

UNIT II

Medium Access Control

The media used in LANs generally convey frames from only one station at a time, although the media themselves are generally shared by a number of stations. In order to overcome the difficulties which may arise through sharing, a **Medium Access Control (MAC)** mechanism or protocol is necessary. A MAC protocol merely regulates how stations may access the medium in an orderly fashion for correct operation and also attempts to ensure that each station obtains a fair share of its use. LAN networks usually have only a single medium over which all messages, represented over a series of frames, are transmitted. If the medium is not being used two, or more, stations may simultaneously attempt an access, leading to a collision.

An MAC technique is therefore required to regulate access by stations to the medium and handle the effect of two, or more, stations simultaneously attempting to access the medium. There is also the danger that once a pair of stations have established communication, all other stations may be excluded, perhaps indefinitely, or at least for a considerable period of time. A LAN does not usually have any separate network control function for operation. Nor is a separate control function required to detect abnormal network conditions, or to control recovery therefrom. Rather, each station is generally equally responsible in

a LAN, in which case control is said to be **fully distributed**.

Three general MAC techniques exist for use within fully distributed networks:

- 1. Contention:** Here there is no regulating mechanism directly to govern stations attempting to access a medium. Rather, two or more stations may contend for the medium and any multiple simultaneous accesses are resolved as they arise.
- 2. Token passing:** A single **token** exists within the network and is passed between

stations in turn. Only a station holding the token may use the medium for transmission. This eliminates multiple simultaneous accesses of the medium with the attendant risk of collision.

3. Slotted and register insertion rings: Similar in principle to token passing, but a unique time interval is granted to a station for transmission.

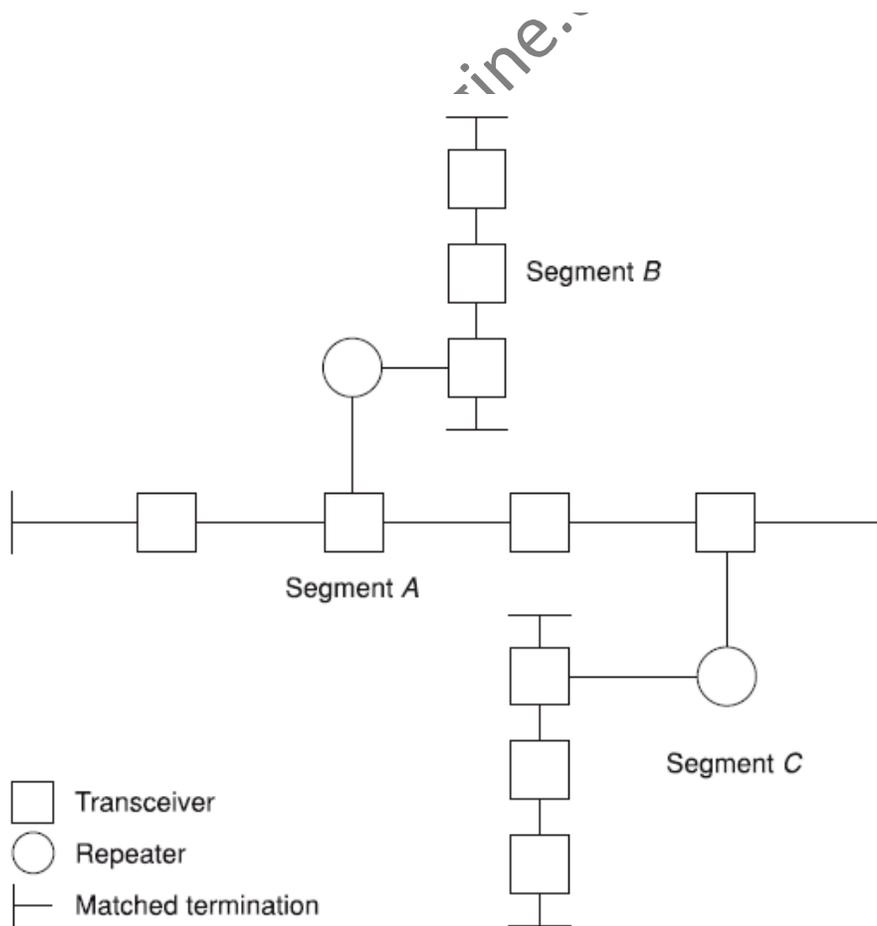
Carrier

IEEE 802.3 CSMA/CD

IEEE 802.3, which, for historical reasons, is also known as Ethernet, is the most popular type of LAN currently in use. The standard defines a range of options for the physical media, which are summarized in Table 10.1. Earlier variants employed a bus-based topology. A star-based topology is also defined which makes use of twisted-pair conductors although the network behaves as a logical bus. The basic, 10Base5 standard allows for a maximum segment length of 500 m and a total bus system not exceeding 2500 m. A **segment** is merely an unrepeated section of cable. This gives rise to the **non-rooted branching tree** structure shown in Figure. It is 'non-rooted' since no head end or master station exists. Segments

Table 10.1 IEEE 802.3 bus topology variants.

	Transmission medium	Signalling technique	Data rate (Mbps)	Maximum segment length (m)
10Base5	Coaxial cable (50 Ω)	Baseband (Manchester)	10	500
10Base2	Coaxial cable (50 Ω)	Baseband (Manchester)	10	185
10BaseT	UTP	Baseband (Baseband)	10	100
10Broad36	Coaxial cable (75 Ω)	Broadband (DPSK)	10	3600
10BaseF	Fibre	N/A	10	up to 2000



not exceeding 500 m may be interconnected via repeaters which reshape and retime data signals to overcome the attenuation and distortion introduced by the medium. In this way unacceptable signal deterioration in networks up to the maximum system length of 2500 m is avoided. However, care must be taken in building IEEE 802.3 networks comprising several segments. A further rule, known as the '5:4:3 rule', must be followed to ensure satisfactory operation. This rule states that no more than five segments may be interconnected using four repeaters and that only three of the segments may be populated with stations.

Figure indicates the frame structure used by IEEE 802.3. The preamble, in conjunction with Manchester line coding, provides a period during which the receiver may synchronize its clock with that of the incoming bit stream. Once synchronized, a receiver monitors the incoming bit stream for the unique start of frame delimiter (SFD) pattern from which the receiver may then correctly align itself with that of the incoming

Preamble	SFD	DA	SA	Length	Pad (46-0)	Data (0-1500)	FCS
(7)	(1)	(6)	(6)	(2)	(0-1500)		(4)

() Number of bytes

SFD Start of frame delimiter

SA Source address

DA Destination address

FCS Frame check sequence

Figure 10.5 IEEE 802.3 frame.

frame structure. Destination and source addresses (DA, SA) are 6 bytes long, although there exists an option for 2-byte addressing. Ethernet addresses, each of which is unique, are burnt into NIC cards during manufacture. The length field indicates to the receiver how many data bytes to expect. The data field carries the LLC PDU and therefore the length depends upon its size, with a maximum of 1500 bytes. To ensure that the frame length is at least equal to twice the

maximum propagation delay, or collision window, for reasons discussed in the previous chapter, the data and pad field must equal at least 46 bytes. Where there are less than 46 bytes of data to be transmitted within a frame sufficient pad bytes are added to ensure that the pad and fields in combination are made equal to 46 bytes.

Contention algorithm

When a packet is going by, a station interface can hear its carrier and so does not initiate a transmission of its own. When a carrier is absent a station may decide to transmit its data. However, it takes about 4ms for an electrical signal to pass from one end of a 1km coaxial cable to the other end. During this 4ms more than one station might decide to transmit data simultaneously, not knowing each other's carrier already existed on the cable. This results in a collision. The transceiver must listen to the cable when it is transmitting.

Figure 5.1. A two-segment Ethernet

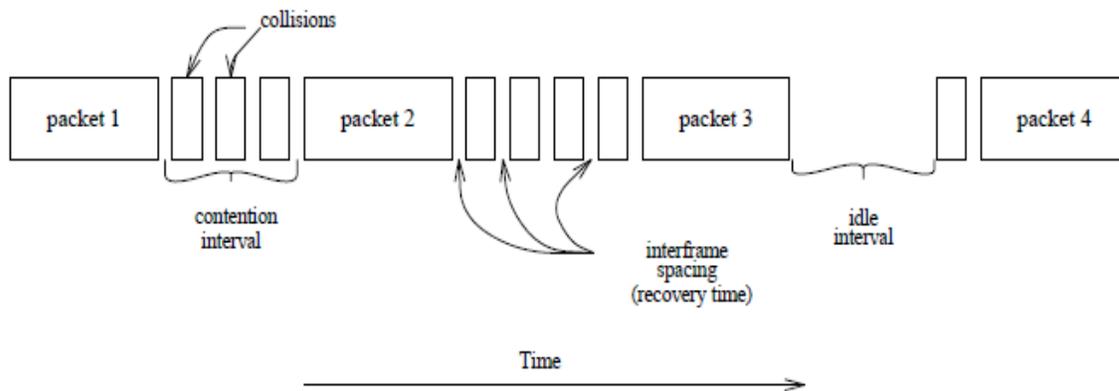
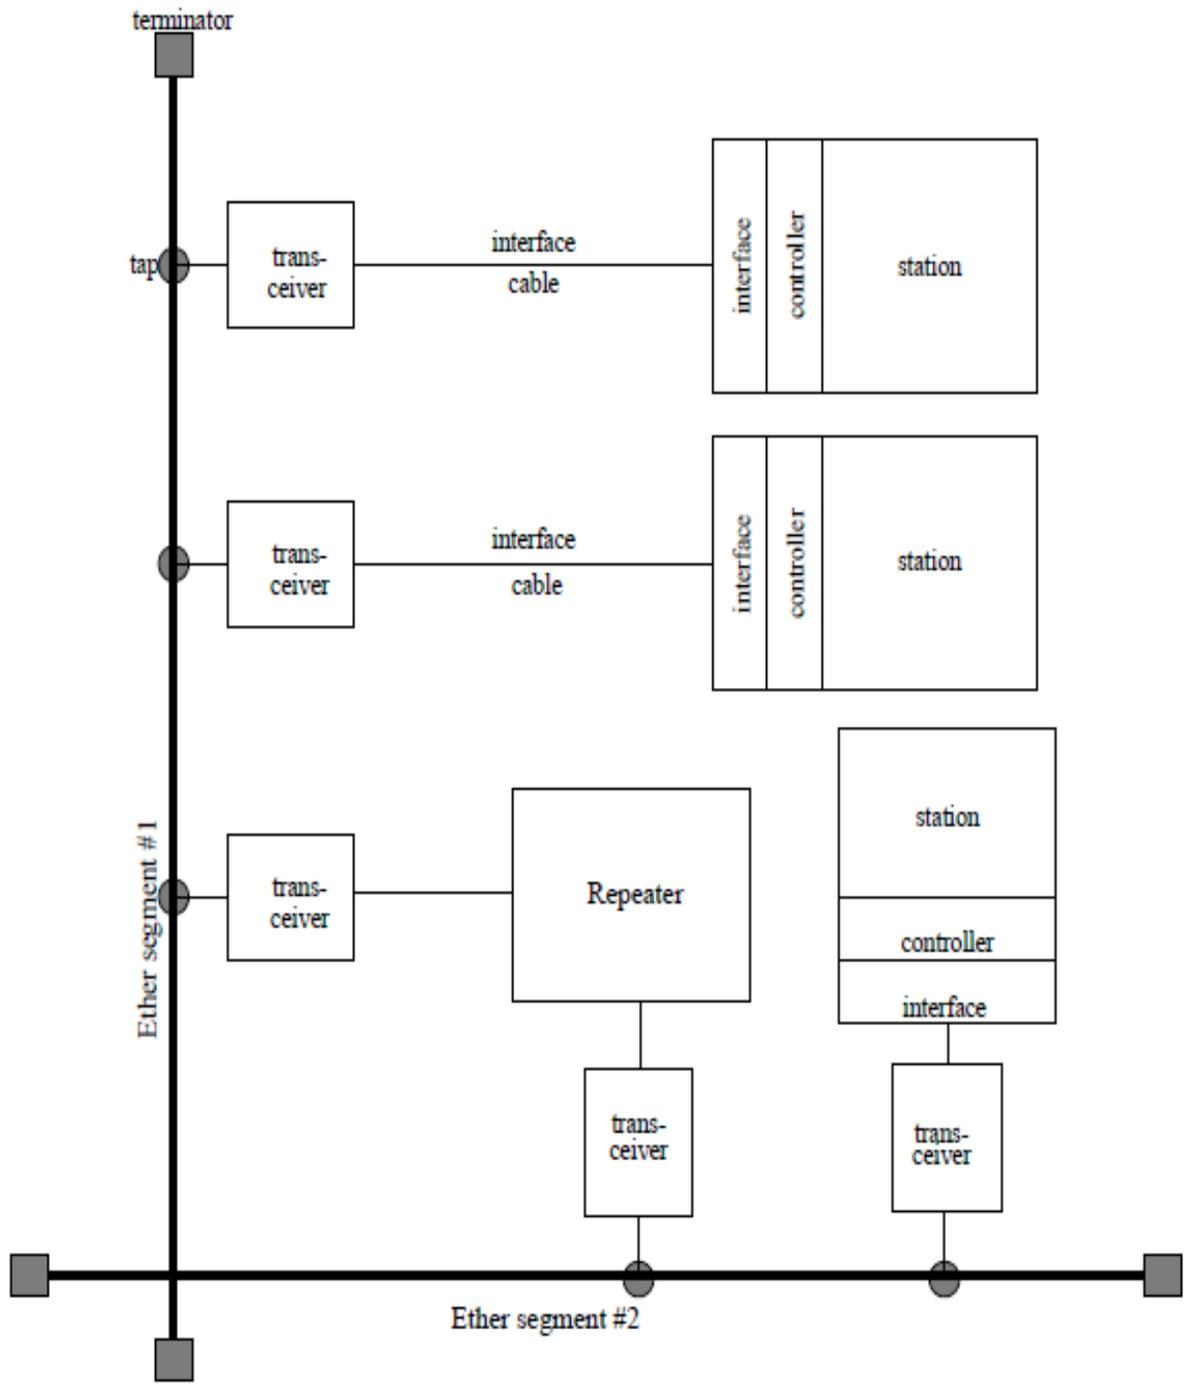
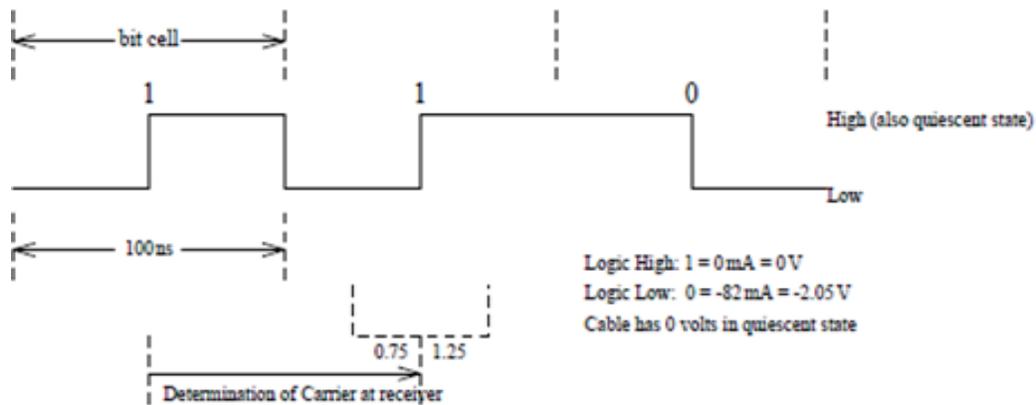


Figure 5.2. Contention based access allows any station to use the bus at any time provided it first checks to see that the bus is not in use. Collisions occur when two stations try to start at the same time.



5.1. Ethernet



MAC Frame

Figure 13.3 depicts the frame format for the 802.3 protocol; it consists of the following fields:

Preamble. A 7-octet pattern of alternating 0s and 1s used by the receiver to establish bit synchronization.

Start frame delimiter. The sequence 10101011, which indicates the actual start of the frame and which enables the receiver to locate the first bit of the rest of the frame.

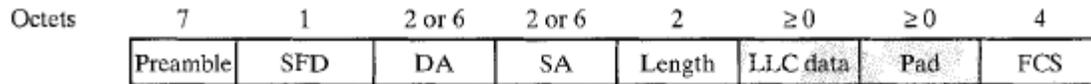
Destination address (DA). Specifies the station(s) for which the frame is intended. It may be a unique physical address, a group address, or a global address. The choice of a 16- or 48-bit address length is an implementation decision, and must be the same for all stations on a particular LAN.

Source address (SA). Specifies the station that sent the frame.

Length. Length of the LLC data field.

LLC data. Data unit supplied by LLC.

Pad. Octets added to ensure that the frame is long enough for proper CD operation.



LEGEND

SFD = Start-frame delimiter
 DA = Destination address

SA = Source address
 FCS = Frame-check sequence

FIGURE 13.3 IEEE 802.3 frame format.

Frame check sequence (FCS). A 32-bit cyclic redundancy check, based on all fields except the preamble, the SFD, and the FCS.

IEEE 802.3 10-Mbgs Specifications (Ethernet)

The IEEE 802.3 committee has been the most active in defining alternative physical configurations; this is both good and bad. On the good side, the standard has been responsive to evolving technology. On the bad side, the customer, not to mention the potential vendor, is faced with a bewildering array of options. However, the committee has been at pains to ensure that the various options can be easily integrated into a configuration that satisfies a variety of needs. Thus, the user that has a complex set of requirements may find the flexibility and variety of the 802.3 standard to be an asset.

To distinguish among the various implementations that are available, the committee has developed a concise notation:

<data rate in Mbps><signaling method><maximum segment length in hundreds of meters>

The defined alternatives are:2

- 10BASE5
- 10BASE2
- 10BASE-T
- 10BROAD36
- 10BASE-F

TABLE 13.2 IEEE 802.3 100BASE-T physical layer medium alternatives.

	100BASE-TX	100BASE-FX	100BASE-T4	
Transmission medium	2 pair, STP	2 pair, Category 5 UTP	2 optical fibers	4 pair, Category 3, 4, or 5 UTP
Signaling technique	MLT-3	MLT-3	4B5B, NRZI	8B6T, NRZ
Data rate	100 Mbps	100 Mbps	100 Mbps	100 Mbps
Maximum segment length	100 m	100 m	100 m	100 m
Network span	200 m	200 m	400 m	200 m

Token ring and FDDI

Token ring is the most commonly used MAC protocol for ring-topology LANs. In this section, we examine two standard LANs that use token ring: IEEE 802.5 and FDDI.

IEEE 802.5 Medium Access Control

MAC Protocol

The token ring technique is based on the use of a small frame, called a token, that circulates when all stations are idle. A station wishing to transmit must wait until it detects a token passing by. It then seizes the token by changing one bit in the token, which transforms it from a token into a start-of-frame sequence for a data frame. The station then appends and transmits the remainder of the fields needed to construct a data frame.

When a station seizes a token and begins to transmit a data frame, there is no token on the ring, so other stations wishing to transmit must wait. The frame on the ring will make a round trip and be absorbed by the transmitting station. The transmitting station will insert a new token on the ring when both of the following conditions have been met:

The station has completed transmission of its frame. The leading edge of the transmitted frame has returned (after a complete circulation of the ring) to the station. If the bit length of the ring is less than the frame length, the first condition implies the second; if not, a station could release a free token after it has finished transmitting but before it begins to receive its own transmission. The second condition is not strictly necessary, and is relaxed under certain circumstances. The advantage of imposing the second condition is that it ensures that only one data frame at a time may be on the ring and that only one station at a time may be transmitting, thereby simplifying error-recovery procedures.

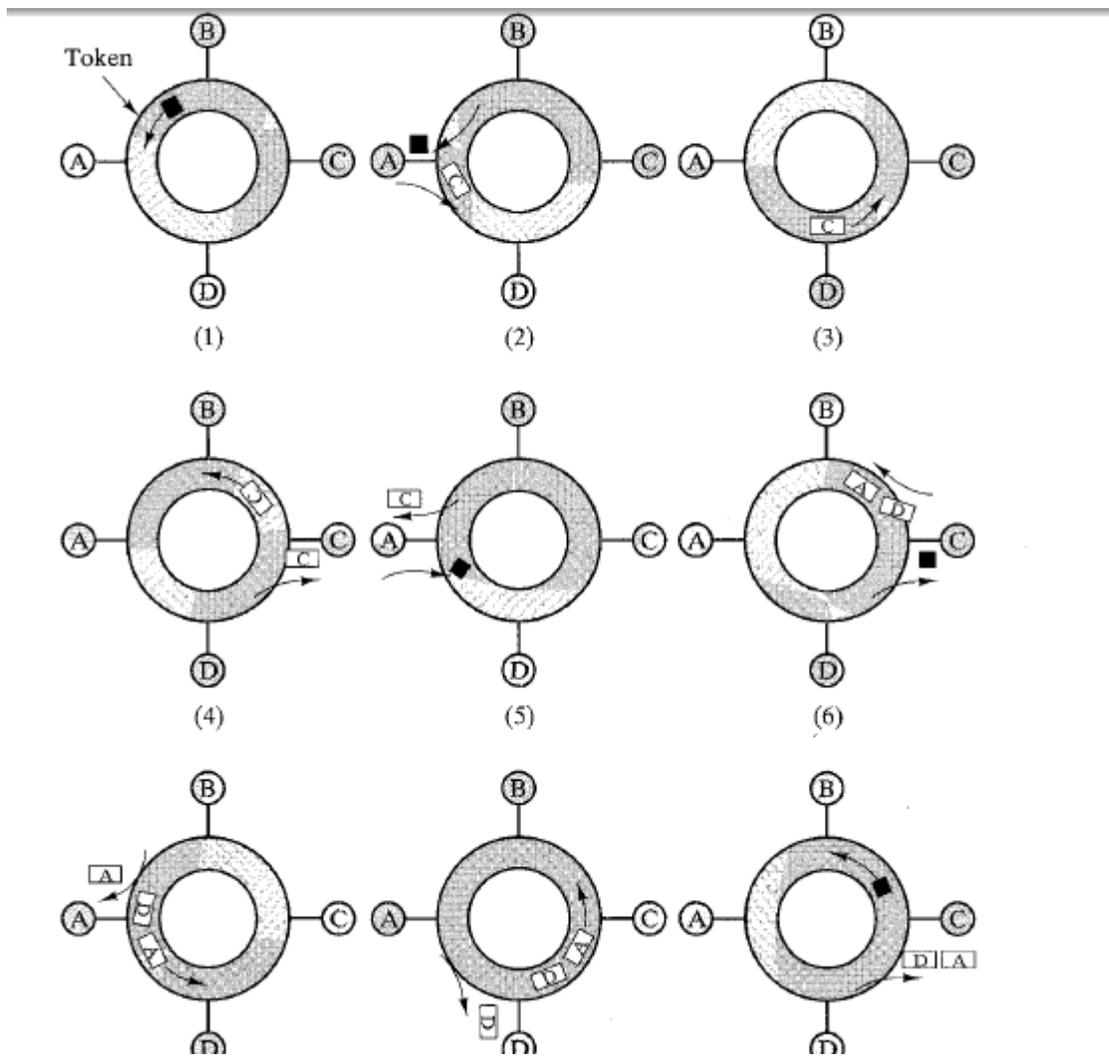
Once the new token has been inserted on the ring, the next station downstream with data to send will be able to seize the token and transmit. Figure 13.5 illustrates the technique. In the example, A sends a packet to C, which receives it and then sends its own packets to A and D.

Note that under lightly loaded conditions, there is some inefficiency with token ring because a station must wait for the token to come around before transmitting. However, under heavy loads, which is when it matters, the ring functions in a round-robin fashion, which is both efficient and fair. To see this, consider the configuration in Figure 13.5. After station A transmits, it releases a token. The first station with an opportunity to transmit is D. If D transmits, it then releases a token and C has the next opportunity, and so on.

The principal **advantage** of token ring is the flexible control over access that it provides. In the simple scheme just described, the access is fair. As we shall see, schemes can be used to regulate access to provide for priority and for guaranteed

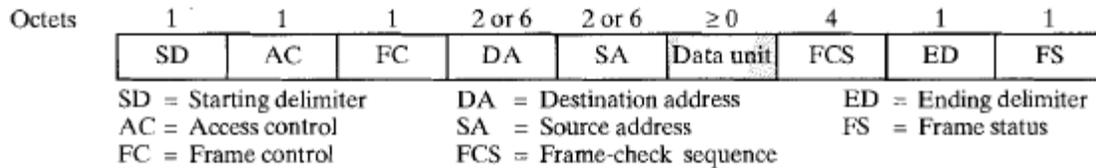
bandwidth services.

The principal **disadvantage** of token ring is the requirement for token maintenance. Loss of the token prevents further utilization of the ring. Duplication of the token can also disrupt ring operation. One station must be selected as a monitor to ensure that exactly one token is on the ring and to ensure that a free token is reinserted, if necessary.



MAC Frame

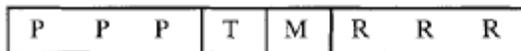
Figure depicts the frame format for the 802.5 protocol. It consists of the following fields:



(a) General frame format



(b) Token frame format



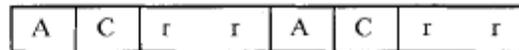
PPP = Priority bits M = Monitor bit
 T = Token bit RRR = Reservation bits

(c) Access control field



J, K = Nondata bits
 1 = Intermediate-frame bits
 E = Error-detected bits

(d) Ending delimiter field



A = Addressed recognized bit
 C = Copied bit
 r = Reserved

(e) Frame status field

FIGURE 13.6 IEEE 802.5 frame format.

Starting delimiter (SD). Indicates start of frame. The SD consists of signaling patterns that are distinguishable from data. It is coded as follows: JKOJKOOO, where J and K are nondata symbols. The actual form of a nondata symbol depends on the signal encoding on the medium.

Access control (AC). Has the format PPPTMRRR, where PPP and RRR are 3-bit priority and reservation variables, and M is the monitor bit; their use is explained below. T indicates whether this is a token or data frame. In the case of a token frame, the only remaining field is ED.

Frame control (FC). Indicates whether this is an LLC data frame. If not, bits 7 in this field control operation of the token ring MAC protocol.

Destination address (DA). As with 802.3.

SA Data unit

(c) Access control field

A = Addressed recognized bit

r A C r

Source address (SA). As with 802.3.

R

Data unit. Contains LLC data unit.

Frame check sequence (FCS). As with 802.3.

End delimiter (ED). Contains the error-detection bit (E), which is set if any repeater detects an error, and the intermediate bit (I), which is used to indicate that this is a frame other than the final one of a multiple-frame transmission. FCS

Frame status (FS). Contains the address recognized (A) and frame-copied (C) bits, whose use is explained below. Because the A and C bits are outside the scope of the FCS, they are duplicated to provide a redundancy check to detect erroneous settings.

Token Ring Priority

The 802.5 standard includes a specification for an optional priority mechanism.

Eight levels of priority are supported by providing two 3-bit fields in each data frame and token: a priority field and a reservation field. To explain the algorithm, let us define the following variables:

P_f = priority of frame to be transmitted by station

P_c = service priority: priority of current token

P_r = value of P , as contained in the last token received by this station

R_c = reservation value in current token

R_r = highest reservation value in the frames received by this station during the last token rotation

The scheme works as follows:

1. A station wishing to transmit must wait for a token with $P_c \geq P_f$.
2. While waiting, a station may reserve a future token at its priority level (P_f).

If a data frame goes by, and if the reservation field is less than its priority ($R_c < P_f$), then the station may set the reservation field of the frame to its priority ($R_c \leftarrow P_f$). If a token frame goes by, and if ($R_c < P_f$ AND $P_f < P_c$), then the station sets the reservation field of the frame to its priority ($R_c \leftarrow P_f$).

This setting has the effect of preempting any lower-priority reservation.

3. When a station seizes a token, it sets the token bit to 1 to start a data frame, sets the reservation field of the data frame to 0, and leaves the priority field unchanged (the same as that of the incoming token frame).
4. Following transmission of one or more data frames, a station issues a new token with the priority and reservation fields.

Early Token Release.

When a station issues a frame, if the bit length of the ring is less than that of the frame, the leading edge of the transmitted frame will return to the transmitting sta

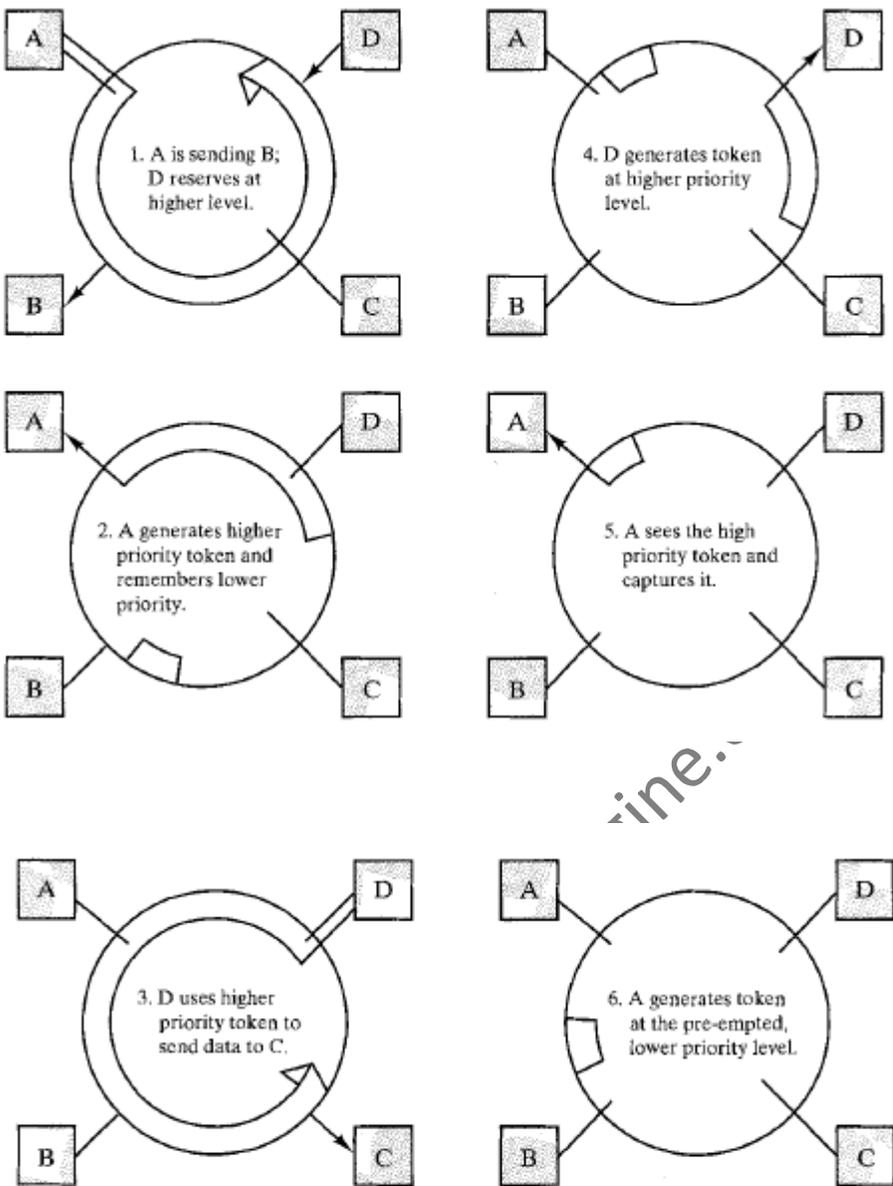


FIGURE 13.7 IEEE token ring priority scheme.

has completed transmission; in this case, the station may issue a token as soon as it has finished frame transmission. If the frame is shorter than the bit length of the ring, then after a station has completed transmission of a frame, it must wait until the leading edge of the frame returns before issuing a token. In this latter case, some of the potential capacity of the ring is unused.

To allow for more efficient ring utilization, an early token release (ETR) option has been added to the 802.5 standard. ETR allows a transmitting station to release a token as soon as it completes frame transmission, whether or not the frame header has returned to the station. The priority used for a token released prior to receipt of the previous frame header is the priority of the most recently received frame.

IEEE 802.5 Physical Layer Specification

The 802.5 standard (Table 13.4) specifies the use of shielded twisted pair with data rates of 4 and 16 Mbps using Differential Manchester encoding. An earlier specification of a 1-Mbps system has been dropped from the most recent edition of the standard. A recent addition to the standard is the use of unshielded twisted pair at 4 Mbps.

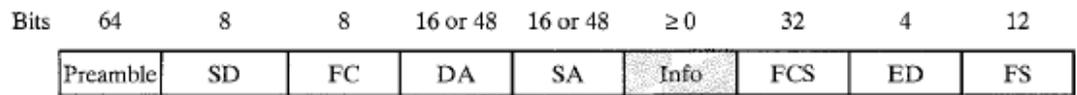
TABLE 13.4 IEEE 802.5 physical layer medium alternatives.

Transmission medium	Shielded twisted pair	Unshielded twisted pair
Data rate (Mbps)	4 or 16	4
Signaling technique	Differential Manchester	Differential Manchester
Maximum number of repeaters	250	72
Maximum length between repeaters	Not specified	Not specified

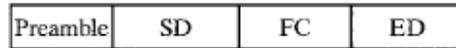
**FDDI
Medium
Access**

Control

FDDI is a token ring scheme, similar to the IEEE 802.5 specification, that is designed for both LAN and MAN applications. There are several differences that are designed to accommodate the higher data rate (100 Mbps) of FDDI.



(a) General frame format



(b) Token frame format

LEGEND

SD = Start-frame delimiter	SA = Source address	ED = Ending delimiter
FC = Frame control	FCS = Frame-check sequence	FS = Frame status
DA = Destination address		

FIGURE 13.8 FDDI frame formats.

MAC Frame

Figure 13.8 depicts the frame format for the FDDI protocol. The standard defines the contents of this format in terms of symbols, with each data symbol corresponding to 4 data bits. Symbols are used because, at the physical layer, data are encoded in 4-bit chunks. However, MAC entities, in fact, must deal with individual bits, so the discussion that follows sometimes refers to 4-bit symbols and sometime to bits. A frame other than a token frame consists of the following fields:

Preamble. Synchronizes the frame with each station's clock. The originator of the frame uses a field of 16 idle symbols (64 bits); subsequent repeating stations may change the length of the field, as consistent with clocking requirements. The idle symbol is a nondata fill pattern. The actual form of a nondata symbol depends on the signal encoding on the medium.

Starting delimiter (SD). Indicates start of frame. It is coded as JK, where J and K are nondata symbols.

Frame control (FC). Has the bit format CLFFZZZZ, where C indicates

whether this is a synchronous or asynchronous frame (explained below); L indicates the use of 16- or 48-bit addresses; FF indicates whether this is an LLC, MAC control, or reserved frame. For a control frame, the remaining 4 bits indicate the type of control frame.

Destination address (DA). Specifies the station(s) for which the frame is intended. It may be a unique physical address, a multicast-group address, or a broadcast address. The ring may contain a mixture of 16- and 48-bit address lengths.

Source address (SA). Specifies the station that sent the frame.

Information. Contains an LLC data unit or information related to a control operation.

Frame check sequence (FCS). A 32-bit cyclic redundancy check, based on the FC, DA, SA, and information fields.

Ending delimiter (ED). Contains a nondata symbol (T) and marks the end of the frame, except for the FS field.

Frame Status (FS). Contains the error detected (E), address recognized (A), and frame copied (F) indicators. Each indicator is represented by a symbol, which is R for "reset" or "false" and S for "set" or "true." A token frame consists of the following fields:

Preamble. As above.

Starting delimiter. As above.

Frame control (FC). Has the bit format 10000000 or 11000000 to indicate that this is a token.

Ending delimiter (ED). Contains a pair of nondata symbols (T) that terminate the token frame.

MAC Protocol

The basic (without capacity allocation) FDDI MAC protocol is fundamentally the same as IEEE 802.5. There are two key differences.

1. In FDDI, a station waiting for a token seizes the token by aborting (failing to repeat) the token transmission as soon as the token frame is recognized. After the captured token is completely received, the station begins transmitting one or more data frames. The 802.5 technique of flipping a bit to convert a token to the start of a data frame was considered impractical because of the high data rate of FDDI.

2. In FDDI, a station that has been transmitting data frames releases a new token as soon as it completes data frame transmission, even if it has not begun to receive its own transmission. This is the same technique as the early token release option of 802.5. Again, because of the high data rate, it would be too inefficient to require the station to wait for its frame to return, as in normal 802.5 operation.

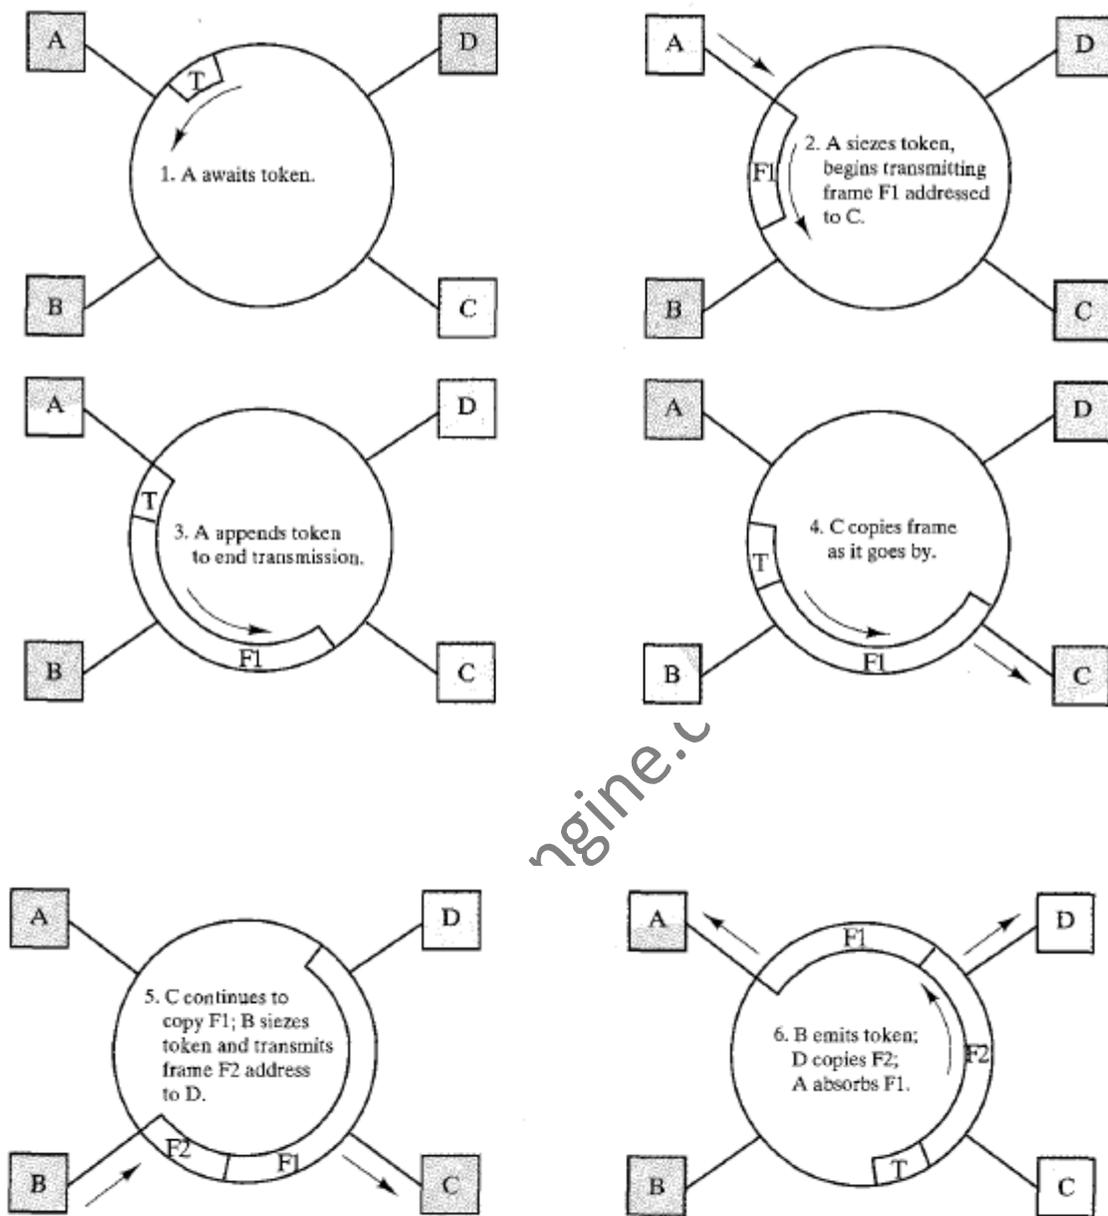


FIGURE 13.9 Example of FDDI token ring operation.

All stations have the same value of TTRT and a separately assigned value of SAI. In addition, several variables that are required for the operation of the capacity allocation algorithm are maintained at each station:

- * Token-rotation timer (TRT)
- * Token-holding timer (THT)

* Late counter (LC)

Each station is initialized with TRT set equal to TTRT and LC set to zero.'

When the timer is enabled, TRT begins to count down. If a token is received before TRT expires, TRT is reset to TTRT. If TRT counts down to 0 before a token is received, then LC is incremented to 1 and TRT is reset to TTRT and again begins to count down. IF TRT expires a second time before receiving a token, LC is incremented to 2, the token is considered lost, and a Claim process (described below) is initiated. Thus, LC records the number of times, if any, that TRT has expired since the token was last received at that station. The token is considered to have arrived early if TRT has not expired since the station received the token-that is, if LC = 0.

FDDI Physical Layer Specification

The FDDI standard specifies a ring topology operating at 100 Mbps. Two media are included (Table 13.5). The optical fiber medium uses 4B15B-NRZI encoding. Two twisted pair media are specified: 100-ohm Category 5 unshielded twisted pair⁶ and 150-ohm shielded twisted pair. For both twisted pair media, MLT-3 encoding is used.

Wireless Lan:

In just the past few years, wireless LANs have come to occupy a significant niche in the local area network market. Increasingly, organizations are finding that wireless LANs are an indispensable adjunct to traditional wired LANs, as they satisfy requirements for mobility, relocation, ad hoc networking, and coverage of locations difficult to wire.

As the name suggests, a wireless LAN is one that makes use of a wireless transmission medium. Until relatively recently, wireless LANs were little used; the reasons for this included high prices, low data rates, occupational safety concerns, and licensing requirements. As these problems have been addressed, the popularity of wireless LANs has grown rapidly.

In this section, we first look at the requirements for and advantages of wireless LANs, and then preview the key approaches to wireless LAN implementation.

Wireless LANs Applications

Four application areas for wireless LANs: LAN extension, cross building interconnect, nomadic access, and ad hoc networks. Let us consider each of these in turn.

LAN Extension

Early wireless LAN products, introduced in the late 1980s, were marketed as substitutes for traditional wired LANs. A wireless LAN saves the cost of the installation of LAN cabling and eases the task of relocation and other modifications to network structure. However, this motivation for wireless LANs was overtaken by events. First, as awareness of the need for LAN became greater, architects designed new buildings to include extensive prewiring for data applications. Second, with advances in data transmission technology, there has been an increasing reliance on twisted pair cabling for LANs and, in particular, Category 3 unshielded twisted pair. Most older buildings are already wired with an abundance of Category 3 cable. Thus, the use of a wireless LAN to replace wired LANs has not happened to any great extent.

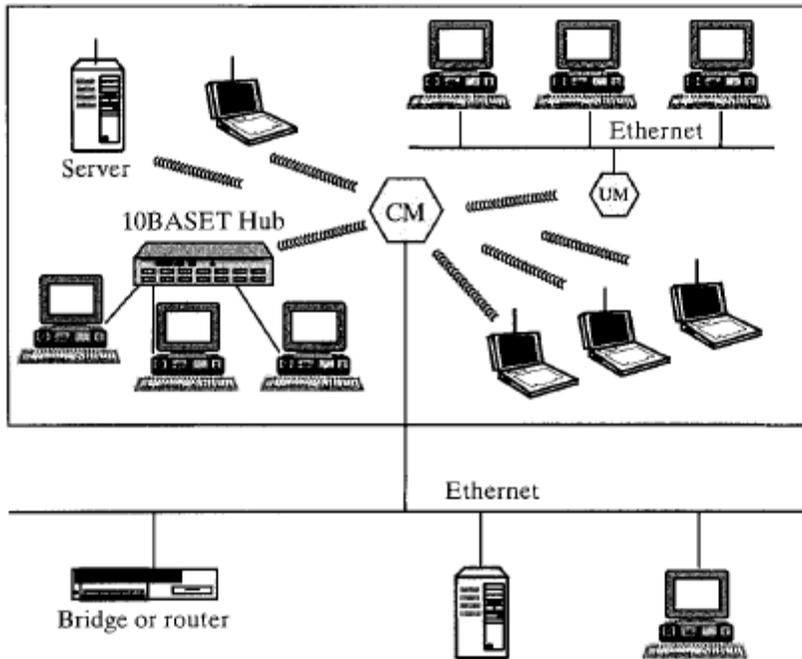


FIGURE 12.17 Example single-cell wireless LAN configuration.

Figure 12.17 indicates a simple wireless LAN configuration that is typical of many environments. There is a backbone wired LAN, such as Ethernet, that supports servers, workstations, and one or more bridges or routers to link with other networks. In addition there is a control module (CM) that acts as an interface to a wireless LAN. The control module includes either bridge or router functionality to link the wireless LAN to the backbone. In addition, it includes some sort of access control logic, such as a polling or token-passing scheme, to regulate the access from the end systems. Note that some of the end systems are standalone devices, such as a workstation or a server. In addition, hubs or other user modules (UM) that control a number of stations off a wired LAN may also be part of the wireless LAN configuration.

The configuration of Figure 12.17 can be referred to as a single-cell wireless LAN; all of the wireless end systems are within range of a single control module.

Another common configuration, suggested by Figure 12.18, is a multiple-cell wireless LAN. In this case, there are multiple control modules interconnected by a wired LAN. Each control module supports a number of wireless end systems within its transmission range. For example, with an infrared LAN, transmission is limited to a single room; therefore, one cell is needed for each room in an office building that requires wireless support

Cross-Building Interconnect

Another use of wireless LAN technology is to connect LANs in nearby buildings, be they wired or wireless LANs. In this case, a point-to-point wireless link is used between two buildings. The devices so connected are typically bridges or routers. This single point-to-point link is not a LAN per se, but it is usual to include this application under the heading of wireless LAN.

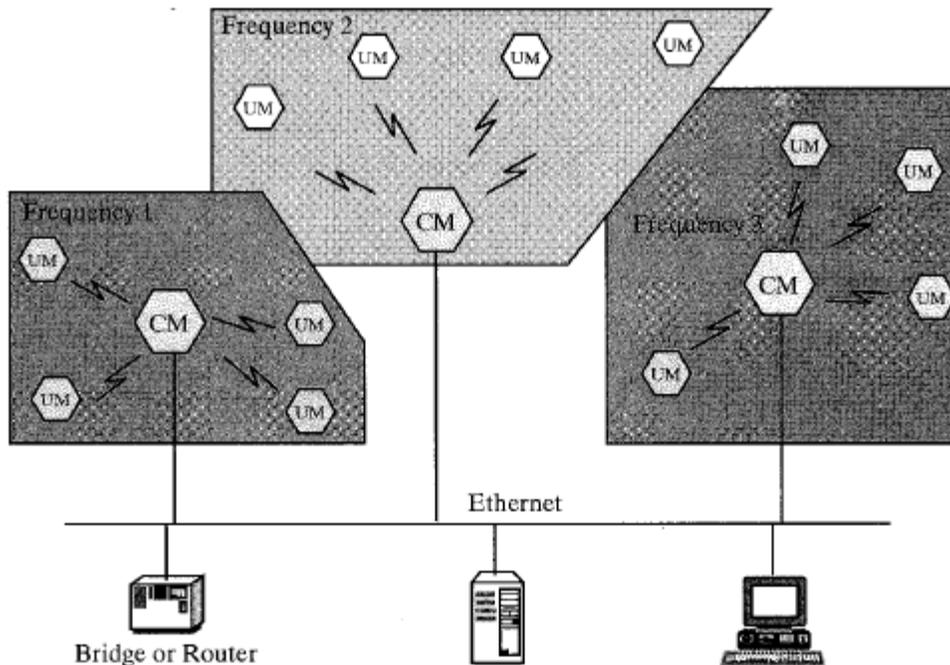


FIGURE 12.18 Example multiple-cell wireless LAN configuration.

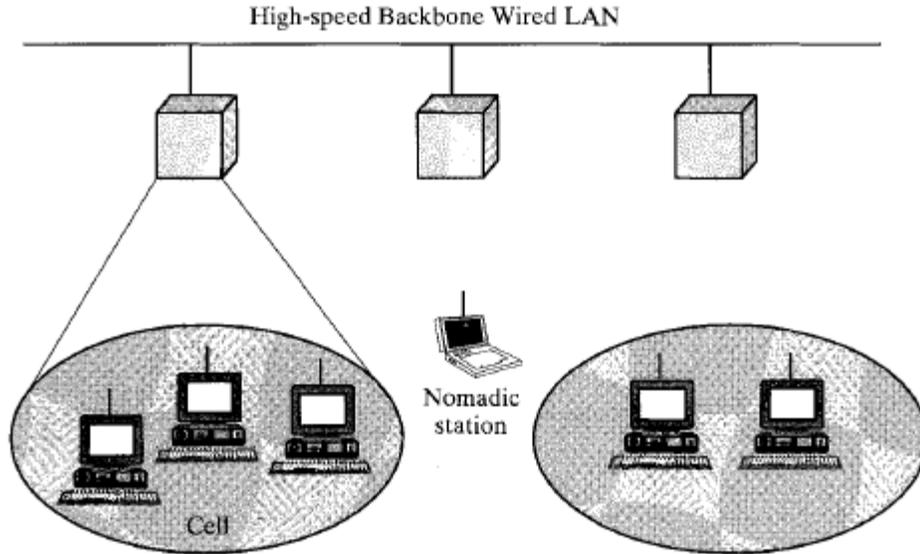
Nomadic Access

Nomadic access provides a wireless link between a LAN hub and a mobile data terminal equipped with an antenna, such as a laptop computer or notepad computer. One example of the utility of such a connection is to enable an employee returning from a trip to transfer data from a personal portable computer to a server in the office. Nomadic access is also useful in an extended environment such as a campus or a business operating out of a cluster of buildings. In both of these cases, users may move around with their portable computers and may wish access to the servers on a wired LAN from various locations.

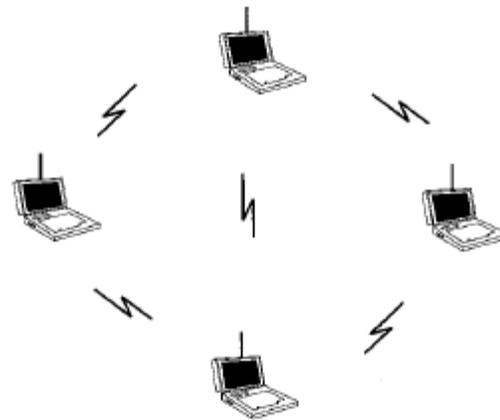
Ad Hoc Networking

An ad hoc network is a peer-to-peer network (no centralized server) set up temporarily to meet some immediate need. For example, a group of employees, each with a laptop or palmtop computer, may convene in a conference room for a business or classroom meeting. The employees link their computers in a temporary network just for the duration of the meeting.

Figure 12.19 suggests the differences between an ad hoc wireless LAN and a wireless LAN that supports LAN extension and nomadic access requirements. In the former case, the wireless LAN forms a stationary infrastructure consisting of one or more cells with a control module for each cell. Within a cell, there may be a number of stationary end systems. Nomadic stations can move from one cell to another. In contrast, there is no infrastructure for an ad hoc network. Rather, a peer collection of stations within range of each other may dynamically configure themselves into a temporary network.



(a) Infrastructure wireless LAN



(b) Ad hoc LAN

Wireless LAN Requirements

A wireless LAN must meet the same sort of requirements typical of any LAN, including high capacity, ability to cover short distances, full connectivity among attached stations, and broadcast capability. In addition, there are a number of requirements specific to the wireless LAN environment. The following are among the most important requirements for wireless LANs:

Throughput. The medium access control protocol should make as efficient use as possible of the wireless medium to maximize capacity.

Number of nodes. Wireless LANs may need to support hundreds of nodes across multiple cells.

Connection to backbone LAN. In most cases, interconnection with stations on a wired backbone LAN is required. For infrastructure wireless LANs, this is easily accomplished through the use of control modules that connect to both types of LANs. There may also need to be accommodation for mobile users and ad hoc wireless networks.

Service area. A typical coverage area for a wireless LAN may be up to a 300 to 1000 foot diameter.

Battery power consumption. Mobile workers use battery-powered workstations that need to have a long battery life when used with wireless adapters. This suggests that a MAC protocol that requires mobile nodes to constantly monitor access points or to engage in frequent handshakes with a base station is inappropriate.

Transmission robustness and security. Unless properly designed, a wireless LAN may be interference-prone and easily eavesdropped upon. The design of a wireless LAN must permit reliable transmission even in a noisy environment and should provide some level of security from eavesdropping.

" **Collocated network operation.** As wireless LANs become more popular, it is

quite likely for two of them to operate in the same area or in some area where interference between the LANs is possible. Such interference may thwart the normal operation of a MAC algorithm and may allow unauthorized access to a particular LAN.

" **License-free operation.** Users would prefer to buy and operate wireless LAN products without having to secure a license for the frequency band used by the LAN.

" **Handoff/roaming.** The MAC protocol used in the wireless LAN should enable mobile stations to move from one cell to another.

" **Dynamic configuration.** The MAC addressing and network management aspects of the LAN should permit dynamic and automated addition, deletion, and relocation of end systems without disruption to other users.

Wireless LAN Tec

Wireless LANs are generally categorized according to the transmission technique that is used. All current wireless LAN products fall into one of the following categories:

Infrared (IR) LANs. An individual cell of an IR LAN is limited to a single room, as infrared light does not penetrate opaque walls.

Spread Spectrum LANs. This type of LAN makes use of spread spectrum transmission technology. In most cases, these LANs operate in the ISM (Industrial, Scientific, and Medical) bands, so that no FCC licensing is required for their use in the U.S.

Narrowband Microwave. These LANs operate at microwave frequencies but do not use spread spectrum. Some of these products operate at frequencies that require FCC licensing, while others use one of the unlicensed ISM bands.

Bridges and switches

The early designs for bridges were intended for use between local area networks (LANs) that use identical protocols for the physical and medium access layers (e.g., all conforming to IEEE 802.3 or all conforming to FDDI). Because the devices all use the same protocols, the amount of processing required at the bridge is minimal. In recent years, bridges that operate between LANs with different MAC protocols have been developed. However, the bridge remains a much simpler device than the router, which is discussed in Chapter 16.

Because the bridge is used in a situation in which all of the LANs have the same characteristics, the reader may ask why one does not simply use one large LAN. Depending on circumstance, there are several reasons for the use of multiple LANs connected by bridges:

Reliability. The danger in connecting all data processing devices in an organization to one network is that a fault on the network may disable communication for all devices. By using bridges, the network can be partitioned into self-contained units.

Performance. In general, performance on a LAN or MAN declines with an increase in the number of devices or with the length of the medium. A number of smaller LANs will often give improved performance if devices can be clustered so that intra-network traffic significantly exceeds inter-network traffic.

Security. The establishment of multiple LANs may improve security of communications.

It is desirable to keep different types of traffic (e.g., accounting, personnel, strategic planning) that have different security needs on physically separate media. At the same time, the different types of users with different levels of security need to communicate through controlled and monitored mechanisms.

Functions of a Bridge

Figure 14.1 illustrates the operation of a bridge between two LANs, A and B. The bridge performs the following functions:

- * Reads all frames transmitted on A, and accepts those addressed to stations on B.
- * Using the medium access control protocol for B, retransmits the frames onto B.
- * Does the same for B-to-A traffic.

Several design aspects of a bridge are worth highlighting:

1. The bridge makes no modification to the content or format of the frames it receives, nor does it encapsulate them with an additional header. Each frame to be transferred is simply copied from one LAN and repeated with exactly the same bit pattern as the other LAN. Because the two LANs use the same LAN protocols, it is permissible to do this.

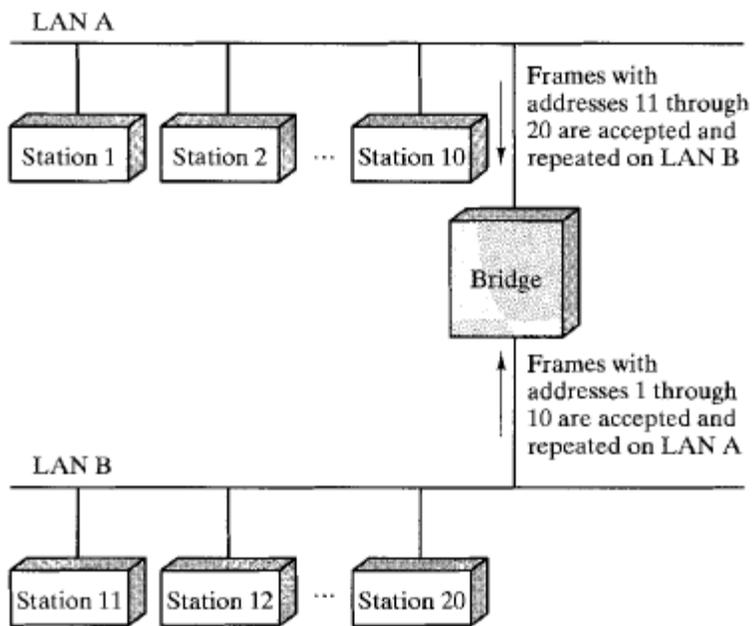


FIGURE 14.1 Bridge operation.

2. The bridge should contain enough buffer space to meet peak demands. Over a short period of time, frames may arrive faster than they can be retransmitted.
3. The bridge must contain addressing and routing intelligence. At a minimum, the bridge must know which addresses are on each network in order to know which frames to pass. Further, there may be more than two LANs interconnected by a number of bridges. In that case, a frame may have to be routed through several bridges in its journey from source to destination.
4. A bridge may connect more than two LANs

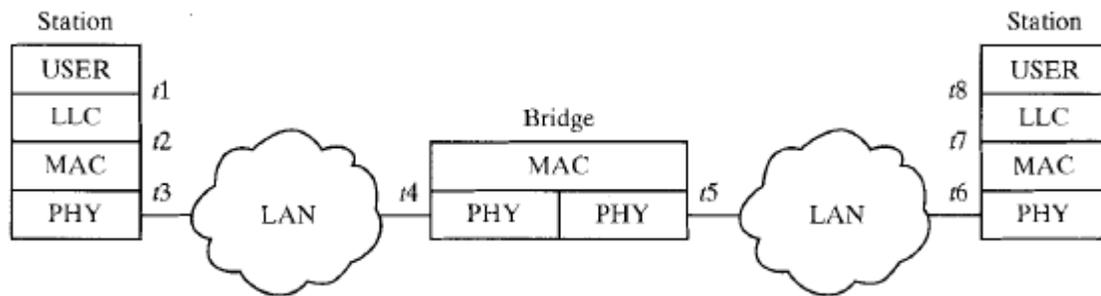
Bridge Protocol Architecture

The IEEE 802.1D specification defines the protocol architecture for MAC bridges. In addition, the standard suggests formats for a globally administered set of MAC

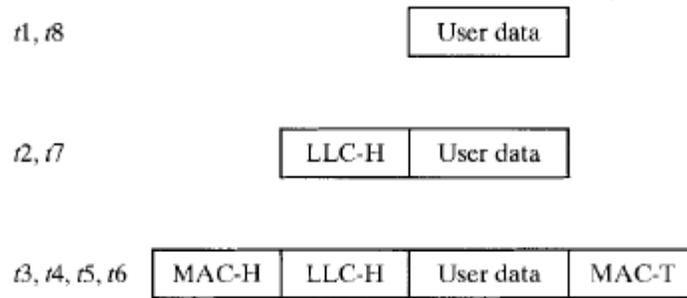
station addresses across multiple homogeneous LANs. In this subsection, we examine the protocol architecture of these bridges.

Within the 802 architecture, the endpoint or station address is designated at the MAC level. Thus, it is at the MAC level that a bridge can function. Figure 14.2 shows the simplest case, which consists of two LANs connected by a single bridge. The LANs employ the same MAC and LLC protocols. The bridge operates as previously described. A MAC frame whose destination is not on the immediate LAN is captured by the bridge, buffered briefly, and then transmitted on the other LAN. As far as the LLC layer is concerned, there is a dialogue between peer LLC entities in the two endpoint stations. The bridge need not contain an LLC layer, as it is merely serving to relay the MAC frames. Figure 14.2b indicates the way in which data is encapsulated using a bridge. Data are provided by some user to LLC. The LLC entity appends a header and

Notesengine.com



(a) Architecture

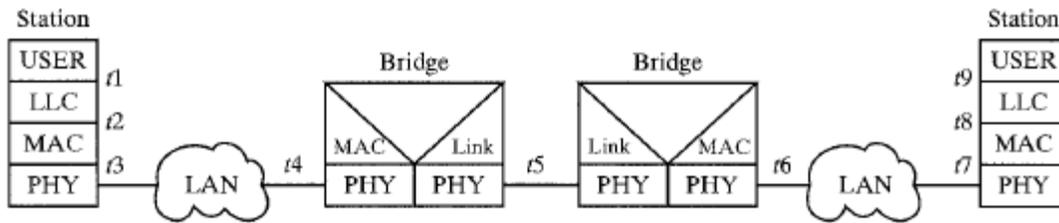


(b) Operation

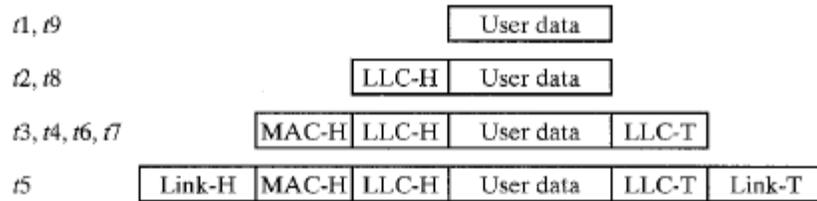
FIGURE 14.2 Connection of two LANs by a bridge.

pas

ses the resulting data unit to the MAC entity, which appends a header and a trailer to form a MAC frame. On the basis of the destination MAC address in the frame, it is captured by the bridge. The bridge does not strip off the MAC fields; its function is to relay the MAC frame intact to the destination LAN. Thus, the frame is deposited on the destination LAN and captured by the destination station.



(a) Architecture



(b) Operation

FIGURE 14.3 Bridge over a point-to-point link.

ROUTING WITH BRIDGES

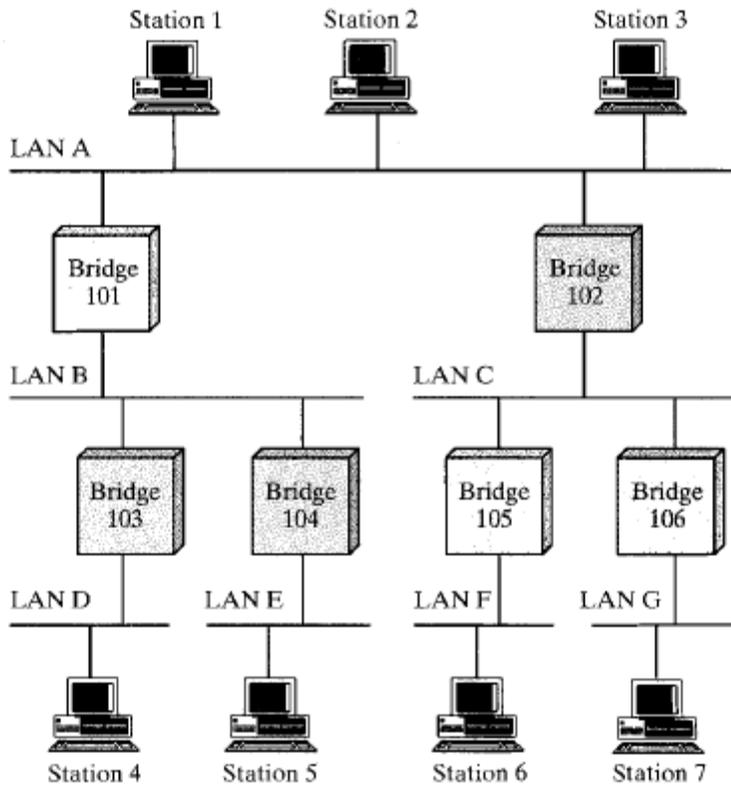


FIGURE 14.5 Configuration of bridges and LANs.

Suppose that station 1 transmits a frame on LAN A intended for station 5. The frame will be read by both bridge 101 and bridge 102. For each bridge, the addressed station is not on a LAN to which the bridge is attached. Therefore, each bridge must make a decision of whether or not to retransmit the frame on its other LAN, in order to move it closer to its intended destination. In this case, bridge 101 should repeat the frame on LAN B, whereas bridge 102 should refrain from retransmitting the frame. Once the frame has been transmitted on LAN B, it will be picked up by both bridges 103 and 104. Again, each must decide whether or not to forward the frame. In this case, bridge 104 should retransmit the frame on LAN E, where it will be received by the destination, station 5.

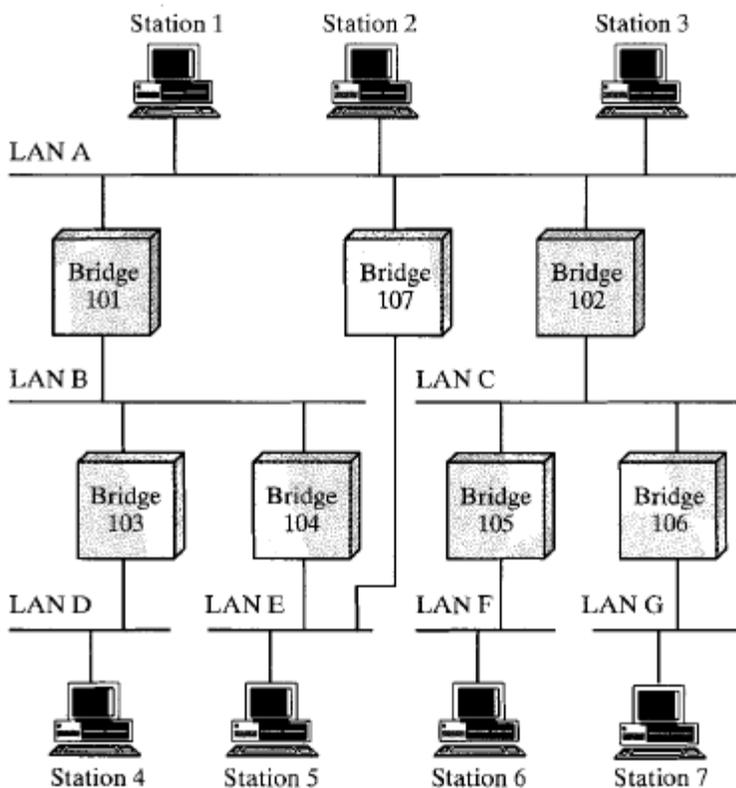


FIGURE 14.6 Configuration of bridges and LANs, with alternate routes.

A variety of routing strategies have been proposed and implemented in recent years. The simplest and most common strategy is *fixed routing*. This strategy is suitable for small LAN collections and for interconnections that are relatively stable. More recently, two groups within the IEEE 802 committee have developed specifications for routing strategies. The IEEE 802.1 group has issued a standard for routing based on the use of a *spanning tree* algorithm. The token ring committee, IEEE 802.5, has issued its own specification, referred to as *source routing*. We examine these three strategies in turn.

Fixed routing was introduced in our discussion of routing for packet-switching networks. For fixed routing with bridges, a route is selected for each source-destination pair of LANs in the internet. If alternate routes are available between two LANs, then typically the route with the least number of hops is selected. The routes are fixed, or at least only change when there is a change in the topology of the internet.

Figure 14.7 shows a fixed-routing design for the configuration of Figure 14.6. A central routing matrix shows, for each source-destination pair of LANs, the identity of the first bridge on the route. So, for example, the route from LAN E to LAN F begins by going through bridge 107 to LAN A. Again, consulting the matrix, the route from LAN A to LAN F goes through bridge 102 to LAN C. Finally, the route from LAN C to LAN F is directly through bridge 105

CENTRAL ROUTING DIRECTORY

Destination LAN	Source LAN						
	A	B	C	D	E	F	G
A	—	101	102	103	107	105	106
B	101	—	102	103	104	105	106
C	102	101	—	103	107	105	106
D	101	103	102	—	104	105	106
E	107	104	102	103	—	105	106
F	102	101	105	103	107	—	106
G	102	101	106	103	107	105	—

Bridge 101 Table
from LAN A from LAN B

Dest	Next	Dest	Next
B	B	A	A
C	—	C	A
D	B	D	—
E	—	E	—
F	—	F	A
G	—	G	A

Bridge 102 Table
from LAN A from LAN C

Dest	Next	Dest	Next
B	—	A	A
C	C	B	A
D	—	D	A
E	—	E	A
F	C	F	—
G	C	G	—

Bridge 103 Table
from LAN B from LAN D

Dest	Next	Dest	Next
A	—	A	B
C	—	B	B
D	D	C	B
E	—	E	B
F	—	F	B
G	—	G	B



Bridge 104 Table
from LAN B from LAN E

Dest	Next	Dest	Next
A	—	A	—
C	—	B	B
D	—	C	—
E	E	D	B
F	—	F	—
G	—	G	—

Bridge 105 Table
from LAN C from LAN F

Dest	Next	Dest	Next
A	—	A	C
B	—	B	C
D	—	C	C
E	—	D	C
F	F	E	C
G	—	G	C

Bridge 106 Table
from LAN C from LAN G

Dest	Next	Dest	Next
A	—	A	C
B	—	B	C
D	—	C	C
E	—	D	C
F	—	E	C
G	G	F	C

Bridge 107 Table
from LAN A from LAN E

Dest	Next	Dest	Next
B	—	A	A
C	—	B	—
D	—	C	A
E	E	D	—
F	—	F	A
G	—	G	A

FIGURE 14.7 Fixed routing (using Figure 14.6).

Spanning Tree Routing

The spanning tree approach is a mechanism in which bridges automatically develop a routing table and update that table in response to changing topology. The algorithm consists of three mechanisms: frame forwarding, address learning, and loop resolution.

Frame Forwarding

In this scheme, a bridge maintains a filtering database, which is based on MAC address. Each entry consists of a MAC individual or group address, a port number, and an aging time (described below); we can interpret this in the following fashion. A station is listed with a given port number if it is on the same side of the bridge as the port. For example, for bridge 102 of Figure 14.5, stations on LANs C, F, and G are on the same side of the bridge as the LAN C port; and stations on LANs A, B, D, and E are on the same side of the bridge as the LAN A port. When a frame is received on any port, the bridge must decide whether that frame is to be forwarded through the bridge and out through one of the bridge's other ports. Suppose that a bridge receives a MAC frame on port x. The following rules are applied (Figure

1. Search the forwarding database to determine if the MAC address is listed for any port except port x.
2. If the destination MAC address is not found, flood the frame by sending it out on all ports except the port by which it arrived.
3. If the destination address is in the forwarding database for some port y \neq x, then determine whether port y is in a blocking or a forwarding state. For reasons

explained below, a port may sometimes be blocked, which prevents it from receiving or transmitting frames.

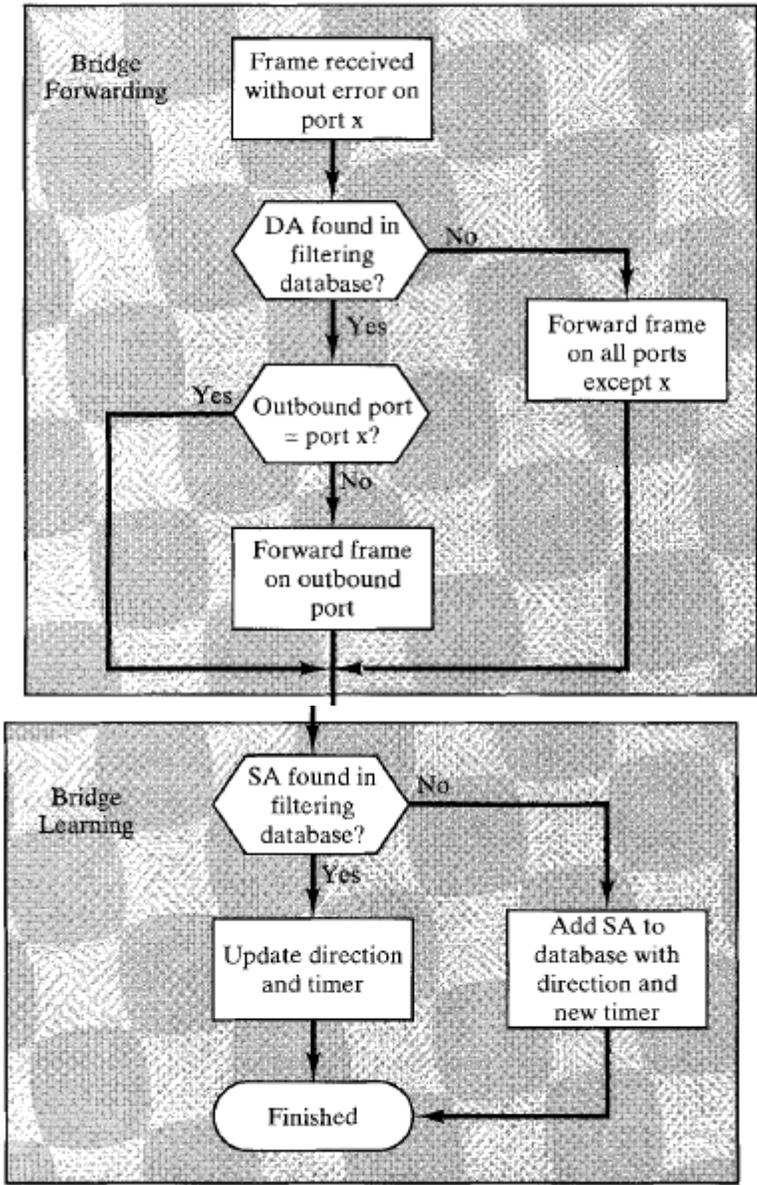


FIGURE 14.8 Bridge forwarding and learning.

Spanning Tree Algorithm

The address learning mechanism described above is effective if the topology of the network is a tree; that is, if there are no alternate routes in the network. The existence of alternate routes means that there is a closed loop.

For example in the following is a closed loop: LAN A, bridge 101, LAN B, bridge 104, LAN E, bridge 107, LAN A. To see the problem created by a closed loop, consider Figure 14.9. At time t_0 , station A transmits a frame addressed to station B. The frame is captured by both bridges. Each bridge updates its database to indicate that station A is in the direction of LAN X, and retransmits the frame on LAN Y. Say that bridge a retransmits at time t_1 and bridge P a short time later, t_2 . Thus, B will receive two copies of the frame. Furthermore, each bridge will receive the other's transmission on LAN Y.

Note that each transmission is a MAC frame with a source address of A and a destination address of B. Each bridge, then, will update its database to indicate that station A is in the direction of LAN Y. Neither bridge is capable now of forwarding a frame addressed to station A.

But the problem is potentially more serious. Assume that the two bridges do not yet know of the existence of station B. In this case, we have the following scenario. A transmits a frame addressed to B. Each bridge captures the frame. Then, each bridge, because it does not have information about B, automatically retransmits a copy of the frame on LAN Y. The frame transmitted by bridge a is captured by station

B and by bridge *P*. Because bridge *P* does not know where B is, it take this frame and retransmits it on LAN X. Similarly, bridge *a* receives bridge p's transmission on LAN Y and retransmits the frame on LAN X. There are now two

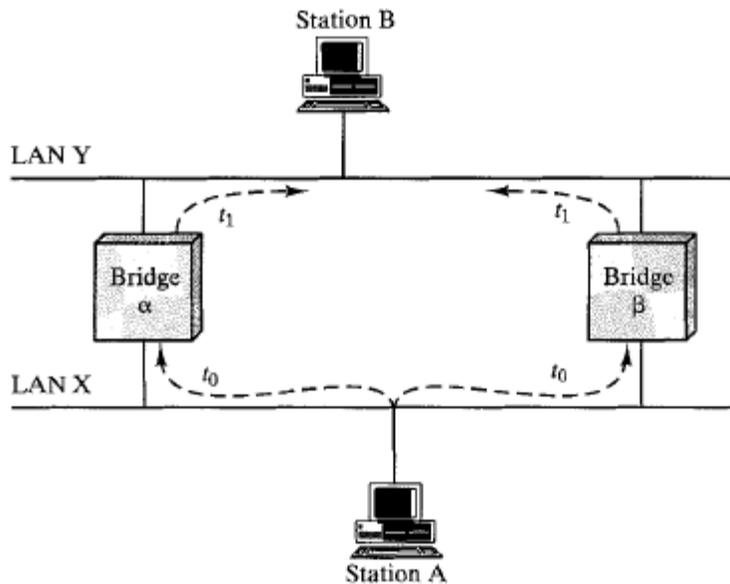


FIGURE 14.9 Loop of bridges.

frames on LAN X that will be picked up for retransmission on LAN Y. This process repeats indefinitely.

To overcome the above problem, a simple result from graph theory is used: For any connected graph, consisting of nodes and edges connecting pairs of nodes, there is a spanning tree of edges that maintains the connectivity of the graph but contains no closed loops. In terms of internets, each LAN corresponds to a graph node, and each bridge corresponds to a graph edge. Thus, in Figure 14.6, the removal of one (and only one) of bridges 107, 101, or 104, results in a spanning tree.

What is desired is to develop a simple algorithm by which the bridges of the internet can exchange sufficient information to automatically (without user intervention)

derive a spanning tree. The algorithm must be dynamic. That is, when a topology change occurs, the bridges must be able to discover this fact and automatically derive a new spanning tree.

The algorithm is based on the use of the following:

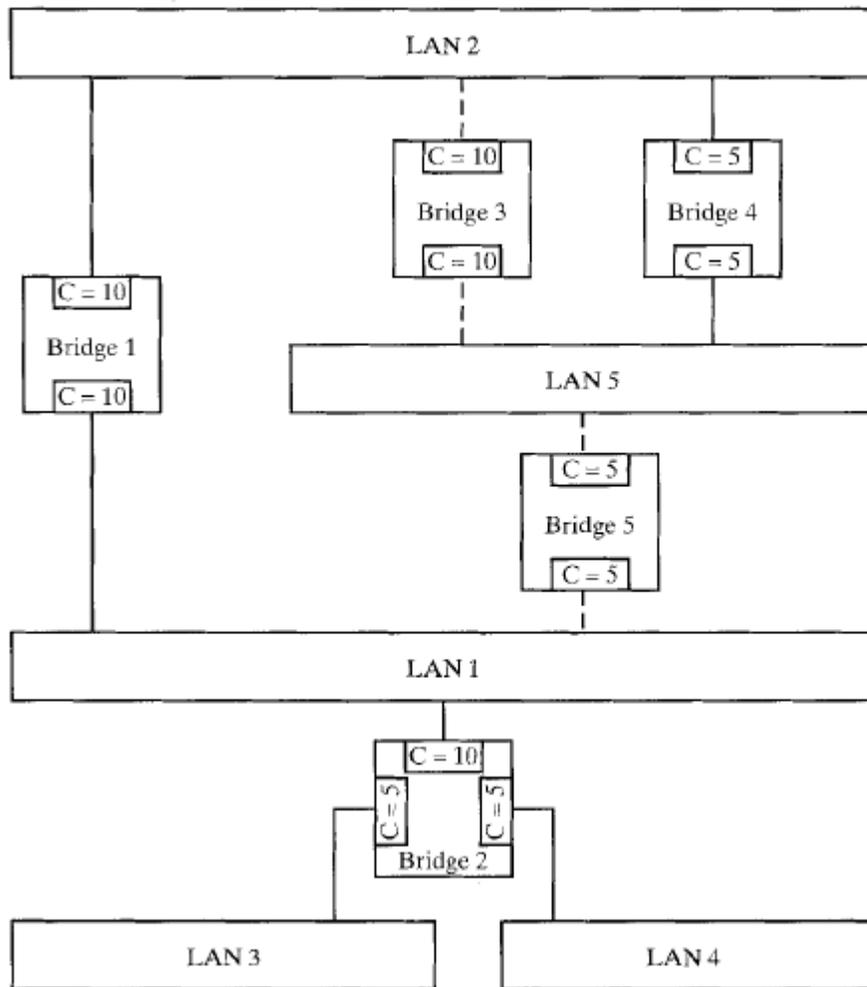
1. Each bridge is assigned a unique identifier; in essence, the identifier consists of a MAC address for the bridge plus a priority level.

2. There is a special group MAC address that means "all bridges on this LAN."

When a MAC frame is transmitted with the group address in the destination address field, all of the bridges on the LAN will capture that frame and interpret it as a frame addressed to itself.

3. Each port of a bridge is uniquely identified within the bridge, with a port *identifier*.

With this information established, the bridges are able to exchange routing information in order to determine a spanning tree of the internet. We will explain the operation of the algorithm using Figures 14.10 and 14.11 as an example. The following concepts are needed in the creation of the spanning tree:

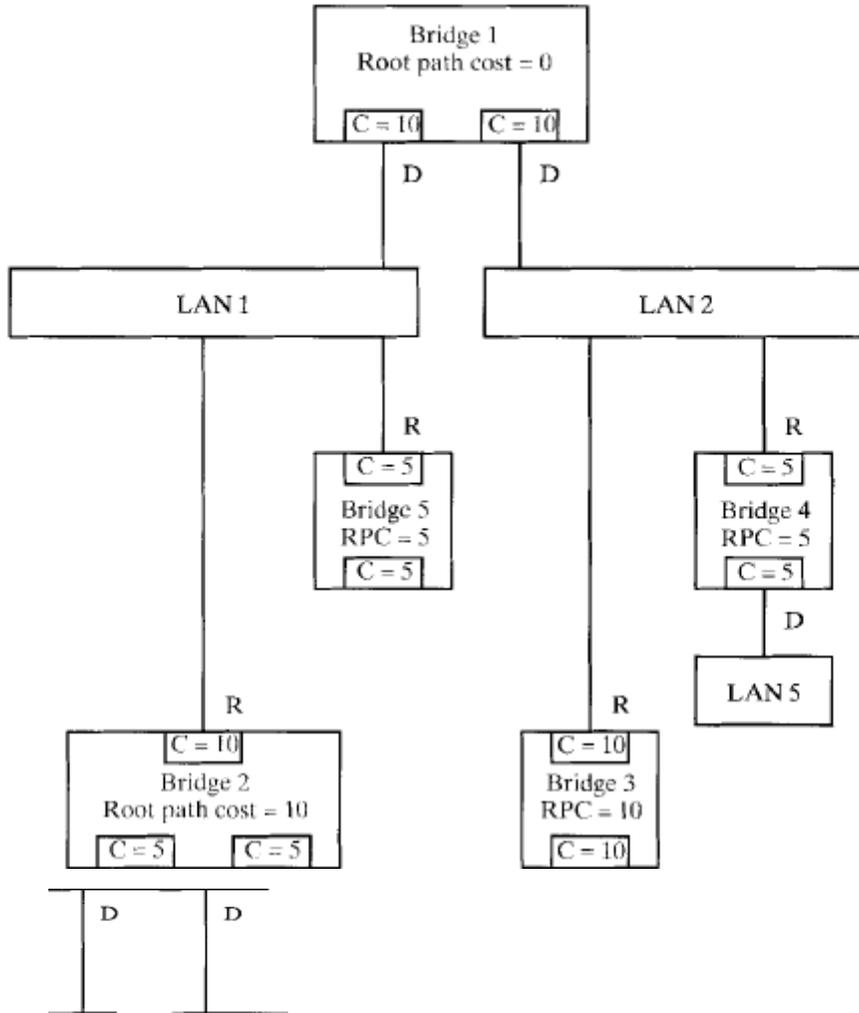


Example configuration for spanning tree algorithm.

Root bridge. The bridge with the lowest value of bridge identifier is chosen to be the root of the spanning tree.

Path cost. Associated with each port on each bridge is a path cost, which is the cost of transmitting a frame onto a LAN through that port. A path between two stations will pass through 0 or more bridges. At each bridge, the cost of transmission is added to give a total cost for a particular path. In the simplest case, all path costs would be assigned a value of 1; thus, the cost of a path would simply be a count of the number of bridges along the path. Alternatively, costs could be assigned in inverse proportion to the data rate of the corresponding LAN, or any other criterion chosen by the network manager.

Root port. Each bridge discovers the first hop on the minimum-cost path to the root bridge. The port used for that hop is labeled the root port. When the cost is equal for two ports, the lower port number is selected so that a unique spanning tree is constructed.



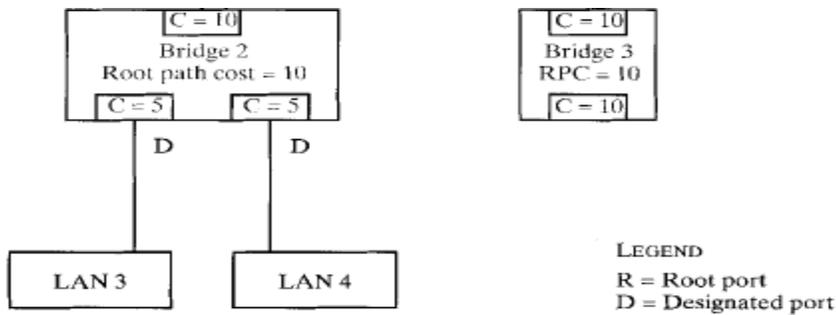


FIGURE 14.11 Spanning tree for configuration of Figure 14.10.

Root path cost. For each bridge, the cost of the path to the root bridge with minimum cost (the path that starts at the root port) is the root path cost for that bridge.

Designated bridge, designated port. On each LAN, one bridge is chosen to be the designated bridge. This is the bridge on that LAN that provides the minimum cost path to the root bridge. This is the only bridge allowed to forward frames from the LAN for which it is the designated bridge toward the root bridge. The port of the designated bridge that attaches the bridge to the LAN is the designated port. For all LANs to which the root bridge is attached, the root bridge is the designated bridge. All internet traffic to and from the LAN passes through the designated port.

1. Determine the root bridge.
2. Determine the root port on all other bridges.
3. Determine the designated port on each LAN. This will be the port with the minimum root path cost. In the case of two or more bridges with the same root path cost, the highest-priority bridge is chosen as the designated bridge. If the designated bridge has two or more ports attached to this LAN, then the port

with the lowest value of port identifier is chosen.

By this process, when two LANs are directly connected by more than one bridge, all of the bridges but one are eliminated. This cuts any loops that involve two LANs. It can be demonstrated that this process also eliminates all loops involving more than two LANs and that connectivity is preserved. Thus, this process discovers a spanning tree for the given internet. In our example, the solid lines indicate the bridge ports that participate in the spanning tree. The steps outlined above require that the bridges exchange information. The information is exchanged in the form of bridge protocol data units (BPDUs). A BPDU transmitted by one bridge is addressed to and received by all of the other bridges on the same LAN. **Each BPDU** contains the following information:

The identifier of this bridge and the port on this bridge

The identifier of the bridge that this bridge considers to be the root

The root path cost for this bridge

To begin, all bridges consider themselves to be the root bridge. Each bridge will broadcast a BPDU on each of its LANs that asserts this fact. On any given LAN, only one claimant will have the lowest-valued identifier and will maintain its belief. Over time, as BPDUs propagate, the identity of the lowest-valued bridge identifier throughout the internet will be known to all bridges. The root bridge will regularly broadcast the fact that it is the root bridge on all of the LANs to which it is attached; this allows the bridges on those LANs to determine their root port and the fact that they are directly connected to the root bridge. Each of these bridges in turn broadcasts a BPDU on the other LANs to which it is attached (all LANs except the one on its root port), indicating that it is one hop away from the root bridge. This activity is propagated throughout the internet. Every time that a bridge receives a BPDU, it transmits BPDUs, indicating the identity of the root bridge and the number of hops to reach the root bridge. On any LAN, the bridge claiming to be the one that is closest to the root becomes the designated bridge.

We can trace some of this activity with the configuration in Figure 14.10. At startup time, bridges 1,3, and 4 all transmit BPDUs on LAN 2, each claiming to be the root bridge. When bridge 3 receives the transmission from bridge 1, it recognizes a superior claimant and defers. Bridge 3 has also received a claiming BPDU from bridge 5 via LAN 5. Bridge 3 recognizes that bridge 1 has a superior claim to be the root bridge; it therefore assigns its LAN 2 port to be its root port, and sets the root path cost to 10. By similar actions, bridge 4 ends up with a root path cost of 5 via LAN 2; bridge 5 has a root path cost of 5 via LAN 1; and bridge 2 has a root path cost of 10 via LAN 2.

Source Routing

The IEEE 802.5 committee has developed a bridge routing approach referred to as source routing. With this approach, the sending station determines the route that the frame will follow and includes the routing information with the frame; bridges read the routing information to determine if they should forward the frame.

Basic Operation

The basic operation of the algorithm can be described by making reference to the configuration in Figure 14.12. A frame from station X can reach station Z by either of the following routes:

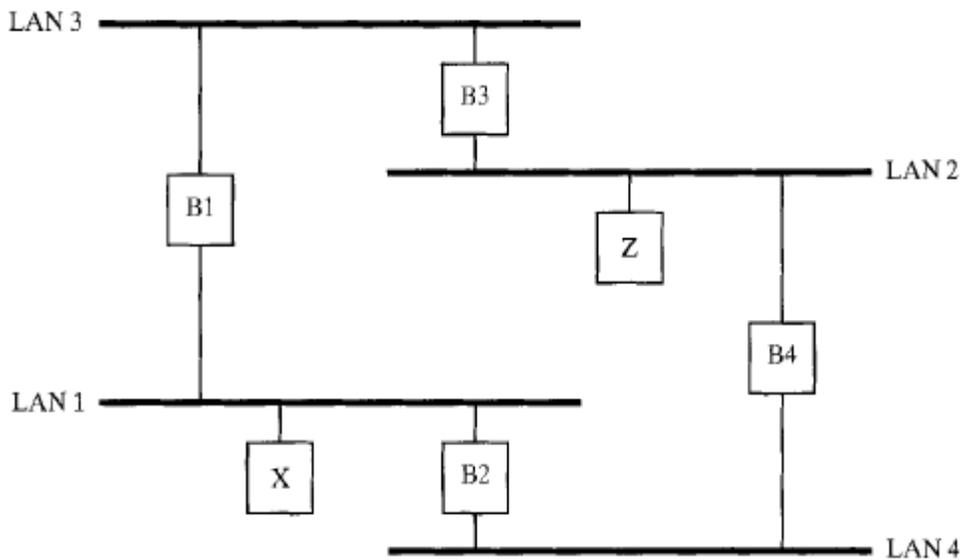
- * LAN 1, bridge B1, LAN 3, bridge B3, LAN 2
- * LAN 1, bridge B2, LAN 4, bridge B4, LAN 2

Station X may choose one of these two routes and place the information, in the form of a sequence of LAN and bridge identifiers, in the frame to be transmitted.

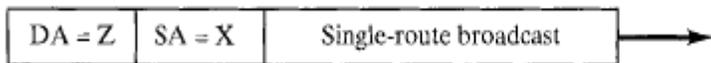
When a bridge receives a frame, it will forward that frame if the bridge is on the designated route; all other frames are discarded. In this case, if the first route above is specified, bridges B1 and B3 will forward the frame; if the second route is specified, bridges B2 and B4 will forward the frame.

Note that with this scheme, bridges need not maintain routing tables. The bridge makes the decision whether or not to forward a frame solely on the basis of the routing information contained in that frame. All that is required is that the

bridge know its own unique identifier and the identifier of each LAN to which it is attached. The responsibility for designing the route falls to the source station. For this scheme to work, there must be a mechanism by which a station can determine a route to any destination station. Before addressing this issue, we need to discuss various types of routing directives.



(a) Configuration



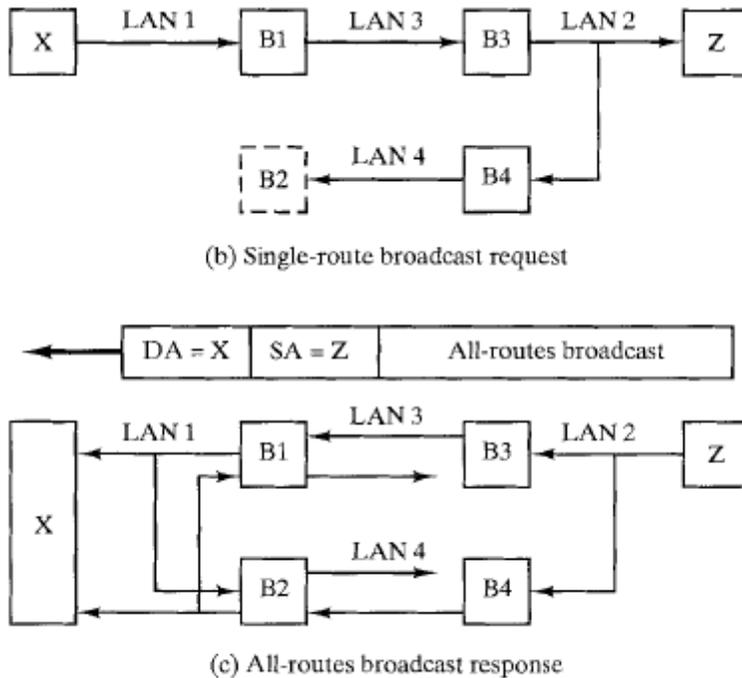


FIGURE 14.12 Route discovery example.

Routing Directives and Addressing Modes

The source routing scheme developed by the IEEE 802.5 committee includes four different types of routing directives. Each frame that is transmitted includes an indicator

of the type of routing desired. The four directive types are

- **0 Null.** No routing is desired. In this case, the frame can only be delivered to stations on the same LAN as the source station.
- **Nonbroadcast.** The frame includes a route, consisting of a sequence of LAN numbers and bridge numbers, that defines a unique route from the source station to the destination station. Only bridges on that route forward the frame, and only a single copy of the frame is delivered to the destination station.
- **All-routes broadcast.** The frame will reach each LAN of the internet by all possible routes. Thus, each bridge will forward each frame once to each of its ports in a direction away from the source node, and multiple copies of the frame may appear on a LAN. The destination station will receive one copy of

the frame for each possible route through the network.

- **Single-route broadcast.** Regardless of the destination address of the frame, the frame will appear once, and only once, on each LAN in the internet. For this effect to be achieved, the frame is forwarded by all bridges that are on a spanning tree (with the source node as the root) of the internet. The destination station receives a single copy of the frame.

Notesengine.com