Internetworking LANs and WANs: Concepts, Techniques and Methods.
Second Edition. Gilbert Held

Print ISBN 0-471-97514-1 Online ISBN 0-470-84155-9

2

LOCAL AREA NETWORKS

In the first chapter of this book we became acquainted with local and wide area networks, comparing the general characteristics of each. In this chapter we will focus our attention upon the technological characteristics of local area networks, examining in detail several types of LANs. In doing so we will examine a portion of the IEEE 802 series of standards applicable to Ethernet and Token-Ring local area networks as well as the operation of ATM as a local area network.

2.1 TECHNOLOGICAL CHARACTERISTICS

Although a local area network is a limited distance transmission system, the variety of options available for constructing such networks is anything but limited. Many of the options available for the construction of local area networks are based upon the technological characteristics which govern their operation. Those characteristics include different topologies, signaling methods, transmission media, access methods used to transmit data on the network and the hardware and software required to make the network operate. In this section we will primarily examine the topologies, signaling methods, transmission media and access methods used to transmit data on local area networks, deferring a discussion of most specific hardware and software to future sections in this chapter and later chapters in this book where discussion of those topics is more appropriate.

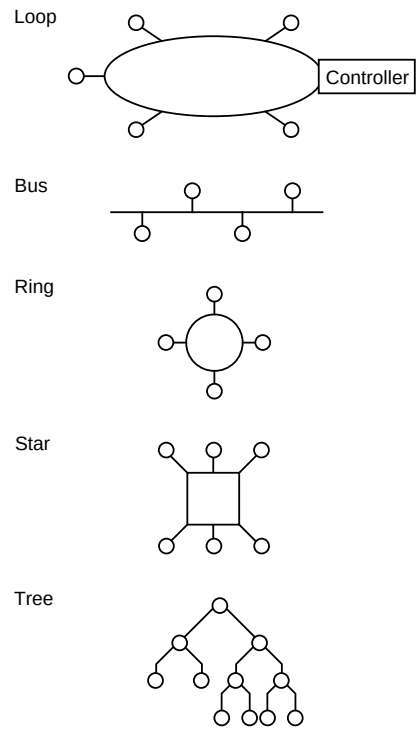

Figure 2.1 LAN topology. The five most common geometric layouts of local area network cabling form a loop, bus, ring, star, or tree structure

2.1.1 Topology

The topology of a local area network means the structure or geometric layout of the cable used to interconnect work stations on the network. Unlike conventional data communications networks that can be configured in a variety of ways by the addition of hardware and software, most local area networks are designed to operate based upon the interconnection of work stations that follow a specific topology. The most common topologies used in LANs include the loop, bus, ring, star, and tree as illustrated in Figure 2.1.

Loop

As mentioned in Chapter 1, IBM introduced a series of transaction processing terminals in 1974 that communicated

through the use of a common controller on a cable formed into a loop. This type of topology is illustrated at the top of Figure 2.1.

The controller employed a poll and select access method. That is, the controller would send a poll message to each terminal device in a predefined order. If the terminal had data to transmit it would wait until polled and upon being polled, transmit. Thus, terminal devices connected to the loop required a minimum of intelligence. Although this reduced the cost of terminals connected to the loop, the controller lacked enough intelligence to distribute the data flow evenly among terminals. A lengthy exchange between two terminal devices or between the controller and a terminal would thus tend to bog down this type of network structure. A second problem associated with this network structure was the centralized placement of network control in the controller. If the controller failed, the entire network would become inoperative. Due to these problems, the use of loop systems is restricted to several niche areas and is essentially considered as a derivative of a local area network. As such, we will eliminate this type of network from further consideration.

Bus

In a bus topology structure, a cable is usually laid out as one long branch onto which branches are used to interconnect each station on the network to the main data highway. Although this type of structure permits any station on the network to talk to another station, rules are required to govern the action necessary to recover from situations such as when two stations attempt to communicate at the same time. Later in this chapter, we will examine the relationship between the network topology, the method employed to access the network and the transmission medium employed in building the network.

Ring

In a ring topology, a single cable that forms the main data highway is shaped into a ring. Similar to the bus topology, branches are used to interconnect stations to one another via the ring. A ring topology can thus be considered to be a looped bus. Typically, the access method employed in a ring topology

requires data to circulate around the ring, with a special set of rules governing when each station connected to the network can transmit data.

Both bus and ring topologies use access methods that enable only one device to transmit at a time. As this results in network devices sharing the media, another term commonly used to refer to bus and ring-based local area networks is 'shared media networks'.

Star

The fourth major local area network topology is the star structure illustrated in the lower portion of Figure 2.1. In a star network, each station on the network is connected to a network controller. Then, access from any one station on the network to another station can be accomplished through the network controller. Here the network controller can be viewed as functioning similarly to a telephone switchboard, since access from one station to another station on the network can occur only through the central device.

The most common form of star-based local area network involves the use of a LAN switch. LANs developed using bus and ring topologies are now supported by the use of applicable LAN switches. When this occurs, it becomes possible to initiate multiple client-server sessions if two or more servers are connected to a switch, alleviating the major constraint of shared media networks in which only one source to destination data flow can be supported at a time. In addition, as we will note later in this chapter, ATM is based upon the use of switches.

Tree

A tree network structure can be considered to represent a complex bus. In this topology the common point of communications at the top of the structure is known as the headend. From that location feeder cables radiate outward to nodes, which in turn provide workstations with access to the network or provide a feeder cable route to additional nodes from which workstations gain access to the network.

Mixed topologies

Some networks, from a topology perspective, are a mixture of topologies. For example, as previously discussed, a tree structure can be considered as a series of interconnected buses. Another example of the mixture of topologies is the IBM Token-Ring Network. That network can actually be considered to be a 'star–ring' topology, since up to eight workstations and servers are first connected to a common device known as a multistation access unit or MAU, which in turn is connected in a ring topology to other MAUs. Later in this chapter and in Chapter 4 we will examine the IBM Token-Ring Network in detail.

Comparison of topologies

Although there is a close relationship between the topology of the network, its transmission media, and the method used to access the net, we can examine topology as a separate entity and make several generalized observations. First, in a star network the failure of the network controller will render the entire network inoperative. This results from the fact that all data flow on the network must pass through the network controller. On the positive side, the star topology is normally in existence within most buildings in the form of telephone wires that are routed to a switchboard. This means that a local area network that can use in-place twisted-pair telephone wires is normally simple to implement and usually very economical.

In a ring network, the failure of any node connected to the ring normally inhibits data flow around the ring. Due to the fact that data travel in a circular path on a ring network, any cable break has the same effect as the failure of the network controller in a star-structured network. Since each network station is connected to the next network station, it is usually easier to install the cable for a ring network. In comparison, if existing telephone wires are not available you would have to cable each station in a star network to the network controller, which could result in the installation of very long cable runs.

In a bus-structured network, data are normally transmitted from one station to all stations located on the network, with a destination address appended to each transmitted data block. As part of the access protocol only the station with the

destination address in the transmitted data block will respond to the data. This transmission concept means that a break in the bus may affect only network stations on one side of the break which wish to communicate with stations on the other side of the break. Thus, unless a network station functioning as the primary network storage device becomes inoperative, a failure in a bus-structured network is usually less serious than a failure in a ring network.

A tree-structured network is similar to a star-structured network in that all signals flow through a common point. In the tree-structured network the common signal point is the headend. In addition to the failure of the headend rendering the network inoperative, this network structure requires the transmission of information between some workstations to traverse relatively long distances. For example, communications between two workstations at the most distant end of the network would require a signal to propagate twice the length of the longest network segment. Due to the propagation delay associated with the transmission of any signal, the use of a tree structure may result in a degree of response time delay when transmission occurs between two workstations located at the most distant node or pair of nodes from the headend.

2.1.2 Signaling methods

The signaling method used by a local area network references both how data are encoded for transmission and the use of the frequency spectrum of the media. To a large degree the signaling method is related to the use of the frequency spectrum of the media.

Broadband versus baseband

Two signaling methods used by LANs are broadband and baseband. In broadband signaling the bandwidth of the transmission medium is subdivided by frequency to form two or more subchannels, with each subchannel permitting data transfer to occur independently of data transfer on another subchannel. In baseband signaling only one signal is transmitted on the medium at any point in time.

In comparison to baseband signaling, broadband is more complex. Broadband signaling requires information to be

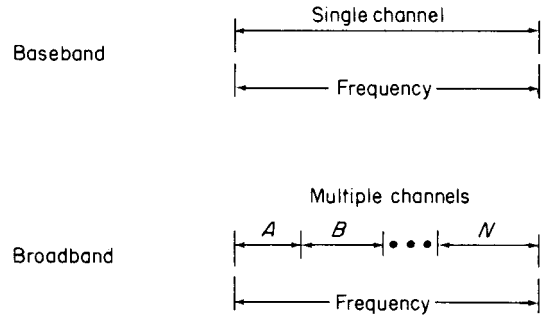

Figure 2.2 Baseband versus broadband signaling. In baseband signaling the entire frequency bandwidth is used for one channel. In comparison, in broadband signaling the channel is subdivided by frequency into many subchannels.

transmitted via the modulation of a carrier signal, requiring the use of special types of modems discussed later in this section.

Figure 2.2 illustrates the difference between baseband and broadband signaling with respect to channel capacity. It should be noted that although a twisted-pair wire system can be used to transmit both voices and data, the data transmission is baseband since only one channel is normally used for data. In comparison, a broadband system on coaxial cable can be designed to carry voice and several subchannels of data as well as facsimile and video transmission.

Broadband signaling

A broadband local area network uses analog technology in which high frequency (HF) modems operating at or above 4 kHz place carrier signals onto the transmission medium. The carrier signals are then modified, a process known as modulation, which impresses information onto the carrier. Other modems connected to a broadband LAN reconvert the analog signal back into its original digital format, a process known as demodulation.

Figure 2.3 illustrates the three primary methods of data encoding used by broadband analog systems: amplitude, frequency and phase modulation. The most common modulation method used in broadband LANs is frequency shift keying (FSK), in which two different frequencies are used, one to represent a binary 'one' and another frequency to represent a binary 'zero'.

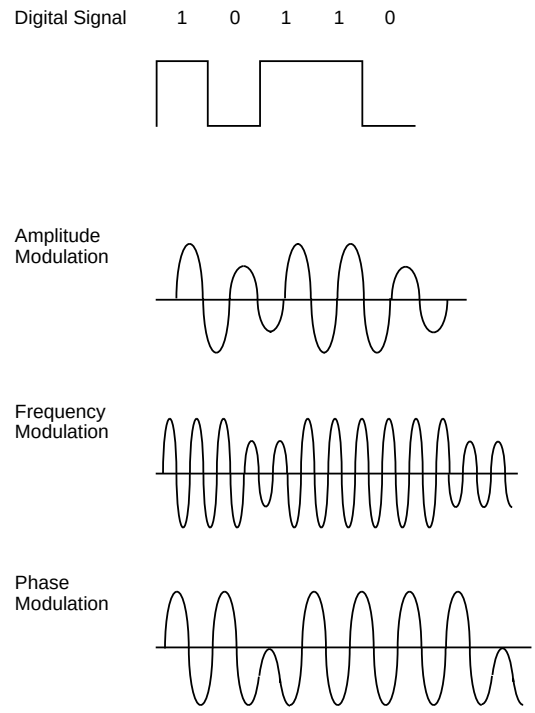

Figure 2.3 Modulation methods. Baseband signaling uses amplitude, frequency or phase modulation, or a combination of modulation techniques to represent digital information

Another popular modulation method uses a combination of amplitude and phase shift changes to represent pairs of bits. Referred to as amplitude modulation phase shift keying (AM PSK), this method of analog signaling is also known as duobinary signaling as each analog signal represents a pair of digital bits.

As it is not economically feasible to design amplifiers that boost signal strength to operate in both directions, broadband LANs are unidirectional. To provide a bidirectional information transfer capability a broadband LAN will use one channel for inbound traffic and another channel for outbound traffic. These channels can be derived by frequency or obtained from the use of a dual cable.

Baseband signaling

In comparison to broadband local area networks that use analog signaling, baseband LANs use digital signaling to convey information.

To understand the digital signaling methods used by baseband LANs let us first review the method of digital signaling used by computers and terminal devices. In that signaling method a positive voltage is used to represent a binary one, while the absence of voltage (zero volts) is used to represent a binary zero. If two successive one bits occur, two successive bit positions then have a similar positive voltage level or a similar zero voltage level. Since the signal goes from zero to some positive voltage and does not return to zero between successive binary ones, it is referred to as a unipolar non-return to zero signal (NRZ). This signaling technique is illustrated at the top of Figure 2.4.

Although unipolar non-return to zero signaling is easy to implement, its use for transmission has several disadvantages. One of the major disadvantages associated with this signaling method involves determining where one bit ends and another begins. To overcome this problem would require synchronization between a transmitter and receiver by the use of clocking circuitry, which can be relatively expensive.

To overcome the need for clocking, two popular types of baseband LANs, Ethernet and Token-Ring, use Manchester and Differential Manchester encoding, respectively. In Manchester encoding a timing transition always occurs in the middle of each bit while an equal amount of positive and negative voltage is used to represent each bit. This coding technique provides a good timing signal for clock recovery from received data due to its timing transitions. In addition, since the Manchester code always maintains an equal amount of positive and negative voltage, it prevents direct current (DC) voltage buildup, enabling repeaters to be spaced further apart from one another.

The lower portion of Figure 2.4 illustrates an example of Manchester coding. Note that a low to high voltage transition represents a binary one, while a high to low voltage transition represents a binary zero. Under Differential Manchester encoding the voltage transition is used only to provide clocking. The encoding of a binary zero or one is represented by the presence or absence of a transition at the beginning of each bit period.

2.1.3 High speed encoding techniques

Although Manchester and Differential Manchester coding have proven to be suitable for Ethernet and Token-Ring networks,

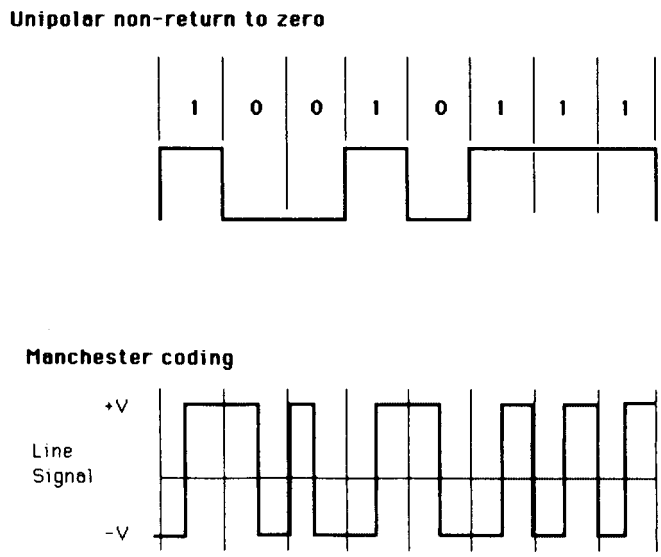

Figure 2.4 Manchester coding. In Manchester coding, a timing transition occurs in the middle of each bit and the line code maintains an equal amount of positive and negative voltage

both require two transitions per bit. This means that their signaling rate in terms of cycles or Hertz per second (c.p.s. or Hz/s) is twice the bit rate. Thus, a 16 Mbps Token-Ring LAN requires a signaling of 32 MHZ.

The necessity to develop local area networks operating at higher data rates would require expensive circuitry if a bi-phase signaling technique such as Manchester or Differential Manchester coding was retained. Rather than attempt to develop circuitry to operate at a signaling rate of 200 MHZ to support a LAN operating rate of 100 Mbps, network designers turned to the use of different encoding schemes. Some of those schemes include MLT-3, 4B5B, 8B6T, 5B6B and 8B10T coding methods.

MLT-3

MLT-3 represents a coding scheme used to support two types of LANs that operate at 100 Mbps. Those LANs include a twisted pair version of the Fiber Distributed Data Interface (FDDI) fiber optic based LAN commonly referred to as Copper Distributed Data Interface (CDDI) and a version of Fast Ethernet standardized as 100BASE-T.

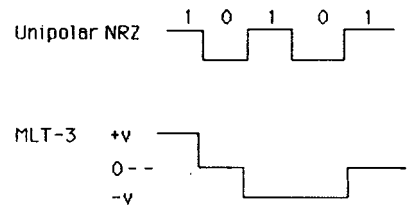

Figure 2.5 MLT-3 encoding

Under MLT-3 encoding, three levels are used for encoding a binary 1: a positive voltage (+v), a negative voltage (−v) and no voltage (0v). The encoding rules are summarized below and illustrated in Figure 2.5. Under MLT-3, encoding occurs based upon the following rules:

1. If the next input bit is 0, then the next output value is the same as the preceding value.
2. If the next input bit is 1, the next output value results in a translation. The translation process examines the preceding output value.
   - If the preceding value was either +v or −v, then the next output value is 0.
   - If the preceding output was 0, then the next output will be non-zero and have the opposite sign to the last non-zero output.

Under MLT-3 encoding bits are encoded as one of three voltage levels. This means that the signaling rate is one-third of the operating rate. Thus, a baud rate of 33.33 MHZ will support a LAN operating rate of 100 Mbps. In comparison, the use of a bi-phase signaling technique would have required a baud rate of 200 MHZ. Thus, MLT-3 coding significantly lowers the signaling rate required to support certain types of high speed local area networks.

Bit transformation

In addition to altering the signaling method, LAN designers use a variety of bit transformation methods to ensure the frequent transition between the two binary digits. The rationale for the use of bit transformation is to obtain synchronization without requiring the use of a bi-phase code. By combining a bit

Table 2.1 Popular LAN encoding schemes

Type of LAN	Encoding
IEEE 802.3 Ethernet (CSMA/CD)	
10BASE-5, 10BASE-2, 10BASE-T	Manchester
100BASE-TX	4B5B/MLT-3
100BASE-T4	8B6T
100BASE-FX	4B5B/NRZI
IEEE 802.5 Token-Ring	Differential Manchester
FDDI	4B5B/NRZI
CDDI	4B5B/IMLT-3
100VG-AnyLAN	5B6B/NRZ
Fiber Channel	8B10B/NRZ

transformation scheme with a coding method LAN designers developed numerous encoding schemes. Table 2.1 lists 11 more popular LANs and the encoding schemes they use.

4B5B

Under 4B5B coding each four bits of data are first encoded into a 5 bit symbol. When used with 200 Mbps Ethernet and FDDI LANs, the 5 bit symbol is then treated on an individual bit by bit basis and encoded using a Non Return to Zero Inverted (NRZI) signaling method. Here NRZI represents a variation of NRZ. Under NRZI a constant voltage pulse is maintained for the duration of the bit time, with a transition from low to high or high to low at the beginning of the bit time used to denote a binary 1, while no transition is used to represent a binary 0. Figure 2.6 provides a comparison of the operation of NRZ and NRZI for encoding a sequence of 8 bits. Note that NRZI is a differential coding method similar to Differential Manchester, in that the signal is decoded by comparing the polarity of adjacent signal elements instead of determining the absolute value of each signal element.

Through the use of 4B5B coding only 16 (2^4) of the 32 (2^5) possible patterns are required for the actual encoding of data. This enabled the codes selected to represent the 16 four-bit data blocks to provide two transitions for each 5-code block. Thus, by

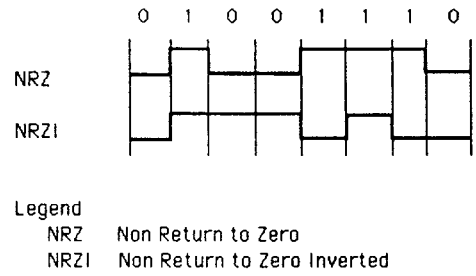

Figure 2.6 NRZ versus NRZI coding

the appropriate selection of data blocks, a sufficient number of transmissions will occur to provide synchronization.

Other coding techniques

In addition to 4B5B, other popular coding methods used by LAN designers include 8B6T, 5B6B and 8B10B. Two of these coding methods are similar to 4B5B, in that a group of *n* bits are used to form an *m* bit symbol, resulting in a coding technique we can classify as nBmB. The exception is 8B6T, which results in 8 bits being mapped into a code group represented by six ternary symbols. As we focus our attention upon specific types of LANs, later in this book we will discuss their coding and signaling method, when appropriate, in additional detail.

2.1.4 Transmission medium

The transmission medium used in a local area network can range in scope from 'twisted-pair' wire, such as is used in conventional telephone lines, to coaxial cable, fiber optic cable and the atmosphere which is used by some transmission schemes to include FM radio and infrared. Each transmission medium has a number of advantages and disadvantages associated with its use in comparison to other media. The primary differences between media concern their cost and ease of installation, the bandwidth of the cable which may permit only one or several transmission sessions to occur simultaneously, the maximum speed of communications permitted and the geographic scope of the network that the medium supports.

Twisted-pair wire

In addition to being the least expensive medium available for LAN installations, twisted-pair wire is very easy to install. Since this wiring uses the RJ45 modular connectors used with the telephone system, once a wire is cut and a connector fastened the attachment of the connector to network devices is extremely simple. Normally, a screwdriver and perhaps a pocket knife are the only tools required for the installation of twisted-pair wire. Anyone who has hooked up a pair of speakers to a stereo set normally has the ability to install this transmission medium.

Although inexpensive and easy to install, unshielded twisted-pair (UTP) wire is very susceptible to noise generated by fluorescent light ballasts and electrical machinery. In addition, a length of twisted-pair wire acts as an antenna. Thus, the longer the wire length the greater the noise it gathers. At a certain length the received noise will obliterate the signal which attenuates or decreases in strength as it propagates along the length of the wire. This noise can affect the error rate of data transmitted on the network, although the utilization of lead-shielded twisted- pair (STP) cable can be employed to provide the cable with a high degree of immunity to the line noise and enable extended transmission distances.

Since the bandwidth of twisted-pair cable is considerably less than that of coaxial or fiber optic cable, normally only one signal is transmitted on this cable at any point in time. As previously explained, this signaling technique is known as baseband signaling and should be compared to the broadband signaling capability of coaxial and fiber optic cable.

It should be noted that, although a twisted-pair wire system can be used to transmit both voice and data, the data transmission is baseband since only one channel is normally used for data. In comparison, a broadband system on coaxial or fiber optic cable can be designed to carry voice and several subchannels of data as well as facsimile and video transmission. Another constraint of unshielded twisted-pair wire is the rate at which data can flow on the network and the distance they can flow. Although data rates up to 155 megabits per second (Mbps) can be achieved, normally local area networks employing data rates beyond 10 Mbps are limited to a transmission distance of 100 meters or less.

In comparison, coaxial and fiber optic cable based systems may be limited in terms of miles. To extend transmission distances over twisted-pair wire, both use shielded wire and

periodically insert repeaters into the cable. The repeater receives a digital signal and then regenerates it; hence it is also known as a data regenerator.

Most high speed twisted-pair based local area networks are hub based, with a maximum cabling distance of 100 meters from the hub port to a participant on the network. In actuality, under cabling standards developed jointly by the Electronic Industries Association and the Telecommunications Industry Association (EIA/TIA), a maximum distance of 90 meters is permitted from a hub port to a wall outlet and 10 meters from the wall outlet to the network device. If cabling is directly from a hub port to a network device, a cabling distance of up to 100 meters is permitted. Later in this book we will examine the EIA/TIA cabling standard as well as the transmission properties of several types of transmission media.

Coaxial cable

Coaxial cable consists of a center conductor copper wire which is then covered by an insulator known as a dielectric. An overlapping woven copper mesh surrounds the dielectric and the mesh which, in turn, is covered by a protective jacket which can consist of polyethylene or aluminum. Figure 2.7 illustrates the composition of a typical coaxial cable; however, it should be noted that over 100 types of coaxial cable are currently marketed. The key differences between such cables involve the number of conductors contained in the cable, the dielectric employed and the type of protective jacket and material used to

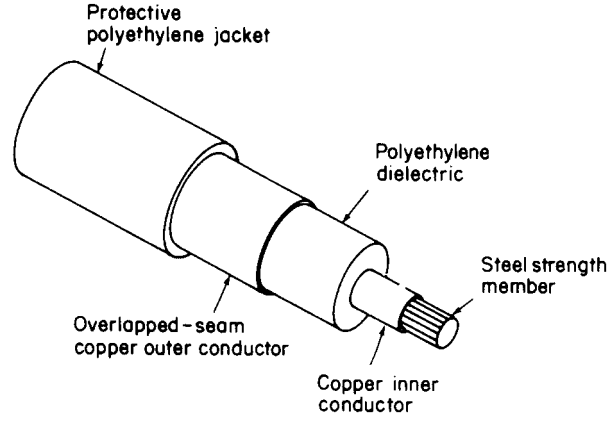

Figure 2.7 Coaxial cable

provide strength to the cable which allows it to be pulled through conduits without breaking.

Two basic types of coaxial cable are used in local area networks, with the type of cable based upon the transmission technique employed: baseband or broadband signaling. Both cable types are much more expensive than twisted-pair wire; however, the greater frequency bandwidth of coaxial cable permits higher data rates for longer distances than can be obtained over twisted-pair wire.

Normally, 50 Ω coaxial cable is used in baseband networks, while 75 Ω cable is used in broadband networks. The latter coaxial is identical to that used in cable television (CATV) applications, including the coaxial cable used in a home. Data rates on baseband networks using coaxial cable range upward to between 50 and 100 Mbps. With broadband transmissions, data rates up to and including 400 Mbps are obtainable.

Hardware interface

A coaxial cable with a polyethylene jacket is normally used for baseband signaling. Data is transmitted from stations on the network to the baseband cable in a digital format and the connection from each station to the cable is accomplished by the use of a simple coaxial T-connector. Figure 2.8 illustrates the hardware interface designed to connect a personal computer to a coaxial cable of a typical baseband local area network. Here the network adapter card is a hardware device that contains the logic to control network access and is inserted into one of the expansion slots in the system unit of the computer. At the rear

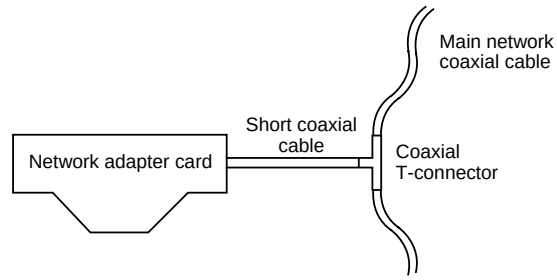

Figure 2.8 Hardware interface to coaxial cable. The network adapter card is installed in the system unit of the PC and connected to the main coaxial cable of the network via a short coaxial cable interfaced to a T-connector

of the computer's system unit a short section of coaxial cable is used to connect the network adapter card to the baseband cable via a T-connector.

Since data on a baseband network travel in a digital form, those signals can be easily regenerated by the use of a device known as a line driver or data regenerator. The line driver or data regenerator is a low-cost device that is constructed to look for a pulse rise; upon detecting the occurrence of the rise, it will disregard the entire pulse and regenerate an entirely new pulse. Thus, you can install low-cost line drives into a baseband coaxial network to extend the distance over which transmission can occur on the cable. Typically, a coaxial cable baseband system can cover several miles and may contain hundreds to thousands of stations on the network.

Broadband coaxial cable

To obtain independent subchannels derived by frequency on coaxial cable broadband transmission requires a method of translating the digital signals from workstations into appropriate frequencies. This translation process is accomplished by the use of radio-frequency (RF) modems which modulate the digital data into analog signals and convert or demodulate received analog signals into digital signals. Since signals are transmitted at one frequency and received at a different frequency, a 'head end' or frequency translator is also required for broadband transmission on coaxial cable. This device is also known as a remodulator as it simply converts the signals from one subchannel to another subchannel.

The requirement for modems and frequency translators normally makes broadband transmission more expensive than baseband. Although the ability of broadband to support multiple channels provides it with an aggregate data transmission capacity that exceeds baseband, in general baseband transmission permits a higher per-channel data flow. While this is an important consideration for mainframe-to-mainframe communications when massive amounts of data must be moved, for most personal computer interactive screen sessions and file transfer operations the speed of either baseband or broadband transmission should be sufficient. This fact may be better understood by comparing the typical transmission rates obtainable on baseband and broadband networks to drive a high-speed dot matrix printer and the differences between the

time required to transmit data on the network and the time required to print the data.

Typical transmission speeds on commonly employed baseband and broadband networks range from 2 to 16 Mbps. In comparison, a high-speed dot matrix printer operating at 120 c.p.s. would require approximately 200 seconds to print 1 second's worth of data transmitted at 2 Mbps and 1600 seconds to print 1 second's worth of data transmitted at 16 Mbps.

Turning our attention to a more modern laser printer capable of printing 8 pages per minute (p.p.m.), if each page contains 200 words with an average of five characters, then in one minute the printer would print 8000 characters. Using 8 bits per character, this would result in 64 000 bits in one minute. Thus, one second's worth of data transmitted at 100 Mbps on a high speed LAN could keep an 8 p.p.m. laser busy for approximately 26 hours!

Based upon the preceding you might question the need for high speed LANs. If your organization makes use of a limited amount of graphics, has no intention of using multimedia application on a network, and does not intend to interconnect a substantial number of devices to the LAN, you probably can use a conventional 10 Mbps Ethernet or 16 Mbps Token-Ring network. However, if your organization uses or intends to use a significant number of graphic intensive applications, multimedia applications, or plans to connect a substantial number of highly active network devices you will more than likely consider the use of a high speed LAN. To illustrate this, consider the use of a LAN to send print jobs to a color laser. One 3 by 5 inch graphic with a resolution of 600 dots per inch (d.p.i.) would require $3 \times 5 \times 600 \times 600$ or 5.4 Mbits without considering the color depth of the image. If the image was stored in what is referred to as true color then 24 bits or 3 bytes are used to represent the color of each bit position. Thus, the full image would require 16.2 Mbits. If your network users transmit a sequence of color print jobs, retrieve color images such as computer aided design (CAD) from network servers, and use or intend to use videoconferencing which can be considered to represent a sequence of 30 images per second, the transmission capacity of conventional local area networks can be expected to become severely taxed. For such situations you will more than likely consider either the use of LAN switches to acquire the ability to obtain multiple simultaneous transmissions or a high speed LAN, both of which will be covered in this book.

Fiber optic cable

Fiber optic cable is a transmission medium for light energy and as such provides a very high bandwidth, permitting data rates ranging up to Gigabits (Gbps) billions of bits per second. The fiber optic cable consists of a thin core of glass or plastic which is surrounded by a protective shield. Several shielded fibers in turn are bundled in a jacket with a central member of aluminum or steel employed for tensile strength.

Digital data represented by electrical energy must be converted into light energy for transmission on a fiber optic cable. This is normally accomplished by a low-power laser or through the use of a light emitting diode and appropriate circuitry. At the receiver, light energy must be reconverted into electrical energy. Normally, a device known as a photo detector, as well as appropriate circuitry to regenerate the digital pulses and an amplifier, are used to convert the received light energy into its original digital format.

In addition to the high bandwidth of fiber optic cables, they offer users several additional advantages in comparison to conventional transmission media. Since data travel in the form of light, they are immune to electrical interference, and building codes that may require expensive conduits to be installed for conventional cables are usually unnecessary. Similarly, fiber optic cable can be installed through areas where the flow of electricity could be dangerous since only light flows through such cables.

Since most fibers only provide a single, unidirectional transmission path, a minimum of two cables is normally required to connect all transmitters to all receivers on a network built using fiber optic cable. Due to the higher cost of fiber optic cable than coaxial or twisted-pair, the dual cable requirement of fiber cables can make them relatively expensive in comparison to other types of cable. In addition, until recently it was very difficult to splice such cable, which usually required sophisticated equipment and skilled installers to implement a fiber optic based network. Similarly, once this type of network was installed, until recently it was difficult to modify the network.

Currently, the cost of the cable, a degree of difficulty of installation and modification make the utilization of fiber optic based local area networks impractical for many commercial applications. Today, the primary use of fiber optic cable is to extend the distance between workstations on a network or to connect two or more networks to one another with the fiber optic

network functioning as a backbone network. The device used to connect a length of fiber optic cable into the LAN or between LANs is a fiber optic repeater. The repeater converts the electrical energy of signals flowing on the LAN into light energy for transmission on the fiber optic cable. At the end of the fiber optic cable, a second repeater converts light energy back into electrical energy. With the cost of the fiber optic cable declining and improvements that simplify the installation and modification of networks using this type of cable continuing to be introduced, the next few years may witness a profound movement toward the utilization of this transmission medium throughout local area networks.

2.1.5 Cabling standards

The Electronics Industry Association/Telecommunications Industry Association Commercial Building Telecommunications Standard commonly referred to as EIA/TIA-568, was ratified in 1992. This standard specifies a variety of building cabling parameters, ranging from backbone cabling used to connect a building's telecommunication closets to an equipment room, to horizontal cabling used to cable individual users to the equipment closet. The standard defines the performance characteristics of both backbone and horizontal cables as well as different types of connectors used with different types of cabling.

Backbone cabling

Four types of media are recognized by the EIA/TIA-568 standard for backbone cabling. Table 2.2 lists the media options supported by the EIA/TIA-456 standard for backbone cabling.

Horizontal cabling

As previously indicated, horizontal cabling under the EIA/TIA-568 standard consists of cable that connects equipment in a telecommunications closet to a user's work area. The media options supported for horizontal cabling are the same as specified for backbone cabling, with the exception of coaxial cable for which 50 ohm thin cable is specified; however, cabling

Table 2.2 EIA/TIA-568 backbone cabling media options

Media type	Maximum cable distance
100 ohm UTP	800 meters (2624 feet)
150 ohm STP	700 meters (2296 feet)
50 ohm thick coaxial cable	500 meters (1640 feet)
62.5/125 μm multimode optical fiber	2000 meters (6560 feet)

distances are restricted to 90 meters in length from equipment in the telecommunications closet to a telecommunications outlet. This permits a patch cord or drop cable up to 10 meters in length to be used to connect a user workstation to a horizontal cabling not exceeding the 10 meters restriction associated with many LAN technologies that use UTP cabling.

UTP categories

One of the more interesting aspects of the EIA/TIA-568 standard is its recognition that different signaling rates require different cable characteristics. This resulted in the EIA/TIA-568 standard classifying UTP cable for five categories. Those categories and their suitability for different types of voice and data applications are indicated in Table 2.3.

In examining the entries in Table 2.3, note that categories 3 through 5 support transmission with respect to indicated signaling rates. This means that the ability of those categories of UTP to support different types of LAN transmission will depend upon the signaling method used by different LANs. For example, consider a LAN encoding technique that results in 6 bits encoded into 4 signaling elements that have a 100 MHZ signaling rate. Through the use of category 5 cable, a data transmission rate of 150 Mbps ((6/4) $\times$ 100) could be supported.

Table 2.3 EIA/TIA-568 UTP cable categories

Category 1	Voice or low-speed data up to 56 Kbps; not useful for LANs.
Category 2	Data rates up to 1 Mbps.
Category 3	Supports transmission up to 16 MHZ.
Category 4	Supports transmission up to 20 MHZ.
Category 5	Supports transmission up to 100 MHZ.

Category 3 cable is typically used for 10 Mbps for Ethernet and 4 Mbps Token-Ring LANs. Category 4 is normally used for 16 Mbps Token-Ring LANs, while category 5 cable supports 100 Mbps Ethernet LANs, such as 100BASE-T and ATM to the desktop at a 155 Mbps operating rate.

2.1.5 Access method

If the topology of a local area network can be compared to a data highway, then the access method might be viewed as the set of rules that enable data from one workstation to successfully reach its destination via the data highway. Without such rules, it is quite possible for two messages sent to the same or a different address by two different workstations to collide, with the result that neither message reaches its destination. There are four access methods primarily used in local area networks, three of which are associated with shared media networks. The three access methods primarily employed in shared media local area networks are Carrier-Sense Multiple Access/Collision Detection (CSMA/CD), Carrier- Sense Multiple Access/Collision Avoidance (CSMA/CA) and token passing. Each of these access methods is uniquely structured to address the previously mentioned collision and data destination problems. A fourth access method involves the creation of a path or channel between the source and destination. This access method is completely different from the preceding methods since it represents a connection-oriented protocol in which data are explicitly transmitted between two stations without other stations receiving the same data and having to check the destination address to determine if they are the recipient. Thus, we can classify CSMA/CD, CSMA/CA, and Token-Ring as connectionless protocols while a fourth access method can be classified as a connection oriented protocol. As connection oriented protocols are usually based upon the use of switches, we will refer to this access method as a switch-based connection-oriented access protocol.

Prior to discussing how access methods work, let us first examine the two basic types of device that can be attached to a local area network to gain an appreciation of the work a shared media access method must accomplish.

Listeners and talkers

We can categorize the operating mode of each device as being a 'listener' or a 'talker'. Some devices, like printers, only receive data and thus operate only as a listener. Other devices, such as personal computers, can either transmit or receive data and are capable of operating in both modes. In a baseband signaling environment where only one channel exists or on an individual channel on a broadband system, if several talkers wish to communicate at the same time a collision will occur unless a scheme is employed that defines when each device can talk and, in the event of a collision, what events must transpire to avoid its recurrence.

For data to correctly reach its destination, each listener must have a unique address and its network equipment must be designed to respond to a message on the net only when it recognizes its address; thus the primary goals in the design of an access method are to minimize the potential for data collision and to provide a mechanism for corrective action when data collide as well as to ensure that an addressing scheme is employed to enable messages to reach their destination.

Carrier-Sense Multiple Access with Collision Detection (CSMA/CD)

Carrier-Sense Multiple Access with Collision Detection can be categorized as a 'listen' then 'send' access method. CSMA/CD is one of the earliest developed access techniques and is the technique used in Ethernet, which is the Xerox Corporation developed local area network whose technology was licensed to many companies and standardized by the IEEE.

Under the CSMA/CD concept, when a station has data to send it first listens to determine if any other station on the network is talking; the fact that the channel is idle is determined in one of two ways based upon whether the network is broadband or baseband.

If a broadband network, the fact that a channel is idle is determined by noting the absence of a carrier tone on the cable. Carrier-sensing thus provides the mechanism to determine whether or not the channel is busy.

Ethernet, like other baseband systems, uses one channel for data transmission and does not employ the use of a carrier. Instead, Ethernet encodes data using a Manchester code in

which a timing transition always occurs in the middle of each bit as previously illustrated in Figure 2.4. Although Ethernet does not transmit data via the use of a carrier, the continuous transitions of the Manchester code can be considered as equivalent to a carrier signal. Carrier-sensing on a baseband network is thus performed by monitoring the line for activity.

In a CSMA/CD network, if the channel is busy, the station will wait until it becomes idle prior to transmitting data. Since it is possible for two stations to listen at the same time and discover an idle channel, it is also possible that the two stations could then transmit at the same time. When this situation arises, a collision will occur. Upon sensing that a collision has occurred, a delay scheme will be employed to prevent a repetition of the collision. Typically, each station will use either a randomly generated or predefined time-out period prior to attempting to retransmit the message that previously collided. Since this access method requires hardware capable of detecting the occurrence of a collision, it is usually more expensive than the hardware required for a similar method that uses collision avoidance.

Figure 2.9 illustrates a CSMA/CD bus-based local area network. Here each workstation is attached to the transmission medium, such as coaxial cable, by a device known as a bus interface unit (BIU). To illustrate the operation of a CSMA/CD network, assume that workstation A is currently using the channel and workstations C and D wish to transmit. The BIUs connecting workstations C and D to the network would listen to the channel and note it was busy. Once workstation A completes its transmission, workstations C and D would attempt to gain

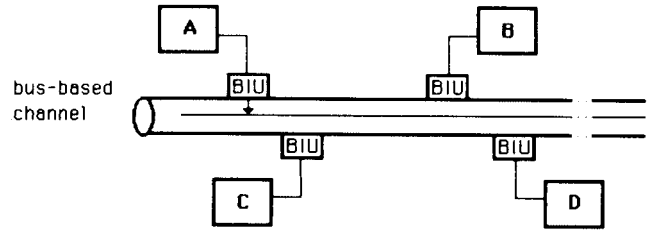

Figure 2.9 CSMA/CD network operation. In a CSMA/CD network, as the distance between workstations increases the resulting increase in propagation delay time increases the probability of the occurrence of collisions

access to the channel. Since workstation A's signal takes longer to propagate down the cable to workstation D than to C, C's BIU notices that the channel is free slightly before workstation D's BIU. However, as workstation C gets ready to transmit, workstation D now assumes the channel is free. Within an infinitesimal period of time C starts transmission followed by D, resulting in a collision. Here the collision is a function of the propagation delay of the signal and the distance between two competing workstations. Due to this, CSMA/CD networks work better as the main cable length decreases.

Although several versions of CSMA/CD are marketed, by far the most common version is based upon licensed technology from Xerox Corporation, which was standardized by the IEEE.

When the IEEE developed its 802.3 standard for CSMA/CD systems it did not precisely follow Xerox's original Ethernet specifications, although there is an extremely high degree of similarity between the original Ethernet specifications and the IEEE 802.3 standard. Recognizing the value of the CSMA/CD access protocol, it forms the basis for the operation of Fast Ethernet that operates at 100 Mbps as well as the evolving Gigabit Ethernet that operates at one billion bits per second (Gbps).

The CSMA/CD access technique is best suited for networks with intermittent transmission, since an increase in traffic volume causes a corresponding increase in the probability of the cable being occupied when a station wishes to talk. In addition, as traffic volume builds under CSMA/CD throughput may decline, since there will be longer waits to gain access to the network as well as additional time-outs required to resolve collisions that occur.

Carrier-Sense Multiple Access with Collision Avoidance (CSMA/CA)

Carrier-Sense Multiple Access with Collision Avoidance represents a modified version of the CSMA/CD access technique. Under the CSMA/CA access technique, each of the hardware devices attached to the talkers on the network estimates when a collision is likely to occur and avoids transmission during those times. Since this technique eliminates the requirement for collision-detection hardware, the cost of hardware to implement this access technique is usually less than that of CSMA/CD hardware. Although this access method appeared to have

promise for adoption it was displaced by the rapid acceptance of Ethernet and its CSMA/CD access protocol.

Token passing

In a token passing access method, each time the network is turned on a token is generated. Consisting of a unique bit pattern, the token travels the length of the network, either around a ring or along the length of a bus. When a station on the network has data to transmit it must first seize a free token. The token is then transformed to indicate that it is in use, and information is added which represents data being transmitted from one station to another. During the time the token is in use the other stations on the network remain idle, eliminating the possibility of collisions occurring. Once the transmission is complete the token is converted back into its original form by the workstation that transmitted the frame and becomes available for use by the next station on the network.

Figure 2.10 illustrates the general operation of a token passing network using a ring topology. Since a station on the network can only transmit when it has a free token, token passing eliminates the requirement for collision-detection hardware. Due to the dependence of the network upon the token, the loss of a station can bring the entire network down. To avoid this, several vendors initially included special backup circuitry in their hardware. When the Token-Ring standard was developed it specified a mechanism whereby the loss of a station results in its disconnection from the network. This enables a failing or purposely powered-off station to be gracefully disconnected from the network without adversely affecting other network devices.

Switch-based, connection-oriented

A switch-based, connection-oriented access protocol requires network devices to request the establishment of a route or path through one or more switches to the recipient. Once a path is established, data are routed on the established path either on a frame or packet basis or for the entire transmission session, with the type of switch used determining the method by which the connection-oriented access between source and destination address occurs.

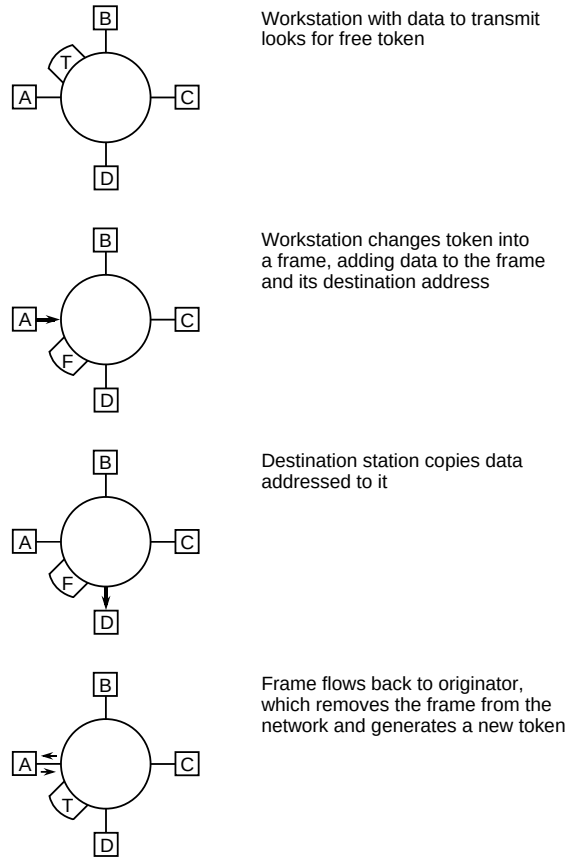

Figure 2.10 Token-Ring operation

The most popular form of switch-based, connection-oriented access protocol is the one used by the Asynchronous Transfer Mode (ATM) protocol. ATM is a switch based LAN and WAN that supports both permanent and switched virtual circuits for the routing of data. A Switched Virtual Circuit (SVC) represents a logical connection established for the routing of data between two endpoints for a transmission session, and a Permanent Virtual Circuit (PVC) represents a similar connection that is established by a network administrator and remains intact until the administrator removes or tears down the connection.

Figure 2.11 illustrates a switch-based, connection-oriented network formed by connecting three switches in a two-tiered hierarchy. In this example note that workstations 1 and 2 connected to a common switch are shown accessing the same

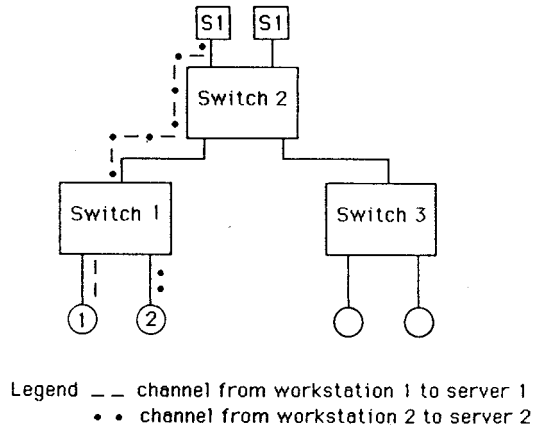

Figure 2.11 Switched-based, connection oriented access

server on a different switch at the same time. This is accomplished by switches subdividing the connection to other switches by time, a technique known as Time Division Multiplexing (TDM). Under ATM terminology the route from switch 1 to switch 2 is referred to as a Virtual Path (VP), and each workstation's connection on the common path is referred to as a Virtual Channel (VC).

Due to the variety of transmission media, network structures and access methods, there is no one best network for all users. Table 2.4 should be used by the reader to obtain a generalized comparison of the advantages and disadvantages of the technical characteristics of local area networks, using the transmission medium as a frame of reference.

2.2 POPULAR TYPES OF LANs

In this section we will examine three popular types of local area networks and one type of LAN whose use can be expected to significantly increase in the near future. The three currently popular types of networks we will examine in this section are Ethernet, Token-Ring and FDDI. Each of these networks presently provides connectivity for millions of workstations. The network whose usage is rapidly increasing and which we will also examine in this section is ATM.

Table 2.4 Technical characteristics of LANs

	Transmission medium			
Characteristic	Twisted-pair wire	Baseband coaxial cable	Broadband coaxial cable	Fiber optic cable
Topology	Bus, star or ring	Bus or ring	Bus or ring	Bus, ring or star
Channels	Single channel	Single channel	Multi-channel	Single, multi-channel
Data rate	Normally up to 16 Mbps, with 100 and 155 Mbps obtainable at distances up to 100 meters	Normally 2 to 10 Mbps, up to 100 Mbps obtainable	Up to 400 Mbps	Up to Gbps
Maximum nodes on net	Usually <255	Usually <1024	Several thousand	Several thousand
Geographical coverage	In hundreds or thousands of feet	In miles	In tens of miles	In tens of miles
Major advantages	Low cost, may be able to use existing wiring	Low cost, simple to install	Supports voice, data, video applications simultaneously	Supports voice, data, video applications simultaneously
Major disadvantages	Limited bandwidth, requires conduits, low immunity to noise	Low immunity to noise	High cost, difficult to install, requires RF modems and headend	Cable cost, difficult to splice

2.2.1 Ethernet

Ethernet is a local area network that uses the CSMA/CD access protocol on a bus structure. The concept behind Ethernet was developed at Xerox Corporation's Palo Alto Research Center (PARC) in Palo Alto, California during the late 1970s. That research center is also considered as the developer of such technological innovations as the graphic interface and the concept of laptop and notebook computers.

Originally Ethernet was developed to provide a mechanism for linking computers located at Xerox's Office Products Division.

Since the intent of Ethernet was to provide a method of connectivity for linking different vendor products, Xerox enlisted Digital Equipment Corporation and Intel Corporation in the development of specifications for the network. In 1980, those three companies jointly published a specification for the Ethernet local area network. That specification was later presented to the IEEE 802 committee and, after several modifications, resulted in the 802.3 standard. Although Ethernet and the 802.3 standard differ in some signaling and formatting methods, many vendors originally introduced equipment capable of supporting both specifications.

Today, when we talk about Ethernet, we typically refer to the IEEE 802.3 standard and the vast majority of Ethernet products currently manufactured support the specifications in that standard. In fact, the terms Ethernet and IEEE 802.3 are considered to be synonymous, although in actuality they are not. In addition, the term Ethernet, as we will shortly note, collectively refer to a family of CSMA/CD access protocol networks that operate on different types of media at data rates from 10 Mbps to 1 Gbps.

Ethernet frame

The transportation of information on an Ethernet local area network occurs in packets formed by eight-bit bytes that are more commonly referenced as frames since data transmission occurs at the data link layer. Figure 2.12 compares the general format of Ethernet and IEEE 802.3 frames. Both are variable in length, since the data field can range in size from a minimum of 46 bytes to a maximum of 1500 bytes.

The frame's preamble provides synchronization and consists of either a seven- or eight-byte sequence of alternating binary ones and zeros. In an IEEE 802.3 frame the start of Frame Delimiter is similar to the preamble; however, its last two bits are always ones. Those bits announce the coming frame and provide additional synchronization for all receivers on the network.

In both the Ethernet and IEEE frame formats the first eight bytes are followed by two six-byte addresses, the first representing the frame's destination, while the second address identifies the originator and is known as the source address. The destination address can specify a single host (unicast), multiple hosts (multicast) or all hosts on the network (broadcast). When

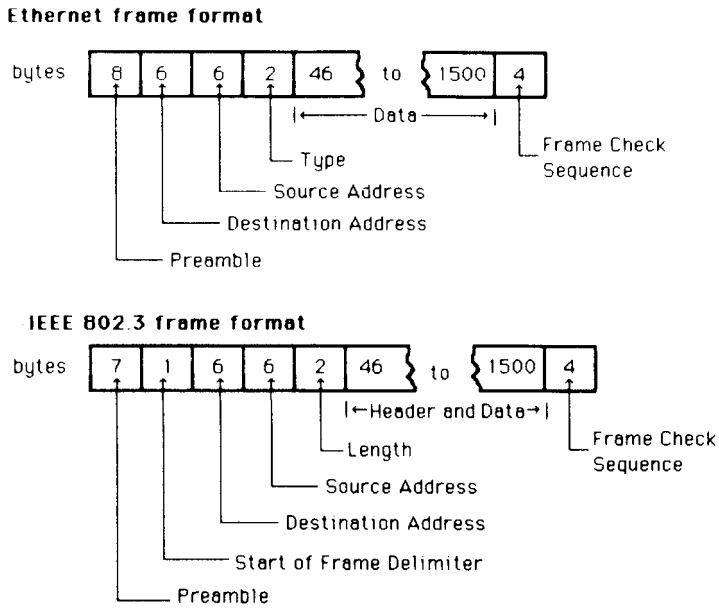

Figure 2.12 Ethernet/IEEE 802.3 frame formats. Both Ethernet and IEEE 802.3 frames can vary in length from 72 to 1526 eight-bit bytes

the first bit in the destination address is set to zero this indicates a unique address. If that bit is set to a one, it indicates either a multicast address to a group or a broadcast to all stations, in which all address bits are then set to one. To ensure the worldwide uniqueness of network addresses the IEEE has assumed the responsibility for assigning the first three bytes of the six-byte address to hardware vendors. Those vendors use that assignment to develop unique addresses in read only memory (ROM) in each network interface card (NIC) they manufacture. The values of the last three bytes are assigned by the network administrator.

The two-byte type field specifies how the data is to be interpreted by the receiver. Under the IEEE 802.3 standard, the type field is replaced by a length field which is also two bytes in length. This field indicates the number of bytes in the data field that are non-pad bytes. Pad bytes are used to fill the data field to its minimum 46-byte size and are also covered by the frame check sequence. Both the type and length fields are followed by the data field, which varies in length from a minimum of 46 bytes to a maximum of 1500 bytes. The last field is the four-byte Frame Check Sequence (FCS). This field

contains a cyclic redundancy check (CRC) that covers the destination field through the data field and provides the error-detection and correction mechanism to ensure data reach the destination correctly.

Although there is a slight difference in frame composition between Ethernet and IEEE 802.3 formats, the key difference between the two is in their support of the network layer. Ethernet specifies the entire link layer through the use of a type field. In comparison, IEEE 802.3 only specifies the lower half of the OSI Reference Model link layer. Readers are referred to Section 2.3 for information on how the link layer is specified under the IEEE 802 series of standards.

From Figure 2.12 it is apparent that both Ethernet and IEEE 802.3 frames can vary in size from a minimum of 72 bytes to a maximum of 1526 bytes. This variance in frame size as well as the fact that both Ethernet and IEEE 802.3 standards require a dead time of 9.6 ms between frames enables us to compute the maximum number of different sized frames that can be transmitted on this type of network per unit of time. From this calculation we can calculate various performance measurements as well as determine such information as the minimum frame processing rate required by bridges and routers that will not degrade network performance and the time required to transfer files between two local area networks connected by a wide area network transmission facility. Readers are referred to Chapter 7 in which several mathematical models of local area network performance are developed and used to compute a variety of performance and throughput statistics.

Types of Ethernet

Ethernet was originally developed as a baseband signaling mechanism in which digital signals are transmitted on a 50 Ω coaxial cable at a data rate of 10 Mbps. The CSMA portion of the access mechanism is built into electronics that reside on a network interface card (NIC) which normally is represented by an adapter card installed in the system unit of a personal computer. The NIC determines whether or not the channel is available for use and either holds the frame in a buffer until the channel is available or transmits the information. Data is transmitted from the NIC to a transmitter/receiver known as a transceiver. The transceiver maintains a balanced electrical

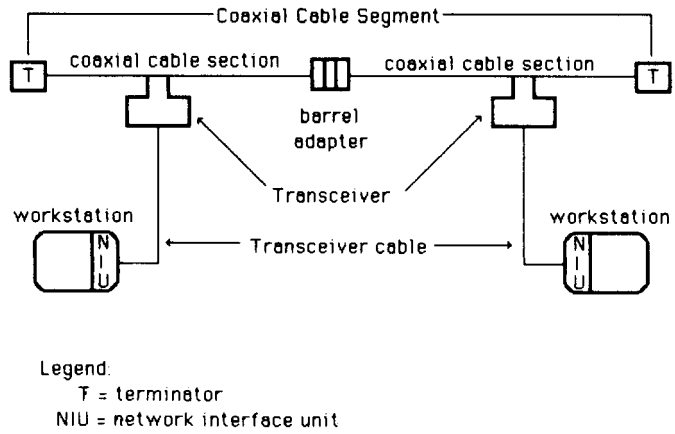

Figure 2.13 Ethernet 50 Ω coaxial cable network structure. An Ethernet LAN using a 50 Ω coaxial bus-based cable has several restrictions. The maximum length of a cable segment is 500 m, the minimum distance between transceivers is 2.5 m, and when segments are joined together by repeaters the maximum separation between nodes is 2.8 km. In addition, there is a maximum of 1024 nodes allowed on this network

signal between the NIC and the coaxial cable and by examining the balance detects collisions, relaying this fact to the NIC.

Figure 2.13 illustrates the relationship between a segment on a main bus-based coaxial cable Ethernet LAN, and its coaxial cable sections, transceivers, NICs and workstations. The maximum length of a 50 Ω coaxial cable segment is 500 m, with no more than 100 transceivers or nodes allowed per segment. The minimum distance between transceivers is 2.5 m, while the total number of nodes on the entire LAN must be less than or equal to 1024 regardless of the number of segments. Although an Ethernet can be extended by joining segments together through the use of repeaters, the maximum separation between nodes is 2.8 km, which represents the maximum length of an Ethernet 50 Ω coaxial cable-based local area network.

An Ethernet LAN using a 50 Ω coaxial bus-based cable has several restrictions. The maximum length of a cable segment is 500 m, the minimum distance between transceivers is 2.5 m and when segments are joined together by repeaters the maximum separation between nodes is 2.8 km. In addition, there is a maximum of 1024 nodes allowed on this network.

The use of 50 Ω coaxial cable was expensive and resulted in several constraints and limitations. In addition to being difficult to bend when installing the cable, a break in the cable would affect all users on a segment. To overcome those limitations

Table 2.5 IEEE 1 Mbps and 10 Mbps Network 802.3 specifications

Feature	10BASE-5	10BASE-2	10BROAD-36	1BASE-5	10BASE-T
Medium	'Thick' 50 Ω coaxial	'Thin' coaxial cable	CATV coaxial cable	Twisted-pair wire cable	Twisted-pair wire
Topology	Bus	Bus	Bus	Star	Star
Segment distance	500 m	200 m	3.6 km	500 m	100 m
Data rate	10 Mbps	10 Mbps	10 Mbps	1 Mbps	10 Mbps

several variations of the original 50 Ω coaxial cable-based Ethernet have been standardized by the IEEE 802.3 committee. Those new versions of the CSMA/CD access protocol involve changes to the media used for the network, its topology, the maximum distance per segment and the operating rate of the network. Table 2.5 summarizes the IEEE 802.3 specifications with respect to media, topology, maximum segment length and data rate for CSMA/CD Networks that operate at 1 Mbps and 10 Mbps. Since there are several types of fast Ethernet networks, we will examine them as a separate entity later in this chapter. Note that the prefix to 'BASE' or 'BROAD' defines the operating rate of the LAN in Mbps. The term BASE or BROAD indicates baseband or broadband, while the suffix denotes the maximum segment length in 100 m multiples. Also note that 10BASE-5 represents the original Ethernet specification based upon the use of 50 Ω coaxial cable referred to as 'thick' coaxial, in comparison to thinner coaxial cable used with 10BASE-2.

Coaxial versus twisted-pair

Although the growth in the use of 50 Ω coaxial cable-based Ethernet networks has substantially diminished, the cause of its decline is not the network but the availability of twisted-pair wire. As a result of the development of 1BASE-5 and 10BASE-T specifications, inexpensive twisted-pair cable that costs less than one-tenth the cost of coaxial cable can be used. In addition, the simplicity of installing twisted-pair in comparison to coaxial cable can halve the cost of installing a network.

Perhaps the key advantage associated with the use of twisted-pair is the greater degree of reliability it provides network users in comparison to a coaxial cable bus-based network. Twisted-

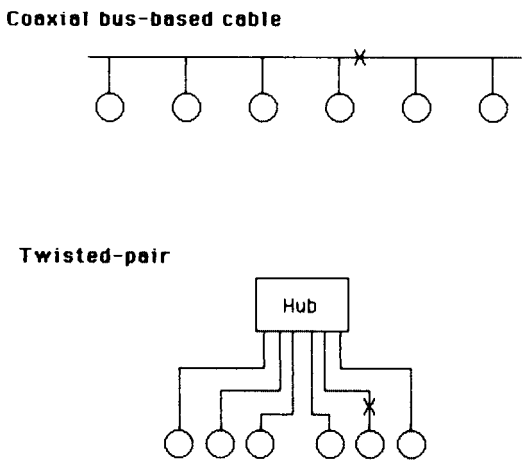

Figure 2.14 Coaxial versus twisted-pair reliability. A twisted-pair-based Ethernet is normally more reliable than a coaxial bus-based network. A cable failure on a star topology only affects one workstation, while the failure of a bus-based cable can affect many users or bring down the entire network

pair based networks use a point-to-point topology in which workstations are connected to a hub to form a star. The resulting star-based network permits both centralized administration and maintenance. To better understand this consider Figure 2.14 which compares a break in a coaxial bus-based Ethernet with a twisted-pair-based Ethernet.

Depending upon where the break occurs on the bus-based cable, the entire network or a portion of the network may become inoperative. In comparison, the failure of a cable on a star-based network only affects the workstation cabled to the hub.

2.2.2 Fast Ethernet

Fast Ethernet is not actually a local area network but a term commonly used to reference a series of three 100 Mbps physical layer LAN specifications in the IEEE 802.3μ addendum. Those specifications include 100BASE-TX, 100BASE-FX, and 100BASE-T4. Each specification maintains the use of the Media Access Control (MAC) protocol used by earlier Ethernet/IEEE 802.3 standards, CSMA/CD.

100BASE-T specifies 100 Mbps operations using the CSMA/CD protocol over two pairs of category 5 unshielded twisted-pair

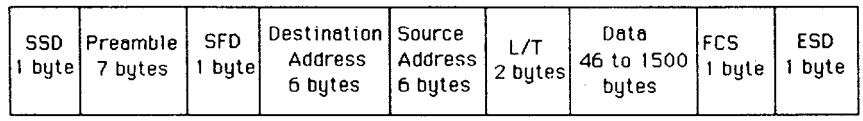

The 100BASE-TX frame differs from the IEEE 802.3 MAC frame through the addition of a byte at each end to mark the beginning and end of the stream delimiter.

SSD Start of stream delimiter
SFD Start of frame delimiter
L/T Length (IEEE 802.3)/Type (Ethernet)
ESD End of stream delimiter

Figure 2.15 Fast Ethernet frame

(UTP) cable. 100BASE-FX changes the LAN transport media to two pairs of fiber, and 100BASE-T4 supports four pairs of category 3, 4, and 5 UTP or shielded-pair (STP) cable.

Frame format

The frame composition associated with each of the three Fast Ethernet standards is illustrated in Figure 2.15. In comparing the composition of the Fast Ethernet frame to Ethernet and IEEE 802.3 frame formats previously illustrated in Figure 2.12, you will note that, other than the addition of starting and ending stream delimiters, the Fast Ethernet frame duplicates the older frames. A third difference between the two is not shown as it is not actually observable from a comparison of frames since this difference is associated with the time between frames. Ethernet and IEEE 802.3 frames are Manchester encoded and have an interpacket gap of 9.6 μs between frames. In comparison, the Fast Ethernet 100BASE-TX frame is transmitted using 4B5B encoding, and IDLE codes (refer to Table 2.7) representing sequences of I (binary 11111) symbols are used to mark a 0.96 μs interpacket gap. Now that we have an overview of the differences between Ethernet/IEEE 802.3 and Fast Ethernet frames, let us focus upon the new fields associated with the Fast Ethernet frame format.

Start of stream delimiter

The Start of Stream Delimiter (SSD) is used to align a received frame for subsequent decoding. The SSD field consists of a sequence of J and K symbols, which defines the unique code 11000 10001. This field replaces the first octet of the preamble in Ethernet and IEEE 802.3 frames whose composition is 10101010.

End of stream delimiter

The End of Stream Delimiter (ESD) is used as an indicator that data transmission terminated normally and a properly formed stream was transmitted. This one-byte field is created by the use of T and R codes (see Table 2.7) whose bit composition is 01101 00111. The ESD field lies outside of the Ethernet/IEEE 802.3 frame and for comparison purposes can be considered to fall within the interframe gap of those frames.

100BASE-T overview

The standardization of 100BASE-T required an extension of previously developed IEEE 802.3 standards. In the definition process of standardization development, both the Ethernet Media Access Control (MAC) and physical layer required adjustments to permit 100 Mbps operational support. For the MAC layer, scaling its speed to 100 Mbps from the 10BASE-T 10 Mbps operational rate required a minimal adjustment, since in theory the 10 BASE-T MAC layer was developed independently of the data rate. For the physical layer, more than a minor adjustment was required, since Fast Ethernet was designed to support three types of media. Using work developed in the standardization process of FDDI in defining 125 Mbps full-duplex signaling to accommodate optical fiber, UTP, and STP through Physical Media Dependent (PMD) sublayers, Fast Ethernet borrowed this strategy. Since a mechanism was required to map the PMD's continuous signaling system to the start-stop half-duplex system used at the Ethernet MAC layer, the physical layer was subdivided. This subdivision is illustrated in Figure 2.16. The PMD sublayer supports the appropriate media to be used, while the convergence sublayer (CS), which was later renamed the physical coding sublayer,

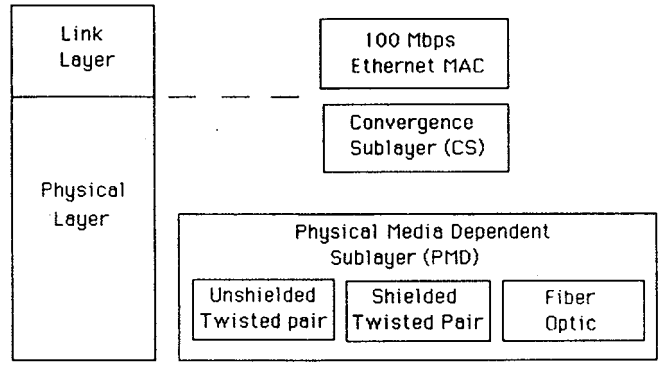

Figure 2.16 Fast Ethernet physical layer subdivision overview

performs the mapping between the PMD and the Ethernet MAC layer.

Although Fast Ethernet represents a tenfold increase in the LAN operating rate from 10BASE-T, to ensure proper collision detection the 100BASE-T network span was reduced to 250 meters, with a maximum of 100 meters permitted between a network node and a hub. The smaller network diameter reduces potential propagation delay. When coupled with a tenfold operating rate increase and no change in network frame size, the ratio of frame duration to network propagation delay for 100BASE-T network is the same as for a 10BASE-T network.

Physical layer

The physical layer subdivision previously illustrated in Figure 2.16, as indicated in the title of the figure, presents an overview of the true layer subdivision. In actuality, a number of changes were required at the physical layer to obtain a 100 Mbps operating rate. Those changes include the use of three wire pairs for data (the fourth is used for collision detection), 8B6T ternary coding (for 100BASE-T4) instead of Manchester coding, and an increase in the clock signaling speed from 20 MHZ to 25 MHZ. As indicated in Table 2.6, in comparison to 10BASE-T the difference at the physical layer resulted in a tenfold increase in the 100BASE-T operating rate.

When the specifications for Fast Ethernet were being developed, it was recognized that the physical signaling layer would incorporate medium dependent functions if support was extended to two pair cable (100BASE-TX) operations. To

Table 2.6 100BASE-T system throughput compared to 10BASE-T

Transmit on 3 pairs versus 1 pair	× 3.00
8B6T coding instead of Manchester	× 2.65
20 to 25 MHZ clock increase	× 1.25
Total throughout increase	10.00

separate medium dependent interfaces to accommodate multiple physical layers, a common interface referred to as the Medium Independent Interface (MII) was inserted between the MAC layer and the physical encoding sublayer. The MII represents a common point of interoperability between the medium and the MAC layer. The MII can support two specific data rates, 10 Mbps and 100 Mbps, permitting older 10BASE-T nodes to be supported at Fast Ethernet hubs. To reconcile the MII signal with the MAC signal, a reconciliation sublayer was added under the MAC layer, resulting in the subdivision of the link layer into three parts: a logical link control layer, a media access control layer, and a reconciliation layer. The top portion of Figure 2.17 illustrates this subdivision.

That portion of Fast Ethernet below the MII, which is the new physical layer, is now subdivided into three sublayers. The lower portion of Figure 2.17 illustrates the physical sublayers for 100BASE-T4 and 100BASE-TX.

The physical coding sublayer performs the data encoding, transmit, receive, and carrier sense functions. Since the data coding method differs between 100BASE-T4 and 100BASE-TX, this difference requires distinct physical coding sublayers for each version of Fast Ethernet.

The Physical Medium Attachment (PMA) sublayer maps messages from the physical coding sublayer (PCS) onto the twisted-pair transmission media, and vice versa.

The Medium-Dependent Interface (MDI) sublayer specifies the use of a standard RJ-45 connector. Although the same connector is used for 100BASE-TX, the use of two pairs of cable instead of four results in different pin assignments.

100BASE-T4

Figure 2.18 illustrates the RJ-45 pin assignments of wire pairs used by 100BASE-T4. Note that wire pairs D1 and D2 are unidirectional. As indicated in Figure 2.18, three wire pairs are

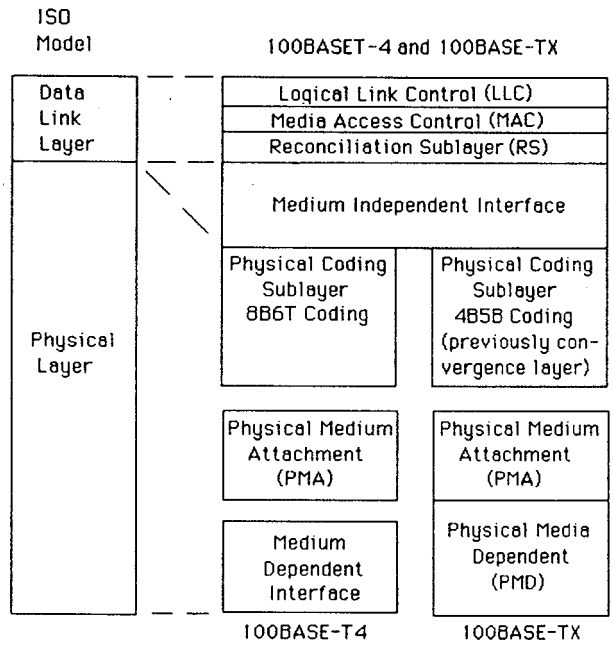

Figure 2.17 100BASE-T4 versus 100BASE-TX physical and link layers

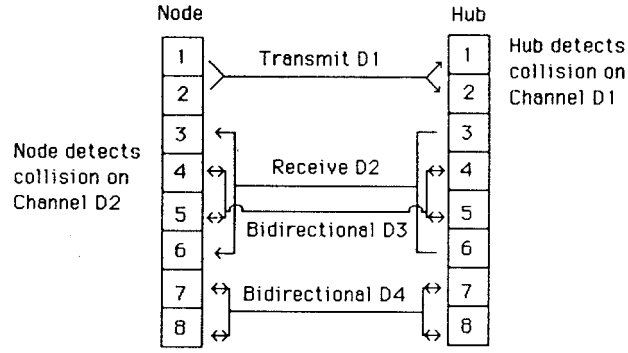

Figure 2.18 100BASE-T4 pin assignments

available for data transmission and reception in each direction, while the fourth pair is used for collision detection.

The 100BASE-T4 Physical coding sublayer implements 8B6T block coding. Under this coding technique, each block of 8 input bits is transformed into a unique code group of 6 ternary symbols. Figure 2.19 provides an overview of the 8B6T coding process used by 100BASE-T4.

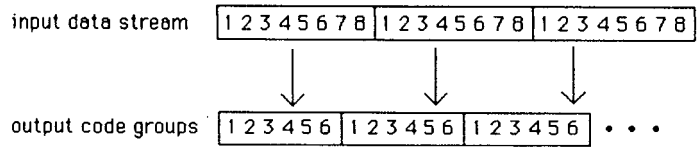

Figure 2.19 8B6T coding process

The output code groups resulting from 8B6T coding flow out to three parallel channels that are placed on three twisted pairs. Thus, the effective data rate on each pair is 100 Mbps/3, or 33.33 Mbps. As 6 bits are represented by 8 bit positions, the signaling rate or baud rate on each cable pair becomes 33 Mbps × 6/8, or 25 MHZ, which is the clock rate used at the MII sublayer.

100BASE-TX

100BASE-TX represents 100BASE-T which supports the use of two pair of category 5 UTP cabling with RJ-45 connectors. A 100BASE-TX network requries a hub, and the maximum cable run is 100 meters from hub port to node, with a maximum network diameter of 250 meters.

Figure 2.20 illustrates the cabling of two pairs of UTP wires between a hub and node to support 100BASE-TX transmission. One pair of wires is used for transmission, while the second pair is used for collision detection and reception of data. The use of a 125 MHZ frequency requires the use of a data grade cable. Thus 100BASE-TX is based upon the use of category 5 UTP.

Although the 100BASE-TX physical layer structure resembles the 100BASE-T4 layer, there are significant differences between the two to accommodate the differences in media used. At the physical coding sublayer, the 100 Mbps start-stop bit stream from the MII is first converted to a full-duplex 125 Mbps bit stream. This conversion is accomplished by the use of the FDDI PMD as the 100BASE-TX PMD. Next, the data stream is encoded

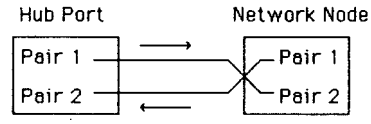

Figure 2.20 100BASE-TX cabling

using a 4B5B coding scheme. The 100BASE-TX PMD decodes symbols from the 125 Mbps continuous bit stream, and converts the stream to 100 Mbps start-stop data bits when the data flow is reversed.

4B5B coding

The use of a 4B5B coding scheme enables data and control information to be carried in each symbol represented by a 5-bit code group. In addition, an inter-stream fill code (IDLE) is defined, as well as a symbol used to force signaling errors. As 4 data bits are mapped into a 5-bit code, only 16 symbols are required to represent data. The remaining symbols not used for control or to denote an IDLE condition are not used by 100BASE-TX and are considered as invalid. Table 2.7 lists the 4B5B 100BASE-TX code groups.

100BASE-FX

100BASE-FX represents the third 100BASE-T wiring scheme, defining Fast Ethernet transmission over fiber optic media. 100BASE-FX requires the use of two-strand 62.5/125 micron multimode fiber media and supports the 4B5B coding scheme, identical to the one used by 100BASE-TX. The use of fiber optics results in longer cable runs than permissible with UTP. This enables 100BASE-FX to be used for connections between bridges, routers, and switches separated by distances greater than supported by UTP cable.

2.2.3 Token passing

Although we normally associate a token-passing local area network with a ring topology, this technology is also standardized for use on a bus.

Bus operation

When used with a bus topology the token bus LAN provides access to the network as if it were a ring. Under the IEEE 802.4 specification for a token bus network a token circulates from end-to-end on the bus and provides a workstation with the

Table 2.7 4B/5B code groups

PCS code group 43210	Name	MII (TXD/RXD) 3210	Interpretation
DATA			
11110	0	0000	Data 0
01001	1	0001	Data 1
10100	2	0010	Data 2
10101	3	0011	Data 3
01010	4	0100	Data 4
01011	5	0101	Data 5
01110	6	0110	Data 6
01111	7	0111	Data 7
10010	8	1000	Data 8
10011	9	1001	Data 9
10110	A	1010	Data A
10111	B	1011	Data B
11010	C	1100	Data C
11011	D	1101	Data D
11100	E	1110	Data E
11101	F	1111	Data F
IDLE			
1111	I		IDLE: Used as inter-Stream fill code
CONTROL			
11000	j		Start-of-Stream Delimiter, Part 1 of 2; always used in pairs with K
10001	K		Start-of-Stream Delimiter, Part 2 of 2; always used in pairs with J
01101	T		End-of-Stream Delimiter, Part 1 or 2; always used in pairs with R
00111	R		End-of-Stream Delimiter, Part 2 of 2; always used in pairs with T
INVALID			
00100	H		Transmit Error; used to force signaling errors
00000	V		Invalid code
00001	V		Invalid code
00010	V		Invalid code
00011	V		Invalid code
00101	V		Invalid code
00110	V		Invalid code
01000	V		Invalid code
01100	V		Invalid code
10000	V		Invalid code
11001	V		Invalid code

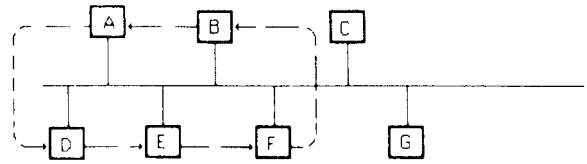

Figure 2.21 Logical ring formation on a physical bus. Token passing on a physical bus enables a logical ring to be formed which may or may not include all stations on the bus

ability to use the bus for a predefined period of time to send or receive data. Under this standard all stations on the network can receive all signals transmitted, a condition known as broadcasting. However, the token access method can be structured to form a logical ring which bypasses one or more workstations as illustrated in Figure 2.21. In this illustration all stations can receive frames; however, stations C and G will not be able to initiate a transmission as they will never receive a token. Here a logical ring is formed consisting of workstations A, D, E, F and B. During normal operations a station which completes its use of a token passes it on to a designated station known as the successor. By passing the token from station to station the logical ring was formed as illustrated in Figure 2.21.

The IEEE 802.4 standard for a physical bus employing token passing as the access method is quite similar to the manufacturing automation protocol (MAP) developed by General Motors. Although this standard has achieved a high degree of usage in industrial applications, its use is considerably overshadowed by the use of a physical ring for token passing and the higher level of requirements to interconnect that type of network. Due to this we will focus our attention upon the Token-Ring.

Ring operation

Each device on a Token-Ring network obtains its connection to the network via a multistation access unit (MAU). The MAU is a passive device that commonly contains 10 connector receptacles as illustrated in Figure 2.22. Eight connectors are used to provide wiring connectivity between network stations, while the connectors on each end of the MAU provide the capability to connect one MAU to another, thus extending the Token-Ring.

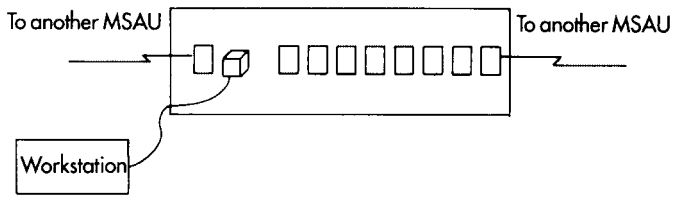

Figure 2.22 MAU connectivity

The MAU functions as a central control point for Token-Ring workstations. Its use facilitates the addition and removal of workstations from a Token-Ring network since the only cabling required is from the workstation to the MAU.

The cable from each workstation to the MAU is known as a lobe. The MAU contains relays that are operated by the voltage supplied by the Token-Ring adapter in each workstation. Thus, if a Token-Ring adapter should fail, the MAU can detect this by the loss of voltage which then opens the relay, in effect shedding the failing lobe from the remainder of the network.

Data flow

Transmissions on a Token-Ring network follow a unidirectional path. Messages flow from one workstation adapter through other workstation adapters on a bit-by-bit basis until they reach their destination, with each intermediate workstation adapter receiving, regenerating and retransmitting each bit that flows on the network. Thus, workstation adapters function as repeaters in addition to performing many other activities that are discussed later in this section.

The MAU functions as a passive device, using the voltage provided by a workstation adapter to switch a relay. If an adapter fails, the relay is opened, disconnecting the lobe from the network.

Network access, token and frame formats

A uniquely coded symbol known as a token is used to provide a workstation with permission to transmit data. To understand how network access is accomplished, let us first examine the token and frame formats used on a Token-Ring network.

Three types of transmission formats are supported on a Token-Ring network: token, abort and frame. The token format as illustrated in the top of Figure 2.23 is the mechanism by

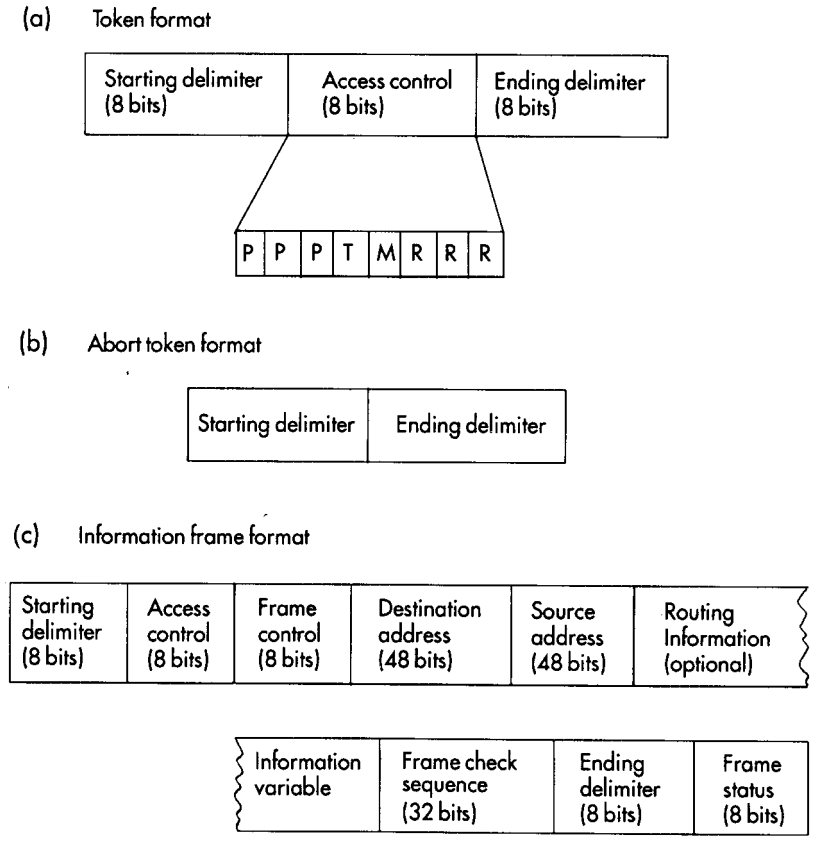

Figure 2.23 Token, abort, and frame formats. (P priority bits, T token bit, M monitor bit, R reservation bits)

which access to the ring is passed from one computer attached to the network to another device connected to the network. Here the token format consists of three bytes, of which the starting and ending delimiters are used to indicate the beginning and end of a token frame. The middle byte of a token frame is an access control byte. Three bits are used as a priority indicator, three bits are used as a reservation indicator, while one bit is used for the token bit and another bit position functions as the monitor bit.

When the token bit is set to a binary zero it indicates that the transmission is a token. When it is set to a binary one it indicates that data are being transmitted.

The second Token-Ring frame format signifies an abort token. In actuality there is no token, since this format is indicated by a

starting delimiter followed by an ending delimiter. The transmission of an abort token is used to abort a previous transmission. The format of an abort token is illustrated in Figure 2.23b.

The third type of Token-Ring frame format occurs when a workstation seizes a free token. At that time the token format is converted into an information frame which includes the addition of frame control, addressing data, an error detection field and a frame status field. The format of the information frame is illustrated in Figure 2.23c. By examining each of the fields in the information frame we will also examine the token and token abort frames due to the commonality of fields between each frame.

Starting/ending delimiters

The starting and ending delimiters mark the beginning and ending of a token or frame. Each delimiter consists of a unique code pattern which identifies it to the network.

Access control

The second field in both token and frame formats is the access control byte. As illustrated at the top of Figure 2.23, this byte consists of four subfields and serves as the controlling mechanism for gaining access to the network. When a free token circulates the network the access control field represents one- third of the length of the frame since it is prefixed by the start delimiter and suffixed by the end delimiter.

The lowest priority that can be specified by the priority bits in the access control byte is zero (000), while the highest is seven (111), providing eight levels of priority. Workstations have a default priority of three, while bridges have a default priority of five. To reserve a token a workstation inserts its priority level in the priority reservation subfield. Unless another workstation with a higher priority bumps the requesting workstation, the reservation will be honored and the requesting station will obtain the token. If the token bit is set to one, this serves as an indication that a frame follows instead of the ending delimiter. The monitor bit is used to prevent a token with a priority exceeding 0 or a frame from continuously circulating on the Token-Ring. This bit is transmitted as a 0 in all tokens and frames, except for a device on the network which functions as an active monitor and thus obtains the capability to inspect and modify that bit.

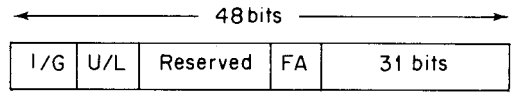

Figure 2.24 Destination address subfields (I/G individual or group bit address identifier, U/L universally or locally administered bit identifier, FA functional address)

The active monitor is the device that has the highest address on the network. All other stations on the network are considered as standby monitors and watch the active monitor.

The function of the active monitor is to determine if a token has been lost, and if so, generate a new one. To accomplish this the active monitor sets the monitor count bit as a frame goes by. If a destination workstation fails or has its power turned off the frame will circulate back to the active monitor, where it is then removed from the network. In the event that the active monitor should fail or be turned off, the standby monitors watch the active monitor by looking for an active monitor frame. If one does not appear within 7 s the standby monitor that has the highest network address then takes over as the active monitor. The last field in the access control byte, reservation bits, enables devices on the network to request the next token to be issued at the priority level of the device.

Frame control

The frame control field informs a receiving device on the network of the type of frame and how it should be interpreted. Frames can be either logical link control (LLC) or reference physical link functions according to the IEEE 802.5 media access control (MAC) standard. As previously noted in this chapter, a logical link control frame includes control information, while a media access control frame carries data.

Destination address

The destination address field is made up of five subfields as illustrated in Figure 2.24. The first bit in the destination address identifies the destination as an individual workstation or as a group of one or more workstations. The latter provides the capability for a message to be broadcast to a group of workstations.

The universally administered address is a unique address permanently encoded into an adapter's ROM. Similar to

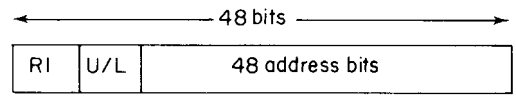

Figure 2.25 Source address field. (RI routing information bit identifier, U/L universally or locally administered bit identifier)

Ethernet, the IEEE assigns blocks of addresses to each vendor manufacturing Token-Ring equipment, which ensures that Token-Ring adapter cards manufactured by different vendors are uniquely defined.

A key problem with the use of universally administered addresses is the requirement to change software coding in a mainframe computer whenever a workstation connected to the mainframe via a gateway is added or removed from the network. To avoid constant software changes, locally administrated addressing can be used. This type of addressing temporarily overrides universally administrated addressing; however, the user is now responsible for insuring the uniqueness of each address.

The functional address subfield in the destination address identifies the function associated with the destination address, such as a bridge, active monitor or configuration report server.

Source address

The source address field always represents an individual address which specifies the adapter card responsible for the transmission. The source address field consists of three subfields as illustrated in Figure 2.25.

The routing information bit identifier identifies the fact that routing information is contained in the routing information field. This bit is set when a frame will be routed across a bridge.

Routing information

The Routing Information Field (RIF) is optional and is included in a frame when the RI subfield of the source address field is set. The RIF is of variable length and contains a control subfield and one or more route fields when included in a frame. This is used to control the flow of frames across one or more bridges.

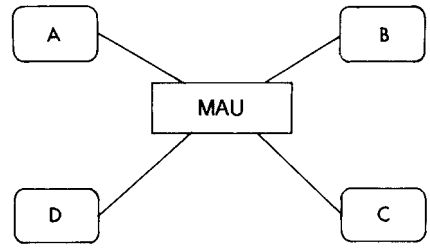

Figure 2.26 Sample network

Information field

The information field is used to contain Token-Ring commands and responses as well as carry user data. This field is of variable length and can be considered to represent the higher level protocol enveloped in a Token-Ring frame.

In the IBM implementation of the IEEE 802.5 Token-Ring standard the maximum length of the information field depends upon the Token-Ring adapter used and the operating rate of the network. Token-Ring adapters with 64 Kbytes of memory can handle up to 4.5 Kbytes on a 4 Mbps network and up to 18 Kbytes on a 16 Mbps network.

Frame check sequence

The frame check sequence field contains four bytes which provide the mechanism for checking the accuracy of frames flowing on the network. The cyclic redundancy check data included in the frame check sequence field covers the frame control, destination address, source address, routing information and information fields. If an adapter computes a cyclic redundancy check that does not match the data contained in the frame check sequence field of a frame, the destination adapter discards the frame information and sets an error bit indicator. This error bit indicator actually represents a ninth bit position of the ending delimiter and serves to inform the transmitting station that the data was received in error.

Frame status

The frame status field contains three subfields that are duplicated for accuracy purposes since they reside outside of CRC checking. One field is used to denote if an address was

recognized, while a second field indicates if the frame was copied at its destination. Each of these fields is one bit in length. The third field, which is two bit positions in length, is currently reserved for future use.

Data flow example

Now that we have examined the token and frame formats, let us examine the flow of data on a Token-Ring network. For simplicity, let us assume our network consists of four workstations labeled A, B, C and D as indicated in Figure 2.26.

Let us assume that station A wants to send data to station C. First, station A must wait until it receives a free token. Then, that workstation changes the token bit from a 0 to a 1 to indicate that a full frame is to flow over the network. Station A then adds destination and source addressing information into those fields and data into the information field, after which station A transmits the frame.

Since data flow is unidirectional, the frame bit flows from the lobe connecting station A to the MAU to station B. Since the destination address contained in the frame is C, the adapter in station B regenerates the frame back to the MAU, from which it flows on the lobe to workstation C.

Station C recognizes its address in the destination address field and copies the frame. To signify this operation, station C changes the address recognized and frame copied subfields in the frame status field and then retransmits the frame back onto the network. When the originating station, which was A, receives the frame it originated, it notes the frame's address recognized and frame copied subfields are set. Then station A removes the frame from the network and transmits a free token onto the network. This enables the next station with data to transmit to grab the token. Since a transmitting station follows frame removal with the generation of a free token, this eliminates a workstation from monopolizing traffic on the network.

Early token release

When IBM introduced its 16 Mbps Token-Ring network in 1988 it introduced a new token passing technique which considerably improves network performance. Known as early token release, this technique is restricted to IBM 16 Mbps Token-Ring networks and was standardized as an extension to the original IEEE 802.5 standard. The basis for early token release is the fact that a

16 Mbps network that has only one frame circulating the network at a time can be considered to have a large amount of unused idle time on the ring. Hence, there is actually room for two different users' sets of data to simultaneously circulate the ring. In recognition of this a station that captures a token and transmits a frame can then generate a free token. This process, known as early token release, improves the performance level obtainable on 16 Mbps networks. However, the use of early token release requires a ring length of approximately 3000 feet to ensure that there is a sufficient delay in circulation time so the token released early does not interfere with a previously transmitted token.

2.2.4 FDDI

Fiber Distributed Data Interface (FDDI) is a local networking standard which provides a 100 Mbps operating rate. In addition, due to the design of FDDI networks which incorporate counter-rotating rings, reliability is increased, since one ring functions as a backup to the other.

Work on FDDI dates from 1982, during which both vendors and standard bodies recognized the need for higher speed LAN products and standards to govern the operation of those products. The FDDI standard was developed by the American National Standards Institute (ANSI) X3T9.5 Task Group.

The original intention of FDDI standards organizations was for the development of specifications for fiber optic media, optical transmitters and receivers, frame formats, protocols, and media access. However, recent developments in the use of twisted-pair have expanded the operation of FDDI to operate over that transmission medium. Known as CDDI, with the C referring to copper, this technique generated a considerable amount of interest. Several vendors now market FDDI over twisted-pair products, however, the use of such products results in a significant limitation on cable distance in comparison with FDDI that can support a ring distance of approximately 100 miles, sufficient to more than cover a large campus or industrial complex.

Advantages

The major advantages of FDDI relate to its operating rate and reliability. FDDI provides an approximate eight- to ten-fold

increase in operating rates over previously developed local area networks. This makes an FDDI network an attractive mechanism to provide an interconnection capability to link lower speed networks as well as to interconnect minicomputers and mainframes via an attachment to their high speed channels. When functioning as a mechanism to interconnect lower speed local area networks, an FDDI LAN serves as a backbone net. One example of its use would be the situation where each floor in a building has its own local area network. An FDDI LAN might then be routed vertically within the building, providing a high speed link between individual networks on each floor.

As previously mentioned, the FDDI standard specified dual fiber optic counter-rotating rings. The dual rings provide an architecture which permits redundancy which can negate the effect of a network failure. In fact, the FDDI standard defines a ring self-heading mechanism which enables stations to identify a failure and take corrective action. In doing so a station that identifies a cable fault would wrap an incoming signal on its healthy side onto an outgoing fiber. Its neighbor on the other side of the fault would also wrap away from the failure, resulting in a dual ring being converted into a single ring which maintains network connectivity. This mechanism will be illustrated later in this section once we review the basic components of an FDDI network. Other advantages of FDDI primarily relate to its use of optical media. Those advantages include the ability to install optical cable without the use of a conduit, the extended transmission distance of an optical system, its immunity to electrical interference, and a high degree of security, since an optical cable is almost impossible to tap.

Hardware components

An FDDI network uses a ring–star topology. Similar to the IEEE 802.5 Token-Ring standard, a rotating token is used to provide stations with permission to transmit data. When an FDDI station wants to transmit information it waits until it detects the token and captures it. Once the station controls the token it can transmit either until it has no more data to send or until a token holding timer expires. When either situation occurs, the station then releases the token onto the ring so that it can be used by the next station that has data to transmit. This token passing technique is more formally known as a timed-token passing

technique and uses bandwidth more efficiently than the 802.5 token passing method. This is because only one token and one frame can be present on a Token-Ring network. In comparison, although only one token is present on an FDDI network at any time, multiple frames from one or more stations can be traversing an FDDI network.

Access to an FDDI network is accomplished through the use of three types of station: a Single Attached Station (SAS) and two types of Dual Attached Station (DAS).

Dual Attached Station

A Dual Attached Station connects to both counter rotating rings used to form an FDDI ring. Each DAS contains two defined optical connection pairs. One pair, called the A interface, contains one primary ring input and the secondary ring output. The second pair, called the B interface, contains the primary ring output and the secondary ring input. Through the use of two optical transceivers each DAS can transmit and receive data on each ring.

A second type of DAS is known as a concentrator. In addition to the previously described A and B interfaces, a DAS concentrator contains a series of extra ports that are called M, or Master ports. The M ports on a DAS concentrator provide connectivity to Single Attached Stations.

Single Attached Station

In comparison to Dual Attached Stations that provide a connection to the dual FDDI rings, a Single Attached Station can only be connected to a single ring. The connection of Single Attached Stations to a DAS concentrator can resemble a star topology, even though the interconnection of DAS and DAS concentrators forms a ring. Since a Single Attached Station only contains a single optical transceiver, its cost is less than that of a Dual Attached Station. However, its inability to connect to the dual ring lowers its reliability in comparison to the connection of workstations to an FDDI network through a Dual Attached Station.

Figure 2.27 illustrates the major components of an FDDI network as well as how a ring can be reconfigured in the event of a cable fault or DAS failure. In this example, it was assumed that a cable fault occurred between the upper right and extreme right Dual Attached Stations. Each of those stations has the

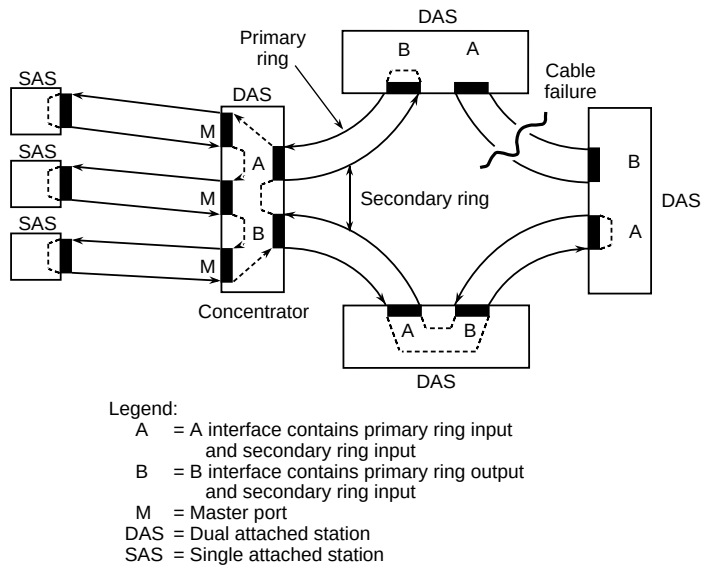

Figure 2.27 FDDI ring operation during cable fault failure. Two dual attched stations can perform a wrap operation to convert a dual ring into a single ring which bypasses a cable or device failure

capability to monitor light levels and recognize a cable failure. By two adjacent stations wrapping away from the failure, the dual ring becomes converted into a single ring and connectivity is restored. When the failure condition is corrected the restoration of an appropriate light level causes each DAS to remove the previously implemented wrap and restores the network to its dual ring operation.

Encoding and signaling

FDDI was the first LAN to use 4B5B encoding with NRZI signaling, enabling a signaling rate of 125 MHZ as well as ensuring that a transition occurs twice for each five bit code. The success of the use of 4B5B encoding with NRZI signaling resulted in the FDDI physical layer being adapted by 100BASE-TX and other local area networks. Although each implementation of 4B5B coding results in the ability to specify 32 unique codes, FDDI's use of this coding technique slightly differs from 100BASE-TX, and other LANs also implement the coding interpretation with slight differences. Table 2.8 lists the FDDI implementation of 4B5B coding. In examining the entries in

Table 2.8 FDDI 4B/5B codes

Function/4-bit group	5B code	Symbol
Starting delimiter		
First symbol of sequential SD pair	11000	J
Second symbol of sequential SD pair	10001	K
Ending delimiter	01101	T
Data symbols		
0000	11110	0
0001	01001	1
0010	10100	2
0011	10101	3
0100	01010	4
0101	01011	5
0110	01110	6
0111	01111	7
1000	10010	8
1001	10011	9
1010	10110	10
1011	10111	11
1100	11010	12
1101	11011	13
1110	11100	14
1111	11101	15
Control indicators		
Logical ZERO (reset)	00111	R
Logical ONE (set)	11001	S
Line status symbols		
Quiet	00000	Q
Idle	11111	I
Halt	00100	H

Table 2.8 you will note that 24 4B5B codes are defined. In comparison, if you examine Table 2.7 you will note that 22 codes are defined. In addition, a comparison of the two tables indicates slight differences in the function of certain codes. When we next examine the FDDI frame formats we will also examine the use of certain 4B5B code functions.

Frame formats

Similarly to Token-Ring networks, there are distinct frames and frame formats that are used on FDDI networks for information transfer. Figure 2.28 illustrates two FDDI frame formats used to transfer information. Like Token-Ring networks, the basic FDDI

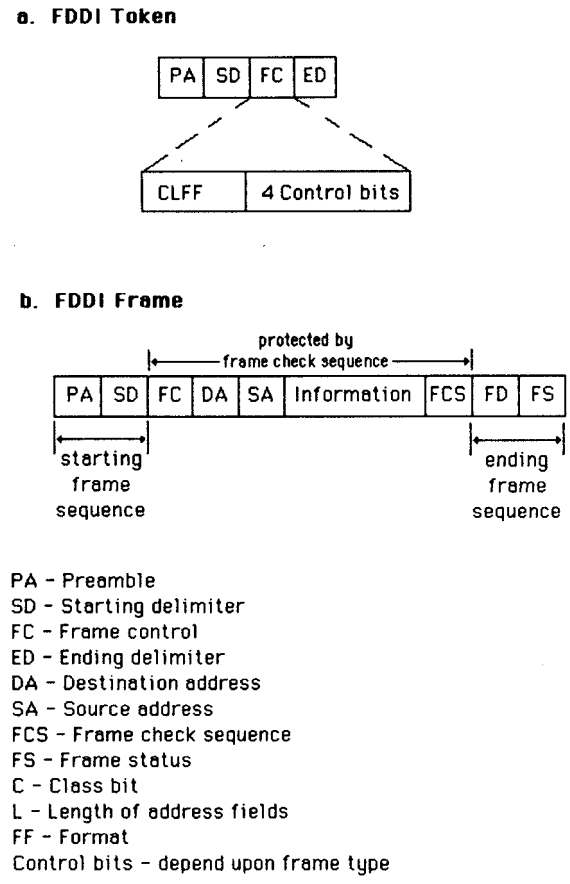

Figure 2.28 FDDI frame formats

frame can convey MAC control data and LLC information. In addition, a station management frame permits management information to be transported between stations and higher level processes. As defined by ANSI, the station management (SMT) standard is used to control the FDDI PMD, PHY, and MAC layers. Services provided by SMT include fault detection, fault isolation, and ring reconfiguration. Data carried by SMT frames can be used by such higher level processes as Simple Network Management Protocol (SNMP) services to permit network administrators to monitor and control each FDDI network node from a central console. In addition to collecting data, SMT provides network administrators with the ability to dynamically alter the network by adding or removing predefined stations. Thus, SMT frames carry both monitoring and control information.

FDDI token

As illustrated in the top portion of Figure 2.28, the FDDI token consists of five fields. The preamble field is variable in length and is formed by 16 or more 5B I symbols. The starting delimiter field consists of the 5B J symbol followed by the 5B K symbol. That field is followed by the frame control field which identifies the type of frame.

The frame control field is eight bits in length, with the class and length of address fields bit positions used to indicate one of two possible settings per bit position. When the class bit is set to 0, this indicates an asynchronous class of transmission, whereas setting the class bit to 1 indicates a synchronous class of transmission. The length of address fields bit indicates the use of 48 bit addressing fields when set to 0 and the use of 16 bit addressing fields when set to 1. The two format bits are used to indicate a MAC or SMT frame when set to 00 or an LLC frame when set to 01. A setting of 10 is implementation dependent, whereas a setting of 11 is currently reserved for future use. The second half of the frame control field consists of four control bits whose values are dependent upon the type of frame defined by the format bits.

There are two special values that can be assigned to the frame control field: hex 80 and hex C0. If the frame control field is set to hex 80 it indicates an unrestricted token, while a value of hex C0 in this eight-bit field indicates a restricted token. The restricted token is generated by a station on an FDDI network that wishes to communicate with another station using all the asynchronous bandwidth available on the network. At the end of this section we will examine the allocation of bandwidth on an FDDI network and the two classes of traffic on that network, asynchronous and synchronous.

When the frame control field is directly followed by the ending delimiter an FDDI token is formed. Here the ending delimiter is the 4B5B T symbol.

FDDI frame

As indicated in Figure 2.28, the first three fields of the FDDI token and frame are the same. Thereafter, the frame contains destination and source address fields which identify the frame recipient and frame originator, respectively. Each address field can be either 16 or 48 bits in length but must be of similar length.

The source address field is followed by a variable information field that can range in length from 0 to 4472 bytes. That field is followed by a frame check sequence (FCS) field 32 bits in length which protects all data from the frame control field through the information field. The ending delimiter and frame status fields function as the ending FDDI frame sequence, with the ending delimiter formed by the use of the 4B5B T symbol which consists of the bit pattern 01101.

Bandwidth allocation

In a Token-Ring network, access is obtained by the setting of priority and reservation bits which enables a station to acquire a token. Once a token has been acquired, it is converted into a single frame to transport a unit of information. In comparison, a token flowing on an FDDI network is removed from the network by a station that has data to transmit, a process referred to as absorption. Once a token is absorbed, the absorbing station can transmit one or more frames prior to returning the token onto the network, with the number of frames that can be transmitted based upon the frame size and the setting of timers within the station. Thus, any discussion of FDDI bandwidth allocation must consider the timers supported by each station in an FDDI network. As those timers, as well as the frame control field of an FDDI token, govern the two classes of traffic that can be carried by an FDDI network, a logical place to initiate an explanation of FDDI bandwidth allocation is by explaining the two classes of traffic supported by this network. Once we have done this we will then examine the timers supported by each FDDI station and then use the preceding to discuss how an FDDI network allocates bandwidth capacity.

Classes of traffic

FDDI defines two classes of traffic, asynchronous and synchronous. These classes of transmission should not be confused with an asynchronous and synchronous mode of transmission. The asynchronous class of transmission is transmission that occurs when the token holding rules of an FDDI network permit transmission. In comparison, the synchronous class of transmission results in a guaranteed percentage of the ring's bandwidth allocated for a particular transmission. Once synchronous bandwidth is allocated, the remaining bandwidth

becomes available for asynchronously transmitted frames. That bandwidth is shared by all stations in a fair and equitable manner based upon the use of timers.

Timers

The control of the amount of asynchronous and synchronous traffic that can be transmitted by a station is governed by FDDI's timed token access protocol. This protocol is based upon the use of timers used by each station to regulate their operation. These timers include a token rotation timer (TRT), token holding timer (THT), and valid transmission timer (TVX).

Token rotation timer

The TRT is used to time the period between the receipt of tokens. Under the timed token access protocol, stations expect to see a token within a specified period of time, referred to as the target token rotation time (TTRT). The value for the TTRT is set when a station initializes itself on the ring and is the same for all stations on the ring.

When a token passes a station, the station sets its TRT to the value of the TTRT and then decrements its TRT timer. If the TRT timer expires prior to the token returning to the station, a counter known as the late counter is incremented. The decision on whether a station can transmit a synchronous or asynchronous class of traffic depends upon the value of the TRT and the value of a counter known as the late counter.

When a token arrives at a station three events occur which govern the allocation of bandwidth. First, upon receiving a token a station can initiate the transmission of synchronous frames. Whether or not it does so, and the number of frames it can transmit, depend upon several factors that will be discussed shortly.

Token holding timer

If the token is received earlier than expected, the token rotation time (TRT) timer will be positive and the station will store that value in its token holding timer (THT). Thus, the value of the THT timer represents the amount of time by which the token was received earlier than expected. Finally, the station resets the TRT timer to the value assigned to the target token rotation timer (TTRT) and begins to decrement that timer.

Synchronous transmission

As previously mentioned, the receipt of a token enables a station to initiate the transmission of synchronous frames. The ability of a station to transmit synchronous frames depends upon whether or not the station was enabled by an application for synchronous transmission. If enabled, the number of synchronous frames that the station can transmit is based upon the size of each frame to be transmitted and the time allocated for synchronous transmission. The frame size governs the amount of time required to place a frame on the ring, while the total time over which the station can transmit synchronously is based upon the value of the station's synchronous allocation timer. That timer is set to zero when a station is not enabled by an application for synchronous transmission. When enabled for synchronous transmission, the value of the synchronous allocation timer can be different for each station on the ring; however, the sum of all synchronous allocation timers on the active stations on the ring must always be less than the target token rotation time.

If enabled for synchronous transmission, a station will either transmit all the frames it has synchronously, or only those frames that can be transmitted within the allocated synchronous allocation timer value. When that timer expires, or all synchronous frames are transmitted and the timer has not expired, the station may then be able to transmit asynchronous frames.

Asynchronous transmission

The decision on whether or not a station can transmit asynchronous frames is based upon the value of the late counter. If the value of the late counter is zero, which means that the TRT timer did not expire, asynchronous frames can be transmitted for the length of time stored in the token holding timer (THT). When the value of that timer reaches zero the token must then be placed back onto the ring.

During both synchronous and asynchronous transmission, the token rotation timer (TRT) continues to decrement. If both synchronous and asynchronous transmissions were stopped due to the expiration of the synchronous allocation timer and the token holding timer, and other stations have data to send, the TRT can be expected to expire prior to the token reappearing at the station. When this occurs, the token will be late, the TRT

will be zero, and the THT will also be set to zero. With a value of zero in the token holding timer the station cannot transmit any asynchronous frames the next time it receives a token. Thus, the timed token access protocol penalizes a station that transmitted its fully allocated amount of traffic; however, the penalty only applies to asynchronous traffic and a station can always transmit synchronous traffic when it receives a token.

If the station is penalized, the next token will arrive early and the station's late counter will be decremented. Once the value of the late counter has reached zero, the station can again begin to transmit asynchronous traffic.

The preceding bandwidth allocation method guarantees an amount of ring capacity to synchronous traffic. Asynchronous traffic is only transmitted when there is spare capacity on the ring, and the use of the previously described counters and timers provides a level of fairness for asynchronous transmission.

In discussing the composition of the frame control field, we indicated that a setting of hex 80 indicates a restricted token. The use of this type of token provides another mechanism for allocating asynchronous transmission by permitting two stations to use all the asynchronous bandwidth available on the ring. When one station wishes to communicate with another station using all of the available asynchronous bandwidth, it transmits its asynchronous frames and then releases a restricted token. Due to FDDI rules, only the last station that receives an asynchronous frame can use a restricted token for asynchronous transmission, thus, this enables two stations to continue transmitting to one another. As the restricted token is only applicable to asynchronous transmission, any station that has synchronous traffic can use that token, ensuring that the guaranteed level of synchronous bandwidth remains available to all stations on the ring.

Transmission example

To illustrate the FDDI capacity allocation algorithm, let us assume that the target token rotation timer was set to 100 milliseconds for all stations, while the synchronous allocation timer was set to 10 milliseconds for our station. Table 2.9 lists the settings of the different station timers and the occurrence of different events during the capacity allocation process for a station on an FDDI network based upon several predefined events occurring on the ring. By examining the entries in Table

Table 2.9 FDDI capacity allocation process example

1. Token arrives at station.
2. TRT is set to value of TTRT (100 ms).
3. Token absorbed by station.
4. Synchronous traffic transmitted for 10 ms (synchronous allocation timer value).
5. Token released onto ring.
6. Token reappears 50 ms later.
7. Token absorbed.
8. TRT now 40 ms due to the 10 ms transmission of synchronous traffic and 50 ms on ring.
9. Token holding timer set to TRT value (40 ms).
10. TRT reset to 100 and begins to decrement.
11. Synchronous traffic again sent for 10 ms.
12. Asynchronous traffic sent for 40 ms (THT value).
13. TRT now has a value of 50 (100 − 10 − 40).
14. Token released.
15. Assume that other stations transmit data and token reappears after 70 ms.
16. TRT expires and late counter incremented to a value of 1.
17. THT set to a value of 0.
18. Assume no synchronous traffic to be sent. Asynchronous traffic cannot be sent since TRT expired and THT now has a value of 0.
19. TRT reset to 100.
20. Assume that token reappears in 30 ms.
21. TRT now set to 70 ms. Although token is early the late count has a value of 1. Thus, token is considered to be late and the station can only transmit synchronous traffic.
22. Token absorbed.
23. Station transmits synchronous traffic for 10 ms (synchronous allocation timer value).
24. Late count value decremented to 0.
25. Token placed back on ring.
26. Assume token reappears 40 ms later.
27. TRT now set to 70 − 40, or 30 ms. Since late count is 0, station can transmit asynchronous traffic for up to 30 ms.

2.9, readers will obtain an appreciation of the method by which timers and the late counter govern the ability of stations to transmit asynchronous and synchronous traffic.

Status

FDDI can be considered to represent a mature backbone LAN technology that has received a high degree of acceptance by industry, academia, and government agencies. Although the use

of Fast Ethernet, LAN switches, and ATM reduced the market for FDDI, its allocation of bandwidth and token rotation predictability will result in a continued market for this product. In addition, as we will shortly note when we next discuss ATM, the use of that technology as a backbone for Ethernet or Token-Ring networks requires a process known as LAN emulation which can have certain performance limitations that are not associated with the use of FDDI.

2.2.5 Logical link control frame format

As mentioned earlier in this chapter, the IEEE subdivided the Data Link Layer of the ISO Reference Model into Media Access Control (MAC) and Logical Link Control (LLC) sublayers.

The logical link control (LLC) sublayer was defined under the IEEE 802.2 standard to make the method of link control independent of a specific access method. Thus, the 802.2 method of link control spans Ethernet (IEEE 802.3), Token Bus (IEEE 802.4), and Token-Ring (IEEE 802.5) local area networks. Functions performed by the LLC include generating and interpreting commands to control the flow of data, including recovery operations, for when a transmission error is detected.

Under the 802.5 Token-Ring standard each type of frame is specified by an appropriate setting in the Token-Ring Access Control Field. When an LLC frame is specified it is transported in the Information field as a sequence of four subfields. Under the 802.3 standard link control information is carried in the data field as an LLC Protocol Data Unit. Although Ethernet does not provide a direct mechanism to identify that the frame transports LLC data, similarly to Token-Ring the Ethernet LLC Protocol Data Unit (PDU) also contains the same four fields. Since Ethernet frame determination requires the examination of the contents of the Length field and the composition of LLC PDUs and LLC frames which are the same, we will primarily focus our attention upon Ethernet's LLC although the discussion of LLC types and classes of service presented in this section is applicable to both networks.

Link control information is carried within the data field of an IEEE 802.3 frame as an LLC Protocol Data Unit. Figure 2.29 illustrates the relationship between the IEEE 802.3 frame and the LLC Protocol Data Unit.

Service Access Points (SAPs) function much like a mailbox. Since the LLC layer is bound below by the MAC sublayer and

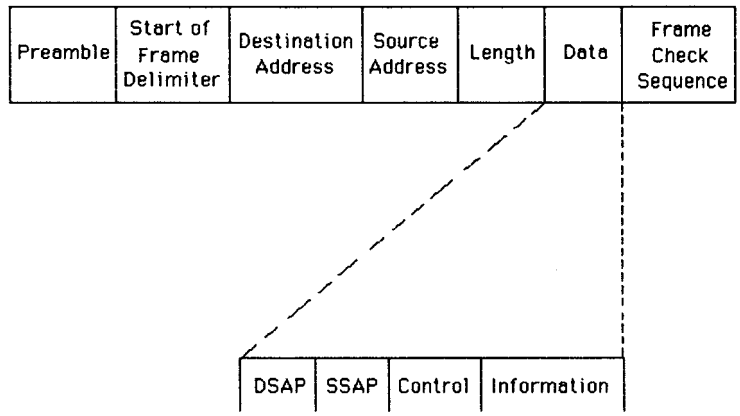

Figure 2.29 Formation of LLC protocol data unit. Control information is carried within a MAC frame

bound above by the network layer, SAPs provide a mechanism for exchanging information between the LLC layer and the MAC and network layers. For example, from the network layer perspective, a SAP represents the place to leave messages about the services requested by an application.

The Destination Services Access Point (DSAP) is one byte in length, and is used to specify the receiving network layer process. As an IEEE 802.3 frame does not include a type field, the DSAP field is used to denote the destination upper-layer protocol carried within the frame. For example, the DSAP hex value EO indicates that the data field contains NetWare data.

The Source Service Access Point (SSAP) is also one byte in length. The SSAP specifies the sending network layer process. Since the destination and source protocols must be the same, the value of the SSAP field will always match the value of the DSAP field. Both DSAP and SSAP addresses are assigned by the IEEE. For example, the hex address FF represents a DSAP broadcast address.

The control field contains information concerning the type and class of service being used for transporting LLC data. For example, a hex value of 03 when NetWare is being transported indicates that the frame is using an unnumbered format for connectionless services.

Types and classes of service

Under the 802.2 standard, there are three types of service available for sending and receiving LLC data. These types are

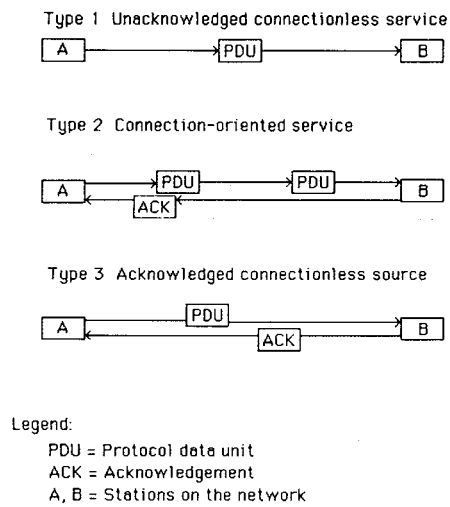

Figure 2.30 Local link control service types

discussed in the next three sections. Figure 2.30 provides a visual summary of the operation of each LLC service type.

Type 1

Type 1 is an unacknowledged connectionless service. The term connectionless refers to the fact that transmission does not occur between two devices as if a logical connection were established. Instead, transmission flows on the channel to all stations; however, only the destination address acts upon the data. As the name of this service implies, there is no provision for flow control or for error recovery. Therefore, this is an unreliable service.

Despite those shortcomings, Type 1 is the most commonly used service for IEEE 802 LANs since most protocol suites use a reliable transport mechanism at the transport layer, thus eliminating the need for reliability at the link layer. In addition, by eliminating the time needed to establish a virtual link and the overhead of acknowledgments, a Type 1 service can provide a greater throughput than other LLC types of service.

Type 2

The Type 2 connection-oriented service requires that a logical link be established between the sender and the receiver prior to

information transfer. Once the logical connection has been established, data will flow between the sender and receiver until either party terminates the connection. During data transfer, a Type 2 LLC service provides all the functions lacking in a Type 1 service, using a sliding window for flow control. When IBM's SNA data is transported on a LAN, it uses connection-oriented services. Type 2 LLC is also commonly referred to as LLC 2.

Type 3

The Type 3 acknowledged connectionless service contains provisions for the setup and disconnection of transmission; it acknowledges individual frames using the stop-and-wait flow control method. Type 3 service is primarily used in an automated factory process-control environment, where one central computer communicates with many remote devices that typically have a limited storage capacity.

Classes of service

All logical link control stations support Type 1 operations. This level of support is known as Class I service. The classes of service supported by LLC indicate the combinations of the three LLC service types supported by a station. Class I supports Type 1 service, Class II supports both Type 1 and Type 2, Class III supports Type 1 and Type 3 service, and Class IV supports all three service types. Since service Type 1 is supported by all classes, it can be considered a least common denominator, enabling all stations to communicate using a common form of service.

Other Ethernet frame types

Two additional frame types that warrant discussion are Ethernet-802.3 and Ethernet-SNAP. In actuality, both types of frames represent a logical variation of the IEEE 802.3 frame in which the composition of the data field varies from the composition of the LLC protocol data unit previously illustrated in Figure 2.29.

Ethernet-802.3

The Ethernet-802.3 frame represents a proprietary subdivision of the IEEE 802.3 data field to transport NetWare. Ethernet-

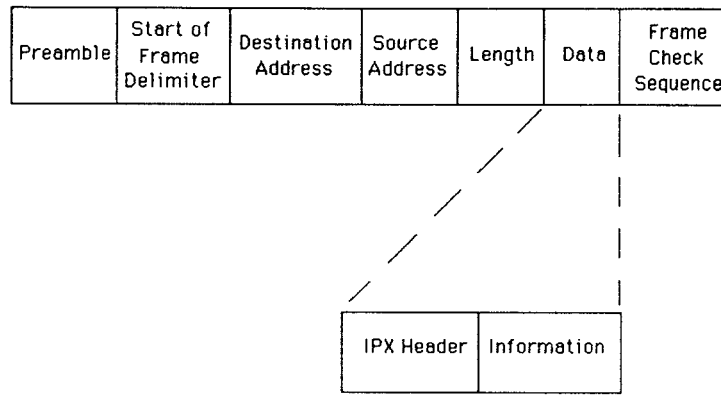

An Ethernet-802.3 frame subdivides the data field into an IPX header field and an information field.

Figure 2.31 Novell's NetWare Ethernet-802.3 frame

802.3 is one of several types of frame that can be used to transport NetWare. The actual frame type used is defined at system setup by binding NetWare to a specific type of frame.

Figure 2.31 illustrates the format of the Ethernet-802.3 frame. Due to the absence of LLC fields, this frame is often referred to as raw 802.3.

For those using or thinking of using NetWare, a word of caution is in order concerning frame types. Novell uses the term Ethernet-802.2 to refer to the IEEE 802.3 frame. Thus, if you set up NetWare for Ethernet-802.2 frames, in effect your network is IEEE 802.3-compliant.

Ethernet-SNAP

The Ethernet Subnetwork Access Protocol (Ethernet-SNAP) frame, unlike the Ethernet-802.3 frame, can be used to transport several protocols. AppleTalk Phase II, NetWare, and TCP/IP protocols can be transported due to the inclusion of an Ethernet type field in the Ethernet-SNAP frame. Thus, SNAP can be considered as an extension that permits vendors to create their own Ethernet protocol transports. Ethernet-SNAP was defined by the IEEE-802.1 committee to facilitate interoperability between IEEE 802.3 LANs and Ethernet LANs. This was accomplished, as we will soon note, by the inclusion of a type field in the Ethernet-SNAP frame.

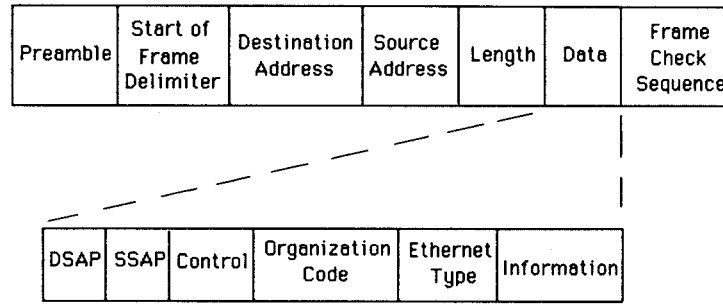

Figure 2.32 Ethernet-SNAP frame format

Figure 2.32 illustrates the format of an Ethernet-SNAP frame. Although the format of this frame is based upon the IEEE 802.3 frame format, it does not use DSAP and SSAP mailbox facilities and the control field. Instead, it places specific values in those fields to indicate that the frame is a SNAP frame.

The value hex AA is placed into the DSAP and SSAP fields, and hex 03 is placed into the control field to indicate that a SNAP frame is being transported. The hex 03 value in the control field defines the use of an unnumbered format, which is the only format supported by a SNAP frame.

The organization code field refers to the organizational body that assigned the value placed in the following field, the Ethernet type field. A hex value of 00-00-00 in the organization code field indicates that Xerox assigned the value in the Ethernet type field. Through the use of the Ethernet-SNAP frame you obtain the ability to transport multiple protocols in a manner similar to the original Ethernet frame that used the type field for this purpose.

Frame determination

Through software, a receiving station can determine the type of frame and correctly interpret the data carried in the frame. To accomplish this, the value of the two bytes that follow the source address is first examined. If the value is greater than 1500, this indicates the occurrence of an Ethernet frame. If the value is less than or equal to 1,500, the frame can be either a pure IEEE 802.3 frame or a variation of that frame. Thus, more bytes must be examined.

If the next two bytes have the hex value FF:FF, the frame is a NetWare Ethernet-802.3 frame. This is because the IPX header has hex FF:FF in the checksum field contained in the

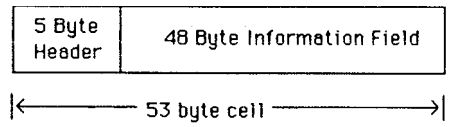

The ATM cell is of fixed length, consisting of a 48-byte information field and five-byte header.

Figure 2.33 The ATM cell

first two bytes in the IPX header. If the two bytes contain the hex value AA:AA, this indicates that it is an Ethernet-SNAP frame. Any other value determined to reside in those two bytes then indicates that the frame must be an Ethernet-802.3 frame.

2.3 ATM

Previously in this chapter we briefly discussed the basic operation of Asynchronous Transfer Mode (ATM), noting that it normally represents a switch-based, connection-oriented technology. In this section we will examine the rationale for the operation and utilization of ATM. In doing so, we will focus our attention upon the use of ATM for both desktop to desktop data transfer as well as its use as a backbone transmission mechanism to interconnect connectionless LANs, such as Ethernet and Token-Ring.

2.3.1 Rationale

ATM was developed as a transmission technology that can support both voice and data on a common network infrastructure. To accomplish this, ATM uses a fixed size transmission unit known as a cell.

The ATM cell

The ATM cell is relatively short, containing a 48 byte information field and a 5 byte header. Figure 2.33 illustrates the basic composition of the ATM cell.

The selection of a relatively short 53 byte cell was based upon the necessity to minimize the effect of data upon such time dependent transmissions as voice and real-time video. For example, if a variable frame length technology such as Ethernet or Token-Ring is used, it becomes possible for a lengthy data transmission that fills a frame to its maximum length and which gains access to a network between a sequence of digitized voice filled frames to delay the receipt of succeeding frames so that, upon their receipt, the converted voice sounds distorted by time. ATM is designed to eliminate such problems as cells are relatively short, resulting in cells transporting voice being able to arrive on a regular basis.

A second significant advantage associated with the use of fixed length cells concerns the design of switching equipment. ATM logic can be developed as firmware embedded in hardware, enabling faster processing at a lower cost than if software was used to perform ATM operations.

In addition to its relatively short cell length facilitating the integration of voice and data, ATM provides three additional benefits. Those benefits are in the areas of scalability, transparency, and traffic classification.

Scalability

ATM cells can be transported on LANs and WANs at a variety of operating rates. This enables different hardware, such as LAN and WAN switches to support a common cell format, a feature lacking with other communications technologies. Within a few years, an ATM cell generated on a 25 Mbps LAN will be able to be transported from the LAN via a T1 line at 1.544 Mbps to a central office where it might be switched onto a 2.4 Gbps SONET network for transmission on the communications carrier infrastructure, with the message maintained in the same series of 53 byte cells, with only the operating rate scaled for a particular transport mechanism.

Transparency

The ATM cell is application-transparent, enabling it to transport voice, data, images, and video. Due to its application transparency, ATM enables networks to be constructed to support any type of application or application mix instead of requiring

organizations to establish separate networks for different applications.

Traffic classification

Five classes of traffic are supported by ATM, including one constant bit rate, three types of variable bit rate, and a user-definable class. As ATM standards were further developed, supports for two traffic classes were merged together into a common ATM adaptation layer (AAL) protocol. Later in this section we will discuss the role of the AAL and its support for different classes of traffic.

By associating such metrics as cell transit delay, cell loss ratio, and cell delay variation to a traffic class, it becomes possible to provide a guaranteed Quality of Service on a demand basis. This enables a traffic management mechanism to adjust network performance during periods of unexpected congestion to favor traffic classes based upon the metrics associated with each class.

The Quality of Service (QoS) is one of the key features of ATM which enables the technology to provide a predefined level of support to different types of data streams. An end point requesting the setup of a connection through an ATM network can request a QoS from the network. Once granted, the end point will be assured that the network will provide the selected QoS for the life of the connection. The ATM Forum presently defines five traffic classes that are summarized in Table 2.10. That table includes the type of each traffic class, a description of its intended use, and an example of its potential utilization.

As we probe further into ATM we will note that the only priority field within an ATM cell is used to indicate whether or not a cell can be dropped. Thus, the method used by an ATM network to provide a QoS is not priority-based. Instead, it is based by a set of traffic parameters that define such metrics as the Peak Cell Rate (PCR), Cell Delay Variation Tolerance (CDVT), Sustainable Cell Rate (SCR), Burst Tolerance (BT), and Minimum Cell Rate (MCR). Only some of these metrics are applicable for certain traffic classes. For example, only the Peak Cell Rate which specifies how often data samples are transmitted and the Cell Delay Variation Tolerance (CDVT) which determines the amount of displacement of a signal from its intended location are applicable for CBR traffic. Later in this

Table 2.10 ATM traffic classes

Traffic class	Description
Continuous Bit Rate (CBR)	Constant bit traffic with a fixed timing relationship between data samples, such as an emulated voice circuit.
Variable Bit Rate—Real Time (VBR/RT)	Variable bit rate traffic that has a fixed timing relationship between data samples, such as compressed video.
Variable Bit Rate—Non-Real Time (VBR/NRT)	Variable bit rate traffic that has no timing relationship between data samples but for which a guarantee of a Qos is required, such as Frame Relay.
Available Bit Rate (ABR)	Variable data transmission that has no timing relationship and can be handled on a best effort basis, such as electronic mail.
Unspecified Bit Rate (UBR)	A class of traffic for which there is no service guarantee. The user can transmit any amount of data up to a specified maximum but the network does not guarantee delay or a cell loss rate.

section we will examine the relationship between ATM Adaption Layers, traffic classes, and traffic definition metrics.

2.3.2 The ATM Protocol Stack

Similar to other networking architectures, ATM is a layered protocol. The ATM protocol stack is illustrated in Figure 2.34 and consists of three layers: the ATM Adaptation Layer (AAL), the ATM Layer, and the Physical Layer. Both the AAL and

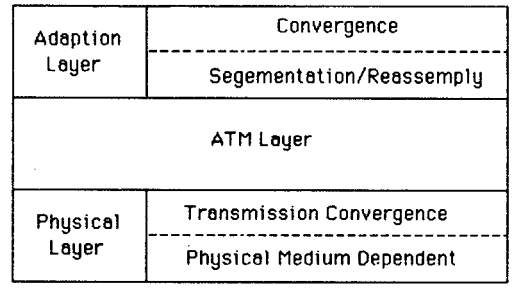

Figure 2.34 The ATM protocol stack

Physical Layers are subdivided into two sublayers. Although the ATM protocol stack consists of three layers, as we will shortly note, those layers are essentially equivalent to the first two layers of the ISO Reference Model. However, since ATM possesses many of the characteristics of a layer 3 or network layer protocol such as a hierarchical address space and a complex routing protocol, some persons consider it to represent a network protocol.

ATM Adaptation Layer

As illustrated in Figure 2.34, the ATM Adaptation Layer consists of two sublayers: a convergence sublayer and a segmentation and reassembly sublayer. The function of the AAL is to adapt higher level data into formats compatible with ATM Layer requirements. To accomplish this task the ATM Adaptation Layer subdivides user information into segments suitable for encapsulation into the 48 byte information fields of cells. The actual adaptation process depends upon the type of traffic to be transmitted, although all traffic winds up in similar cells. Currently there are four different AALs defined, referred to as AAL classes, which are described later in this chapter.

When receiving information, the ATM Adaptation Layer performs a reverse process. That is, it takes cells received from the network and reassembles them into a format the higher layers in the protocol stack understand. This process is known as reassembly. Thus, the segmentation and reassembly processes result in the name of the sublayer that performs those processes.

The ATM Layer

As illustrated in Figure 2.34, the ATM Layer provides the interface between the AAL and the Physical Layer. The ATM Layer is responsible for relaying cells both from the AAL to the Physical Layer and to the AAL from the Physical Layer. The actual method by which the ATM Layer performs this function depends upon its location within an ATM network. As an ATM network consists of endpoints and switches, the ATM Layer can reside at either location. Similarly, a Physical Layer is required at both ATM endpoints and ATM switches.

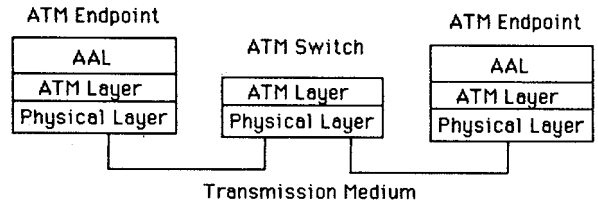

Figure 2.35 The ATM protocol stack within a network

As a switch examines the information within an ATM cell to make switching decisions, it does not perform any adaptation functions. Thus, the ATM switch operates at layers 1 and 2, while ATM endpoints operate at layers 1 through 3 of the ATM protocol stack as shown in Figure 2.35.

When the ATM Layer resides at an endpoint, it will generate idle or empty cells whenever there are no data to send, a function not performed by a switch. Instead, in the switch the ATM Layer is concerned with facilitating switching functions, examining cell header information which enables the switch to determine where each cell should be forwarded to. For both endpoints and switches, the ATM Layer performs a variety of traffic management functions to include buffering incoming and outgoing cells as well as monitoring the transmission rate and conformance of transmission to service parameters that define a Quality of Service (QoS). At endpoints the ATM Layer also indicates to the AAL whether or not there was congestion during transmission, permitting higher layers to initiate congestion control.

Physical Layer

Although Figures 2.34 and 2.35 illustrate an ATM Physical Layer, a specific physical layer is not defined within the protocol stack. Instead, ATM uses the interfaces to existing physical layers defined in other protocols, which enables organizations to construct ATM networks on different types of physical interfaces which in turn connect to different types of media. Thus, the omission of a formal physical layer specification results in a significant degree of flexibility which enhances the capability of ATM to operate on LANs and WANs.

2.3.3 ATM operation

As previously discussed, ATM represents a cell-switching technology that can operate at speeds ranging from the T1 1.544 Mbps to the gigabit per second rate of SONET. In doing so the lack of a specific physical layer definition means that ATM can be used on many types of physical layers, which makes it a very versatile technology.

Components

ATM networks are constructed upon the use of five main hardware components. Those components include ATM network interface cards, LAN switches, ATM routers, ATM WAN switches and ATM service processors.

ATM network interface cards

An ATM network interface card (NIC) is used to connect a LAN based workstation to an ATM LAN switch. The NIC converts data generated by the workstation into cells that are transmitted to the ATM LAN switch and converts cells received from the switch into a data format recognizable by the workstation.

LAN switch

A LAN switch is a device used to provide interoperability between older LANs, such as Ethernet, Token-Ring or FDDI as well as from those networks to ATM. To provide connectivity to ATM, the LAN switch supports a minimum of two types of interface, with one being an ATM interface which enables the switch to be connected to an ATM switch that forms the backbone of the ATM infrastructure. The other interface or interfaces represent connections to older types of LANs.

The LAN switch functions as both a switch and protocol converter. Data received on one port destined to the ATM network are converted from frames to cells and transferred to the switch port providing a connection to the ATM switch. One of the key functions that must be performed by the switch in conjunction with the ATM switch it is connected to is a mapping between the MAC addresses used on a LAN switch and the virtual path/virtual channel (VP/VC) identifiers used by ATM. This mapping process is accomplished through a technology

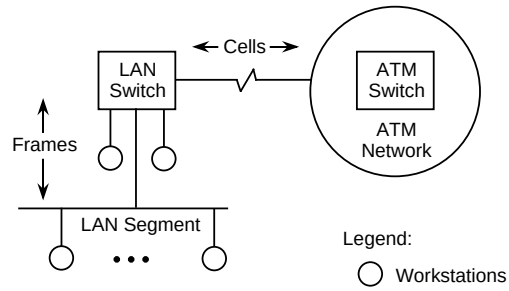

A LAN switch provides both a switching and protocol conversion function, allowing non-ATM devices to access an ATM network

Figure 2.36 Using a LAN swtich

known as LAN Emulation (LANE), which is described in Chapter 3 in the section which covers LAN switches.

As one LAN switch port can be capable of servicing a LAN segment, the use of a switch can minimize an organization's investment in ATM NICs. This is illustrated in Figure 2.36 which illustrates the use of a LAN switch with a single ATM port to provide access to an ATM network for individual workstations connected directly to individual switch ports as well as a group of workstations on a LAN segment. Through the use of the LAN switch, an organization can selectively upgrade existing LANs to ATM while obtaining a connection to the ATM network.

ATM router

An ATM router, or perhaps a more correct terminology, an ATM supportable router, is a router containing one or more ATM NICs. As such, it can provide a direct or indirect capability for LAN workstations to access an ATM network or for two ATM networks to be interconnected. For example, a network segment or individual workstations could be connected to a router which in turn is connected to a LAN switch or directly to an ATM switch.

ATM switches

An ATM switch is a multi-port device which forms the basic infrastructure for an ATM network. Unlike a LAN switch, an ATM switch only permits a single end station to be connected to each switch port. By interconnecting ATM switches an ATM network can be constructed to span a building, city, country, or the globe.

Table 2.11 ATM communications operating rates

Operating rate (Mbps)	Transmission media
25–31	unshielded twisted pair category 3
100	multi-mode fiber
155	shielded twisted pair
622	single-mode fiber

The basic operation of an ATM switch is to route cells from an input port onto an appropriate output port. To accomplish this the switch examines fields within each cell header and uses that information in conjunction with table information maintained in the switch to route cells. Later in this section we will examine the composition of the ATM cell header in detail.

One of the key features of ATM switches reaching the market during the late 1990s is their rate adaptation capability which in general is a function of the transmission media used to connect endpoints and to connect switches to other switches. Table 2.11 lists some of the communications rates associated with different transmission media.

ATM service processor

An ATM service processor is a computer operating software or firmware embedded in an ATM switch which performs services required for ATM network operations. For example, an ATM network address can have one of three formats, with one format similar to a telephone number. In comparison, the NIC in an IEEE 802 standardized LAN has a hardware (MAC) address burnt into the adapter. Stations on a LAN can register their addresses using the facilities of a LAN Emulation Server (LES). That server would then act as a translator between the burnt-in LAN specific hardware addresses and ATM public or private network addresses that could considerably differ from the LAN addressing scheme. The LAN Emulation process, including a description of the operation of the LES, is described in the LAN switch section in Chapter 3.

2.3.4 Network interfaces

ATM support two types of basic interface: User-to-Network Interface (UNI) and a Network-to-Node Interface (NNI).

User-to-network interface

The UNI represents the interface between an ATM switch and an ATM endpoint. As the connection of a private network to a public network is also known as a UNI, the terms Public and Private UNI were used to differentiate between the two types. That is, a Private UNI refers to the connection between an endpoint and switch on an internal private ATM network, such as an organization's ATM-based LAN. In comparison, a Public UNI would refer to the interface between either a customer's endpoint or switch and a public ATM network.

Network-to-node interface

The connection between an endpoint and switch is simpler than the connect between two switches. This results from the fact that switches communicate information concerning the utilization of their facilities as well as passing setup information required to support endpoint network requests.

The interface between switches is known as a Network-to-Node or Network-to-Network Interface (NNI). Similarly to the UNI, there are two types of NNI. A Private NNI describes the switch-to-switch interface on an internal network such as an organization's LAN. In comparison, a Public NNI describes the interface between public ATM switches, such as those used by communications carriers. Figure 2.37 illustrates the four previously described ATM network interfaces.

2.3.5 The ATM cell header

The structure of the ATM cell is identical in both public and private ATM networks, with Figure 2.38 illustrating the fields within the five byte cell header. As we will soon note, although the cell header fields are identical throughout an ATM network, the use of certain fields depends upon the interface or the presence or absence of data being transmitted by an endpoint.

Generic flow control field

The Generic Flow Control (GFC) field consists of the first four bits of the first byte of the ATM cell header. This field is used to

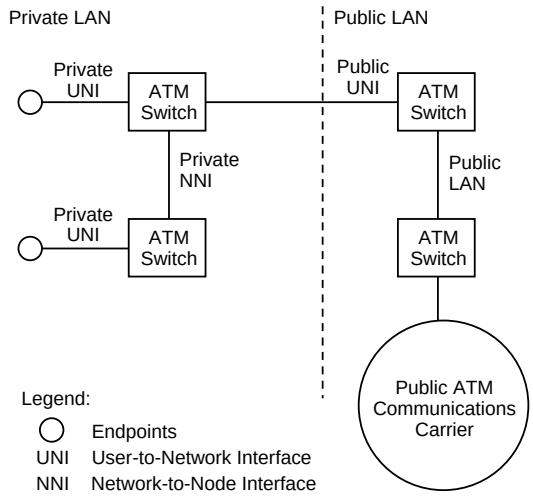

Figure 2.37 ATM network interfaces

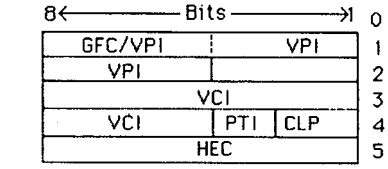

GFC Generic Flow Control
VPI Virtual Path Identifier
VCI Virtual Circuit Identifier
PTI Payload Type Identifier
CLP Cell Loss Priority
HEC Header Error Check

Figure 2.38 The ATM cell header

control the flow of traffic across the User-to-Network interface (UNI) and is used only at the UNI. When cells are transmitted between switches, the four bits become an extension of the Virtual Path Identifier (VPI) field, permitting a larger VPI value to be carried in the cell header.

Virtual path identifier field

The Virtual Path Identifier (VPI) identifies a path between two locations in an ATM network that provides transportation for a

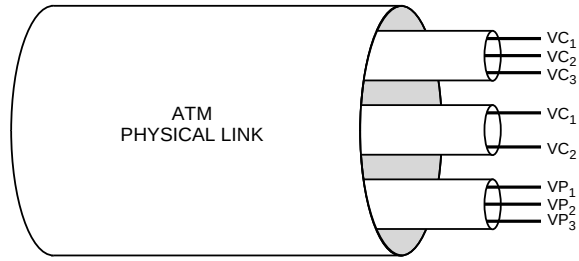

Figure 2.39 Relationship between virtual paths and virtual channels

group of virtual channels, where a virtual channel represents a connection between two communicating ATM devices. When an endpoint has no data to transmit, the VPI field is set to all zeros to indicate an idle condition. As previously explained, when transmission occurs between switches, the GFC field is used to support an extended VPI value.

Virtual channel identifier field

The Virtual Channel Identifier (VCI) can be considered to represent the second part of the two-level routing hierarchy used by ATM, where a group of virtual channels are used to form a virtual path.

Figure 2.39 illustrates the relationship between virtual paths and virtual channels. Here the virtual channel represents a connection between two communicating ATM entities, such as an endpoint to a central office switch, or between two switches. The virtual channel can represent a single ATM link or a concatenation of two or more links, with communications on the channel occurring in cell sequence order at a predefined Quality of Service. In comparison, each virtual path (VP) represents a group of VCs transported between two points that can flow over one or more ATM links. Although VCs are associated with a VP, they are not unbundled nor processed. Thus, the purpose of a virtual path is to provide a mechanism for bundling traffic routed towards the same destination. This technique enables switches to examine the VPI field within the cell header to make a decision concerning the relaying of the cell instead of having to examine the entire three byte address formed by the VPI and the VCI. When an endpoint is in an idle condition, the VPI field is set to all zeros. Although the VCI field will also be set to all zeros to

indicate the idle condition, other nonzero VCI values are reserved for use with a VPI zero value to indicate certain predefined conditions.

As we will note later in this section, all VPIs and VCIs have only local significance on a particular connection between an end point and switch or between two switches. At each switch in an ATM network, VPIs and VCIs can be remapped to different VPIs and VCIs. The actual route through an ATM network is established by signaling packets that are transmitted on a well known virtual channel, VPI=0, VCI=5.

Payload type identifier field

The Payload Type Identifier (PTI) field consists of three bits in the fourth byte of the cell header. This field is used to identify the type of information carried by the cell. Values 0–3 are reserved to identify various types of user data, 4 and 5 denote management information, while 6 and 7 are reserved for future use.

Cell loss priority field

The last bit in the fourth byte of the cell header represents the Cell Loss Priority (CLP) field. This bit is set by the AAL layer and used by the ATM layer throughout an ATM network as an indicator of the importance of the cell. If the CLP is set to 1, this indicates that the cell can be discarded by a switch experiencing congestion. If the cell should not be discarded due to the necessity to support a predefined quality of service, the AAL layer will set the CLP bit to 0. The CLP bit can also be set by the ATM layer if a connection exceeds the quality of service level agreed to during the initial communications handshaking process when setup information is exchanged.

Header error check field

The last byte in the ATM cell header is the Header Error Check (HEC). The purpose of the HEC is to provide protection for the first four bytes of the cell header against the misdelivery of cells due to errors affecting the addresses within the header. To accomplish this the HEC functions as an error detecting and

correcting code. The HEC is capable of detecting all single and certain multiple bit errors as well as correcting single bit errors.

2.3.6 ATM connections and cell switching

Now that we have a basic understanding of the ATM cell header to include the virtual path and virtual channel identifiers, we can turn our attention to the methods used to establish connections between endpoints as well as how connection identifiers are used for cell switching to route cells to their destination.

Connections

In comparison to most LANs that are connectionless, ATM is primarily a connection-oriented communications technology. This means that a connection has to be established between two ATM endpoints prior to actual data being transmitted between the endpoints. The actual ATM connection can be established as a Permanent Virtual Circuit (PVC) or as a Switched Virtual Circuit (SVC).

A PVC can be considered as being similar to a leased line, with routing established for long term use. Once a PVC has been established, no further network intervention is required any time a user wishes to transfer data between endpoints connected via a PVC. In comparison, a SVC can be considered as being similar to a telephone call made on the switched telephone network. That is, the SVC requires network intervention to establish the path linking endpoints each time a SVC occurs.

Both PVCs and SVCs obtain the V as they represent virtual rather than permanent or dedicated connections. This means that through statistical multiplexing, an endpoint can receive calls from one or more distant endpoints.

Cell switching

As previously mentioned, signaling packets are transmitted on the well-known virtual channel VPI=0, VCI=5 to make a connection. That connection results in each switch allocating a VPI and VCI to route data between switches or from a switch to

an end point. As a switch can support numerous simultaneous connections as cells arrive, the switch examines the VIP and VCI fields in the cell header to determine the output port for relaying or transferring a cell. To determine the output port, the ATM switch first reads the incoming VPI, VCI, or both fields, with the field read dependent upon the location of the switch in the network. Next, the switch will use the connection identifier information to perform a table lookup operation. That operation uses the current connection identifier as a match criterion to determine the output port that the cell will be routed onto as well as a new connection identifier to be placed into the cell header. The new connection identifier is then used for routing between the next pair of switches or from a switch to an endpoint.

Types of switch

There are two types of ATM switch, with the differences between each related to the type of header fields read for establishing cross-connections through the switch. A switch limited to reading and substituting VPI values is commonly referred to as a VP switch. This switch operates relatively fast. A switch that reads and substitutes both VPI and VCI values is commonly referred to as a Virtual Channel switch (VC Switch). A VC switch generally has a lower cell operating rate than a VP switch as it must examine additional information in each cell header. You can consider a VP switch as being similar to a central office switch, while a VC switch would be similar to end office switches.

Using connection identifiers

To illustrate the use of connection identifiers in cell switching, consider Figure 2.40, which illustrates a three switch ATM network with four endpoints. When switch 1 receives a cell from device A connected to port 2 with VPI=0, VCI=10, it uses the VPI and VCI values to perform a table lookup, assigning VPI=1, VCI=12 for the cell header and switching the cell onto port 1. Similarly, when switch 1 receives a cell on port 3 with VPI=0, VCI=18 its table lookup operation results in the assignment of VPI=1, VCI=15 to the cell's header and the forwarding of the cell onto port 1. If we assume that switch 2 is a VP switch, it only

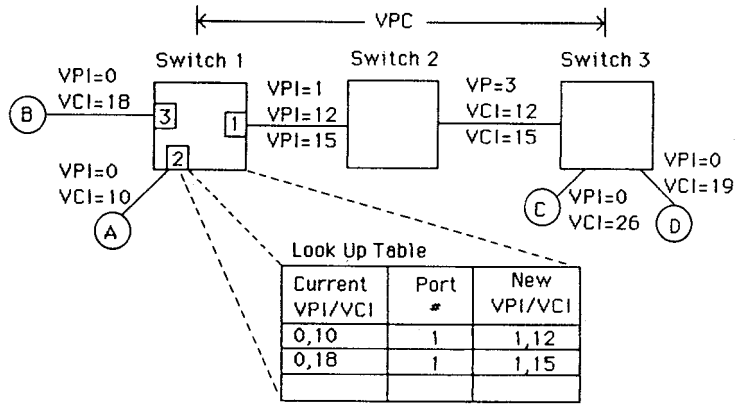

Figure 2.40 ATM cell switching example

reads and modifies the VP; thus, the VCIs are shown exiting the switch with the same values they had upon entering the switch. At switch 3, the VPI is broken down, with virtual channels assigned to route cells to endpoints C and D that were carried in a common virtual path from switch 1 to switch 3.

The assignment of VPI and VCI values is an arbitrary process which considers those already in use, with the lookup tables being created when a connection is established through the network. Concerning that connection, it results from an ATM endpoint requesting a connection setup via the User-Network Interface through the use of a signaling protocol which contains an address within the cell. That address can be in one of three formats. One known as E.164 is the same used in public telephone networks, while the other two address formats include domain identifiers that allow address fields to be assigned by different organizations. The actual signaling method is based upon the signaling protocol used in ISDN, and it enables a Quality of Service to be negotiated and agreed to during the connection setup process. The quality of service is based upon metrics assigned to different traffic classes, permitting an endpoint to establish several virtual connections where each connection transports different types of data with different performance characteristics assigned to each connection. Figure 2.41 summarizing the relationship between the five types of specified ATM traffic classes, the ATM Adaption Layer (AAL) that will support each class, the timing relationship between source and destination, and bit rate per traffic class, as well as summarizes seven metrics used to provide a Quality of Service for each traffic class.

	Traffic Classes				
	Constant Bit Rate	Variable Bit Rate Real Time	Variable Bit Rate Now Real Time	Available Bit Rate	Unspecified Bit Rate
ATM Adaption Layer	AAL 1	AAL 2	AAL 3/4, 5	AAL 3/4	Unspecified
Timing Relationship Source-Destination					
Bit Rate					
Traffic Definition Metrics	Specified	Specified	Specified	Specified	Unspecified
Cell Loss Ratio					
Cell Transfer Delay	Maximum Speficied	Maximum	Mean Specified	Unspecified	Unspecified
Cell Delay Variation	Specified	Specified	Unspecified	Unspecified	Unspecified
Peak Cell Rate	Specified	Specified	Specified	Specified	Specified
Sustainable Cell Rate	N/A	Specified	Specified	N/A	N/A
Minimum Cell Rate	N/A	N/A	N/A	Specified	N/A

Figure 2.41 Relationship of traffic classes, AAL support, and traffic definition metrics

General operation

As indicated in the beginning of this section, ATM can be used as a desktop to desktop transport mechanism or can function as a backbone to connect existing IEEE 802 networks together and into an ATM infrastructure. Figure 2.40 illustrated an example of desktop to desktop or end point to end point ATM operation. In actuality the cost of ATM adapters as well as the considerable investment previously made in Ethernet, Token-Ring, and FDDI infrastructures will preclude the widespread adoption of ATM to the desktop for many years. In the interim the scalabilty of ATM as well as its ability to move from LAN to WAN to LAN has resulted in a growing use of the technology as a backbone for interconnecting older local area networks as well as providing for a migration strategy to ATM. Readers are referred to the LAN switch section in Chapter 3 for detailed information on how ATM is used as a backbone network.