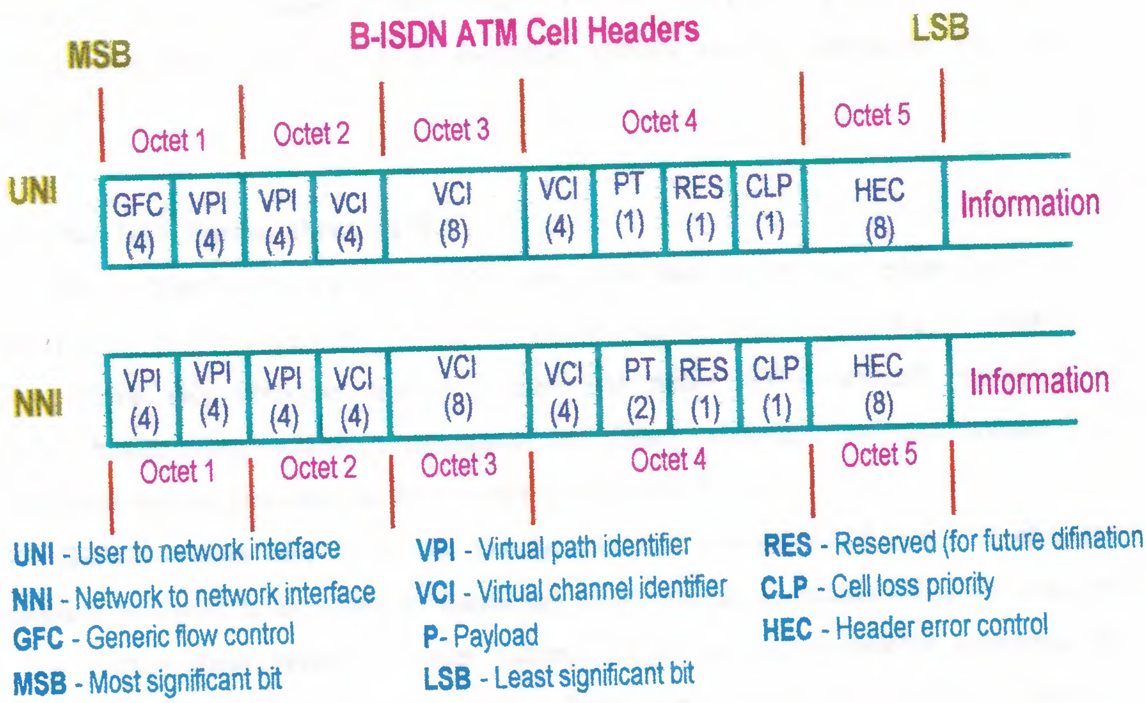

Figure 4.2 B-ISDN ATM Cell Headers

This is sometimes called "positional multiplexing." For example, a particular DSO may be fragmented in a way that its fragments are assigned to every tenth octet in a frame. Demultiplexing can consist of collecting every tenth octet and reassembling these octets to recreate the DSO. By defining the channel in terms of the frame bit, no reference is made specifically to the physical signal, so the DSO is referred to as a virtual channel. The positional relationship with respect to each frame bit in this case "every tenth octet is a packet for a particular DSO" is the implicit VCI.

Similarly, each DS3 signal may be thought of as a pipe through which there are 28 different paths, each represented by a DS1 channel. Each DS1 is position multiplexed into the DS3, in much the same way as the DSO is multiplexed into a DS1. Thus a DS1 may be thought of as a virtual path within a DS3, and there is a VPI defined in terms of the 4rame bit.

The concept of VPIs and VCJs is carried over into ATM, with the important difference that in ATM, there is no such thing as a frame or a frame bit. Instead, explicit VPIs and VCIs are contained in each header. It is this absence of a frame bit that makes ATM asynchronous. VPIs and VCIs essentially provide routing information for each B-ISDN cell.

4.2.5 What Is So Special about ATM?

The technical description of ATM just given fails utterly to explain just why so many people in the computer, LAN, and public carrier communities are so excited by ATM. There are, however, several factors that make ATM uniquely suited as a switching/multiplexing technology to support broadband networking and applications. It is these factors that are generating all the excitement about ATM.

The key reason why ATM will be the transport mechanism of choice in all-future broadband networks is its ability to handle all the numerous applications that are expected to flow across those networks. Older circuit-switched or packet-switched networks are usually engineered to deal efficiently with a fairly narrow range of applications, although with some tweaking they can usually handle a wider range. But, almost by definition, broadband networks must cope with a frighteningly large variety of networking circumstances. Future broadband networks will, for example, be required to:

- Switch HDTV traffic at OC-3 and above, yet be optimized to handle voice, fax, and low-speed data traffic.
- Handle calls that last as little as one second, as may be the case in some telemetry applications, as well as more or less continuous applications such as video surveillance.
- Handle highly bursty traffic (telemetry, again) as well as less bursty applications such as videoconferencing.

These examples illustrate the extreme requirements that will be required of broadband networks. To make matters even more complex, broadband networks must accommodate applications falling at almost every point between these extremes. This puts an even greater burden on broadband switching. After all, if broadband services fell neatly into a few performance categories, switching might most easily be accomplished by using different kinds of switches for different classes of applications, as is essentially the case today Future broadband networks will have to cope with a full spectrum of traffic characteristics.

Clearly classical circuit- or packet-switching techniques could never live up to the requirements just outlined, but because of the high data rates of which it is capable and because it supports both continuous and bursty data streams, ATM also provides a natural means for integrating services in common switching and transmission facilities. In particular:

- ATM can support, voice, data, and video services individually and also provides the opportunity to integrate multiple types of information into multimedia services.
- Although ATM remains essentially a packet-switching technique, inherently circuit-switched applications, the classic example of which is telephony, can be supported through circuit emulation.
- Although classical public networking based on time-division techniques is limited to providing connections at specific bit rates, ATM can provide a flexible mixture of connections up to the capacity of the broadband channel.

In addition to providing support for a wide variety of services, ATM also offers a high degree of flexibility when it comes to changing network architectures. This is important because it is widely expected that customers (especially large business customers) will need to change the architectures of their networks both fairly regularly and at relatively short notice. In addition, it is far from clear what architectures will be most suitable for broadband networking.

An ATM network is likely to be able to respond to the need for flexible architectures more economically than today's cross connect systems by establishing virtual trunk groups between two ATM switches. This minimizes the call processing at intermediate nodes. Unlike trunk groups established through the current generation of cross connects, an individual trunk on a virtual path consumes transmission bandwidth only when actually carrying traffic. The savings that can result from this are obvious. In some views of how the ATM network will develop, ATM cross connects will play an important early role.

Yet another reason why networkers are attracted to ATM, although perhaps a relatively minor one, is the ease with which ATM can handle network synchronization. It can be expected that broadband networks will develop as a series of "islands," each with their own network clock. Even when broadband communications becomes ubiquitous (sometime during the 21st century), there will still not be a single B-ISDN, but rather a collection of separate, Internet worked B-ISDNs, each operated by a separate carrier. In some countries, such as the United States and the United Kingdom, where there are multiple public carriers, this diversity will appear within the country concerned, but even where basic telephony is maintained as a monopoly, there will be a need to interconnect a national network with other national networks.

Each of the B-ISDN islands, national networks, or separate carrier networks will employ different network clocks, a fact that will require some coordination in order to provide internetworking between these separate networks. ATM switches are especially suited to responding to this need for coordination because, since ATM is not based on the establishment of fixed-rate connections, ATM can provide more network synchronization-circuits that are not synchronized to the network clock can be carried over an ATM circuit. In addition to getting around the difficulties associated with internetworking networks with uncoordinated network clocks, this characteristic of ATM networks facilitates the mixing of asynchronous and synchronous traffic in the same network.

4.2.6 The Downside of ATM

But there is a "flip side" to ATM too. The main reason why ATM is so suitable for broadband applications is because of the advanced technology that it incorporates, and the adoption of such technology invariably means that some technical problems remain unresolved.

The adoption of ATM switching in public telephone networks and in private corporate networks is revolutionary rather than revolutionary. It is a true "paradigm shift" because ultimately it means a wholesale abandonment of circuit switching for packet switching in the public networks and LANs on customer premises, and the problems that may arise in this area will probably not be fully known until ATM is well into its~ trial phase.

One problem frequently mentioned in the literature with regard to the handling of isochronous services by ATM networks is the problem of delay, an issue that we met earlier. In ATM networks, the average delay in passing the cell through the network is extremely small-in fact, smaller than for conventional time-division techniques. However, the creation of the cell itself can lead to significant delays. For example, in order to create an ATM cell payload of 48 bytes, 64 PCM speech samples must be taken, and this results in a total delay of 16 ins, which is too great a delay to be tolerated in a public network.

Several techniques have now been developed to deal with the problem of long delays in ATM networks, but delay time is not the only delay-related issue that can create problems when sending voice over ATM networks. The variation in delay time is also important. Such variations, which are typical in packet-switched networks, have no significance in data networks but can seriously disrupt real-time voice or video communications. Again, there are techniques for dealing with the variation in delay time, but this problem is one that does not occur in older networks and represents a design problems for both switch engineers and network engineers.

Although delay is a problem that confronts mainly isochroous transmission in ATM, some kind of congestion control is important in order to accommodate efficient data communications. This is because in networks with particularly bursty traffic, switch buffers may overflow and cells will be lost. This problem can be particularly acute for data services because buffer overflows lead to data being retransmitted, and in this way even more data is lost.

Congestion control in ATM networks has been the topic of many technical papers in the past few years, and a variety of approaches have been suggested. It also seems to be increasingly acknowledged that congestion will not be the problem in real-world ATM networks that some observers once thought it would be. Nevertheless, congestion can still be a problem when a high-resolution video application is up and running with multiple workstations accessing the application.

One critical factor retarding the growth of the ATM market is the huge cost of developing and deploying new central office switches. These costs may be particularly high in the case of ATM switches because of their revolutionary nature. The high cost and low profitability of manufacturing central office switching gear has forced a number of companies out of this business in the past decade. And even when switches are built, there are often significant barriers to entry in the form of old and unbreakable ties between carriers and equipment suppliers, although the ability of switch vendors to come to market first with equipment embodying a new technology is a proven way to break through such barriers. Being among the first to offer ATM switches has already allowed foreign telecommunications companies (Northern Telecom and Fujitsu, among others) to make their mark on the U.S. central office switch market.

Although many vendors now have well-developed schedules for bringing ATM switches to market, the cost factor seems to be shaping the market in that few vendors have plans for very large switches, even for public network deployment. Rather, ATM switching modules are being prepared as adjunct switches for the main central office switches, which continue to deal with conventional traffic. Broadband traffic is routed to the ATM modules. Other vendors are offering small ATM switches for use on customer premises.

This strategy minimizes the costs associated with the development of ATM products and the deployment of ATM switches at a time when the market for broadband services is far from proven. Generally speaking, ATM switches are also scalable in the sense that switches can be increased in capacity through the addition of extra cards, rather than the replacement of the whole switch.

Yet another constraint on the design and manufacture of ATM switches is that large ATM switches and those with very high throughput strain the capabilities of today's microelectronics. In order to operate at SONET rates, ATM switches require very fast processing rates.

TABLE 4.2 CELL RELLAY SPEEDS REQUIRED FOR SONET TRANSMISSION

SONET Rate	Data Rates	Thousands of Cells/Second
OC -1	52 Mbps	135
OC -3	155 Mbps	405
OC -12	620 Mbps	1600
OC -24	1.2 Gbps	3200

This is not just because of the high data rates involved, but also because ATM cells are quite small. For example, even to operate at OC-1 (approximately 52 Mbps), which is the lowest point in the SONET hierarchy, cells will have to be shifted at a rate of 135,000 cells per second. For OC-24, the rate would be 3.2 million cells per second. When one compares conventional packet switches, which offer around 100,000 packets per second, it is easy to see that chip level developments are a key enabler in the penetration of ATM. Also, it is not just a matter of processor speed, but bus bandwidth and memory speed as well. (Table 4.2) shows the ATM cell relay speeds required for selected levels of the SONET hierarchy

5. BROADBAND STANDARDS

5.1 Summary

This chapter surveys the main broadband networking standards, the organizations that make them, and the standards-making process itself. We begin with a look at the work done by the ANSI Ti and X3 committees. Ti has produced important standards for broadband networking including work on SONET, ATM, and Broadband ISDN itself. X3 is devoted to work on computer networking, channels, and interfaces, and has standardized various high-speed technologies falling into this category including FDDI, HIPPI, and Fibre Channel.

Having examined the work of ANSI in the broadband area, we examine its international counterparts the International Telecommunications Union and the International Standards Organization. We then go on to look at the broadband standards-making endeavors of the Institute of Electrical and Electronic Engineers (IEEE), Bellcore, the broadband forums administered by the Interop organization, and the Internet Engineering Task Force.

Broadband standards organizations, which overlap in both purpose and membership. At the international level, the main standards organization concerned with broadband standards is now the International Telecommunications Union's Standards Sector, which has the American National Standards Institute (ANSI). However, as we discuss later, only certain groups within these organizations are pursuing the broadband endeavor. In addition, a variety of other organizations, although not officially designated as standards groups, effectively act as standards groups. The organizations include the Bell companies' central service organization, Bellcore, the IEEE, the various broadband forums associated with the Interop company and the Internet Engineering Task Force (IETF).

We therefore focus on the development of broadband standards within ANSI. However, the reader should note the CCITT has played a leading role in the development of broadband standards, most notably in the area of Broadband ISDN, a set of standards, which it largely initiated. With regard to the development of public network standards, overlapping memberships and formal contact between the relevant ANSI and CCITT committees help to smooth out the differences between broadband standards making at the international and national level. As a consequence, there are fewer differences between international and national telecommunications standards than there used to be. One important change has been that ANSI standards are no longer published in advance of the equivalent international standards has been agreed upon completely by interested parties in the United States.

5.1.1 ANSI

ANSI approves standards for virtually every kind of product used in American homes and businesses. Obviously, no one organization could carry out that task unaided, and the way ANSI works is to delegate each area to a responsible organization, which it then accredits for standards-making purposes. In the telecommunications industry, the organization that is accredited for overseeing the task of setting standards for public telecommunications in the United States is the Exchange Carriers Standards Association (ECSA). ECSA has replaced AT&T as the main standards setter in this area since the breakup of the Bell System and operates its standards-making activities through the ANSI Ti committee.' Another ANSI –committee-the X3 Committee is also important, as the developer of local broadband networking standards. This group is administered by the computer and Business Equipment Manufacturers Association (CBEMA).

Finally, the Institute of Electrical and Electronic Engineers (IEEE) operates an ANSI-accredited effort to standardize LANa operating under 20 Mbps through its 802 Project.

5.1.2 ANSI TI

The ANSI Ti committee is divided into several technical committees dealing with different matters. These technical committees are further divided into subcommittees. Broadband issues are dealt with. by a number of groups within Ti, but are focused on the T1S1 technical committee, which deals with services, architecture, and signalling. Within the T1S1 committee, the T1S1.5 group specializes in matters to do with Broadband ISDN. Other groups within Ti are also important to the broadband endeavor. TIA1 is concerned with performance issues in broadband and narrowband networks. T1E1 is concerned with user—network interfaces and addresses some of the issues surrounding the physical layer interface for Broadband ISDN networks. T1M1 is the TI group that deals with network management issues (or operations, administration, and maintenance as it is called in the public networks). Finally, T1X1 is the key group concerned with the lower-layer SONET transmission fabric for Broadband ISDN. It is this fabric that we will address first.

5.1.3 SONET

One of the essentials for the development of broadband networks is a standardized hierarchy of digital transmission rates, formats, and interfaces for transmission in the multi megabit to gigabit range. This is pro-vided by the SONET

TABLE 5.1 The Sonet Hierarchy

SONET Rate	64K Channel Capacity	Data Rate
OC-1	672	51.84 Mbps
OC-3	2.016	155.52 Mbps
OC-9	6.048	466.56 Mbps
OC-12	8.064	622.08 Mbps
OC-18	12.096	966.12 Mbps
OC-24	116.128	1.244 M bps
OC-36	24.192	1.866 M bps
OC-48	32.256	2.488 Mbps
OC-96	64.512	4.974 Mbps
OC-192	129.024	9.952 Mbps

Hierarchy. SONET was originally designed by Bell-core to supplement the old Bell System "T" hierarchy, which only extended to T3 rates (i.e., 45 Mbps).

SONET has become an international standard in the form of the CCITT's synchronous digital hierarchy (SDH), which differs in a minor ways from SONET, but which 5 fully compatible with SONET.

The current SONET/SDH hierarchy is profiled in (Table 5.1). The top five rates in the SONET hierarchy are gigabit rates, and SONET is potentially extendible up to at least 13 Mbps.

The basic information unit in SONET is a 90-by 9-byte frame the STS1 frame (Figure 5.1). In the STS-1 signal, one of these frames can be transmitted every 125 microseconds. From this it follows that the basic STS1 signal operates at 1,000,0001125 = 8,000 frames per second or 8,000 x 8 x 90 x 9/1,000 = 5L8400 Mbps. The rest of the SONET hierarchy is essentially derived by multiplexing together these basic SONET frames. In B-ISDN, ATM cells are carried in these frames.

The SONET hierarchy allows network owners, especially carries the community from different vendors and have them work together. SONET also has a neat feature called

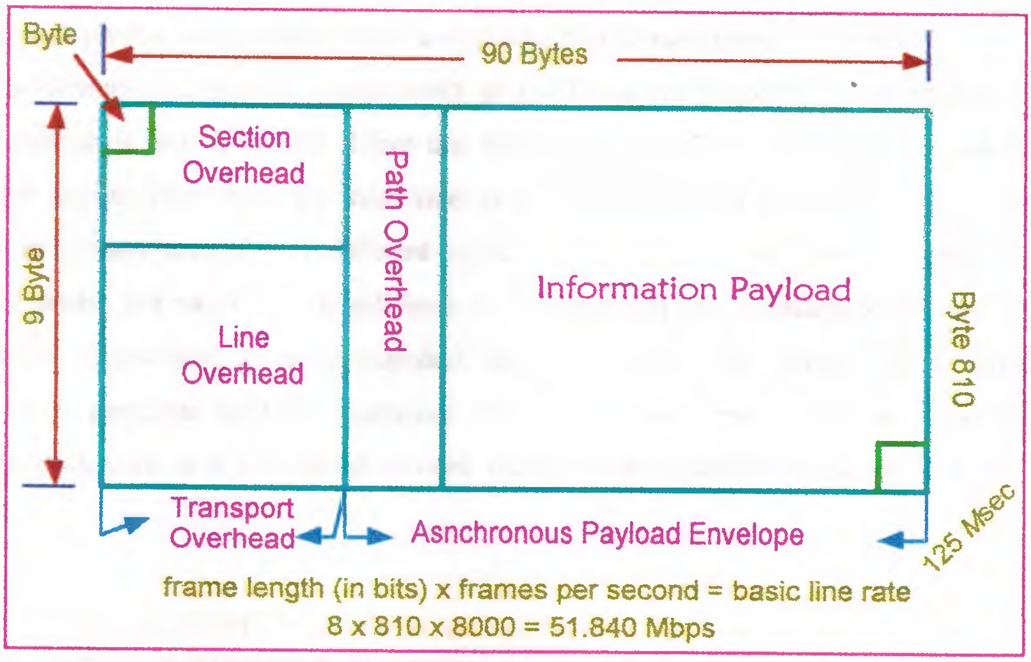

Figure 5.1 The Sonet STS -1 Frame

add-drop multiplexing which makes it possible to strip off a low-speed circuit from a SONET pipe without demultiplexing all the circuits in the pipe. This feature can be a major cost saver, eliminating the need for back-to-multiplexers in many applications.

Despite the advantages that SONET apparently offers the deployment of SONET gear by U.S. public carriers has been lower than originally predicted, and private networkers have all but ignored SONET. SONET equipment sales by companies such as AT&T and Northern Telecom have been disappointing, and the emphasis has been on SONET gear operating at the lower points in the hierarchy. The main factors that have held back the deployment of SONET are:

(1) The need for intelligent operational Support systems (OSSs) to be deployed to support the management of SONET network elements

(2) Outstanding SONET standards issues.

As far as the first point is concerned, the problem is to embody enough intelligence in the network so that it can manage SONET based network elements. What is required here varies among network types. In particular, the Bell company networks are particularly well endowed with embedded OSSs. But these were originally intended to work with older dumber network elements. The Bells have made considerable progress in upgrading these older OSSs to provide configuration, fault, and performance management for SONET.

The network management environment at public carries outside of what used to be the Bell System is less developed. There tire fewer existing OSSs to be upgraded, and this gives these carriers more room for innovation in the management of their SONET gear. The situation in private networks is different again. SONET is not yet widely deployed in private networks, but when it is, it will have to conform with the prevailing private network management environment a vendor-supplied network management system that is tightly coupled to a particular kind of equipment but at the same time serves as a network management element in a centralized network management architecture such as IBM's Net View.

TABLE 5.2 ANSI T1 SONET STANDARDS

ANSI T1.105-1988	Optical Interface Rates & Formats
ANSI Ti. 106-1988	Optical Interface Specifications: Single-mode (interoffice
ANSI T1.105-1990a	Addendum
ANSI T1.105-1991	Optical interface specifications, changes to Phase 1 timing and synchronization, short reach optical parameters, clarification of overhead usage, DCC protocol details, FDDI and ATM mappings, automatic protection switching details, STS- 1 and STS-3 electrical specifications

Table 5.2 shows the SONET standards developed to date by the Ti committee. Originally,

the standards makers on Ti divided their task into three phases. This approach has proved useful. In particular, the completion of Phase 1 has allowed chip makers to carve the basics of SONET into silicon, and this has ultimately led to the carriers being able to deploy basic SONET equipment now. More sophisticated functionality has arrived (and will arrive) mostly in the form of software upgrades.

SONET Phase 1 was embodied in two standards released in 1988 that specified optical parameters for interoffice connections (ANSI T1.106-1988) and the basic rates and formats of the SONET signal (ANSI T1.105—1988). With these standards in place, the basic SONET hardware could be constructed. SONET Phase 2 included the development of mappings from signals not discussed in Phase 1. For example, mappings from FDDI and ATM signals into a OC-3 SONET channel were defined. Phase 2 also covered special standards intended to facilitate the interconnection of SONET equipment located at a single central office. Finally, Phase 2 began the work of defining operations, administration, maintenance, and provisioning (OAM&P) for SONET networks. OAM&P functionality for multipoint SONET networks was left to SONET Phase 3, but OAM&P functionality essentially what most people would recognize as network management proved extraordinarily hard to settle.

There are at least three reasons for this. First, there have been delays in defining a generic model for public carrier network management at T1M1, and this has slowed up the work specifically aimed at SONET.

Second, the CCI'19? has continued to work on management functionality for SDH, and the ongoing coordination of this work and the work of SONET OAM&P has been a drag on the pace of standardization in this area. Third, Ti has decided to base its network management vision on an object-oriented approach. This decision will almost certainly be appreciated by future generations of network managers and users, but at present object-oriented software development is a new discipline. Finding people who know a lot about both SONET and object-oriented programming is not an easy task.

The first of these three issues is probably the major drag on the deployment of SONET. The key-missing element is the Common Management Information Service Element (CMISE). This is a language that is used to communicate operations support information between SONET network elements and OSSs. CMISE is not itself a SONET standard, but the SONET standards makers have specified CMISE as a means of conforming with the practice that will be followed throughout the public network. However, the CMISE standards will not be finalized until 1994 at the earliest, which means that complete SONET functionality will not be available in the public network until 1995 or 1996.

For current deployment of SONET, an alternative language Bell-core's Transaction Language 1 (TL1) is being used. This is an older language, dating back to the mid-1980s, but it is sufficient to fill the gap until CMISE is ready. Currently, most OSS vendors are supporting SONET through a three-layer protocol derived from the old X.25 standard with TL1 as the message language. This will be the standard approach until CMISE is ready, but it fails to provide any end-to-end services, such as addressing, routing, or message-delivery assurance. This functionality must be added by nonstandard methods.

Nevertheless, much of the groundwork has been laid for establishing SONET standards in the network, and escalating SONET deployment can be expected through the niid-1990s. Where SONET is currently being deployed, the use of SONET rings, which provide enhanced network reliability, has become something of a de facto standard for deployment. T1X1 has been developing an official standard for switched SONET rings, and Bellcore has also published in this area. In addition, T1XI has considered a project that would investigate the relationship of SONET and network architecture in a more general sense.

5.2 Broadband ISBN

SONET forms an integral part of the Broadband ISDN model but has evolved and has been standardized somewhat independently of the rest B-ISDN. In addition to SONET, the other key standard technology in B-ISDN is ATM.

ATM was adopted as the target solution for B-ISDN in the late 1980s by both Ti and the CCITT. The main reasons for its adoption were that ATM:

- Allowed the integration of voice, video and data services
- Provided "bandwidth on demand"
- Seems likely to be able to support future services of many kinds
- Supported the switching/multiplexing of very-high-bandwidth channels

ANSI Ti standards relating to B-ISDN are summarized in (Table 5.3). Included in this work, ANSI Ti has developed standards for the ATM adaptation layer (AAL),

TABLE 5-3 ANSI T1ATM/B-ISDN Related Standards

Ti.105	Synchronous Optical Network
TiE1.2/92-020	UNI PMD Specifications
TiS1/92-185	UNI Rates and Formats
TiS1.5/92-002	ATM Layer Specifications
TiS1.5/92-007	Generic Flow Control Baseline Text
T1S1.5/92-XXX	AAL Architecture for Class CR) and Signalling
T1S1.5/92-003	AAL 3/4 Common Part
TiS1.5/92-OiO	AAL5 Common Part
T1S1.5/92-OO8	Variable Bit Rate AAL Service Specific Part
TiS1.5/92-004	AAL for Constant Bit Rate (Class A) Services
TiS1.5/92-029	OAM Aspects (Technical Report)
TiS1.5/92-005	Support of Connectionless Services
TiS1.5/92-009	Traffic and Resource Management
T1S1.5/92-006	Services Baseline Document

Which was discussed briefly in Chapter 2. As mentioned in that chapter, the AAL might be thought of as a protocol conversion process that acts an intermediary between the user's information flow and the ATM process. Its goal is to preserve the general characteristics of the user's information flow while taking account of the underlying ATM technology.

Accounts of ATM in the trade press sometimes tend to ignore the AAL, which makes

ATM look better than it is. As long as we ignore AAL, ATM has the appearance of a unified technology/standard that can work in the local, metropolitan, and wide-area environment and can handle almost any kind of data. Once the AAL is brought into the picture, however, this unified picture of ATM is broken.

This is because there are five "flavors" of AAL, each designed to work with a different kind of data stream. Each type of AAL can handle a wide variety of service types, but each is meant to address specific service requirements. Type 1 AAL is intended for circuit emulation (essentially voice) and constant bit rate video. Type 2 AAL is associated with variable rate video. However, Type 2 AAL is not yet well defined and little or no research work is being concentrated on this area.

This type of AAL has been defined and is widely implemented by switch vendors planning to introduce products for the public network. Vendors offering ATM equipment for private networks have almost invariably implemented AAL Type 5, which was developed in 1992, originally for carrying signalling and certain data applications. The relatively simple structure of AAL 5 makes it suitable for enterprise networks, which, though they may be quite complex, do not have the complexity of large public network.

The mention of signalling takes us to another area in which Ti is involved in important standardization work. This is the area of signalling for Broadband ISDN. For public networks at least, this is both a vital and a difficult topic. It is vital if broadband networks are ever to live up to their promise of delivering a multiplicity of service types. Without signalling, the user cannot tell the network what she requires of it. But signalling is a difficult area of standardization, because broadband signalling must support entirely new features of future broadband networks such as management of virtual channel connections and specification of ATM parameters while maintaining a consistency with the old ways of doing signalling in the narrowband network.

Faced with such difficulties, the standards makers on T1 have chosen a phased approach to standardizing broadband signalling. In the first phase, existing narrowband standards will be extended to suit the requirements of Broadband ISDN. In particular, T1 is taking the CCITTs Q.931 protocol, which is used in ISDN to specify the user network interface (UNI) signalling and is extending it to provide an initial set of user broadband call control capabilities. This extension is known as the Q.93B Recommendation. It is also extending the "ISDN User Part" (ISUP) of Signalling System No.7 for setting up trunks and controlling "bearer" connections. Signalling System No. 7 is the latest generation of signalling system used in the public telephone network. Its most important characteristic is that it is capable of functioning as the "nervous system" of an intelligent network. Bearer services are essentially raw bandwidth services, with little or no value-added content

Phase 2 of the Ti Committee's plans for signalling in the broadband network involves establishing a new call model and protocol architecture that will support multimedia and multiparty services and establishing a new signalling protocol to support sophisticated call control. Most of the work for Phase 1 of the Broadband ISDN signalling endeavor is completed, but it could be well into the mid-1990s before we see practical deployment of the standards resulting from Phase 2.

CONCLUSION

The future of B-ISDN relies mainly on the evolution of fiber optics and on the demand of these services from the consumers. Because of this, no one can be sure when the potential of Broadband-ISDN will be realized.

For now we can only imagine a day when all of our communication tools and home entertainment devices will be merged into one, transmitted through one line, into our homes and offices. Wouldn't it be great if television programs and movies could be stored as files on our computers, and we will be able to do things like edit and email these files to our friends?

Broadband ISDN differs internally from a Narrowband ISDN in a number of ways that reflects the technological migration from narrowband ISDN to fast packet switching and then to ATM broadband ISDN.ATM, which is being accepted as the fundamental technology for future development of B-ISDN.

When ISDN is referred to as a network it is to be considered a telephone network, not a computer network. Broadband ISDN allows its users to communicate over high speed, high quality digital channels. The media is supports include Telex, fax, voice telephone, video telephone, audio, high definition TV and computer networking.

Video telephony is the transfer of voice, moving pictures and scanned images and documents between two points. Areas utilizing such technology are sales, consulting, teaching and legal services. The problem with gaining widespread use of video telephony is the prohibitive costs of terminal equipment. In the future as demand and competition increases the cost of such equipment will fall, and the service will become more widespread.

Most of the applications for ISDN have reached the extent of their development, and now the focus has shifted to services that an be provided across broadband ISDN cables. The ITU-T defines the services and associated standards of ISDN communications, have recommended the two-service area for application with BISDN, Interactive Services, and Distribution Services.

References

[1] Raier Handel, Integrated Broadband Networks, MANFAREDN N. HABUR.

[2] Lawrence Gasman, Broadband Networking,1996

[3] www.ouc.bc.ca.cs.uioa.edu

[4] www.layrinth.net.edu

[5] www.teltone.com