but, allowing for additional gate and register transfer delays, this is normally rounded up to 1 μs per ring interface.

Frame transmission and reception

The short latency time (and hence number of bits in circulation) of the basic token ring means that a station, after initiating the transmission of an information frame, can wait until the FS field at the tail of the frame has been received before transmitting a new token without any noticeable loss in ring utilization. However, from Example 8.4 we can deduce that with an FDDI ring the loss in ring utilization will be significant if this mode of operation is adopted. With an FDDI ring, therefore, the early token release method is utilized; that is, a new token is transmitted immediately after the station has transmitted the FS symbol at the tail of a frame. The station then follows the token with IDLE symbols until it receives the SD symbols indicating the start of a new frame or token. The basic scheme is shown in Figure 8.24.

As we can see, as with the basic token ring, the source station removes a frame after it has circulated the ring. However, because of the long latency of an FDDI ring, more than one frame may be circulating around the ring at one time. Although not shown in Figure 8.24, the ring interface must repeat the SD, FC, and DA fields (symbols) of any received frames before it can determine if its own address is in the SA field. This can result in one or more frame fragments – comprising SD, FC, and DA fields – circulating around the ring. This means that a station, on receipt of the token, starts to transmit a waiting frame and concurrently receives and discards any frame fragments that may be circulating around the ring.

Timed token rotation protocol

Unlike the transmission control method used with the basic token ring – which is based on the use of the priority and reservation bits in the access control field of token and information frames – an FDDI ring uses a scheme that is based on a timed token rotation protocol that is controlled by a preset parameter known as the **target token rotation time (TTRT)**.

For each rotation of the token, each station computes the time that has expired since it last received the token. This is known as the **token rotation time** (**TRT**) and includes the time taken by a station to transmit any waiting frames plus the time taken by all other stations in the ring to transmit any waiting frames for this rotation of the token. Clearly, if the ring is lightly loaded then the TRT is short but, as the loading of the ring increases, so the TRT measured by each station increases also. The TRT is thus a measure of the total ring loading.

The timed token rotation protocol ensures access to the ring is shared fairly between all stations by allowing a station to send any waiting frames only if its measured TRT is less than the preset TTRT for the ring. Hence, when a station has frames waiting to send, on receipt of the token it computes the difference between the TTRT and its current TRT. This is known as the **token hold time** (**THT**) since it determines for how long the station may transmit

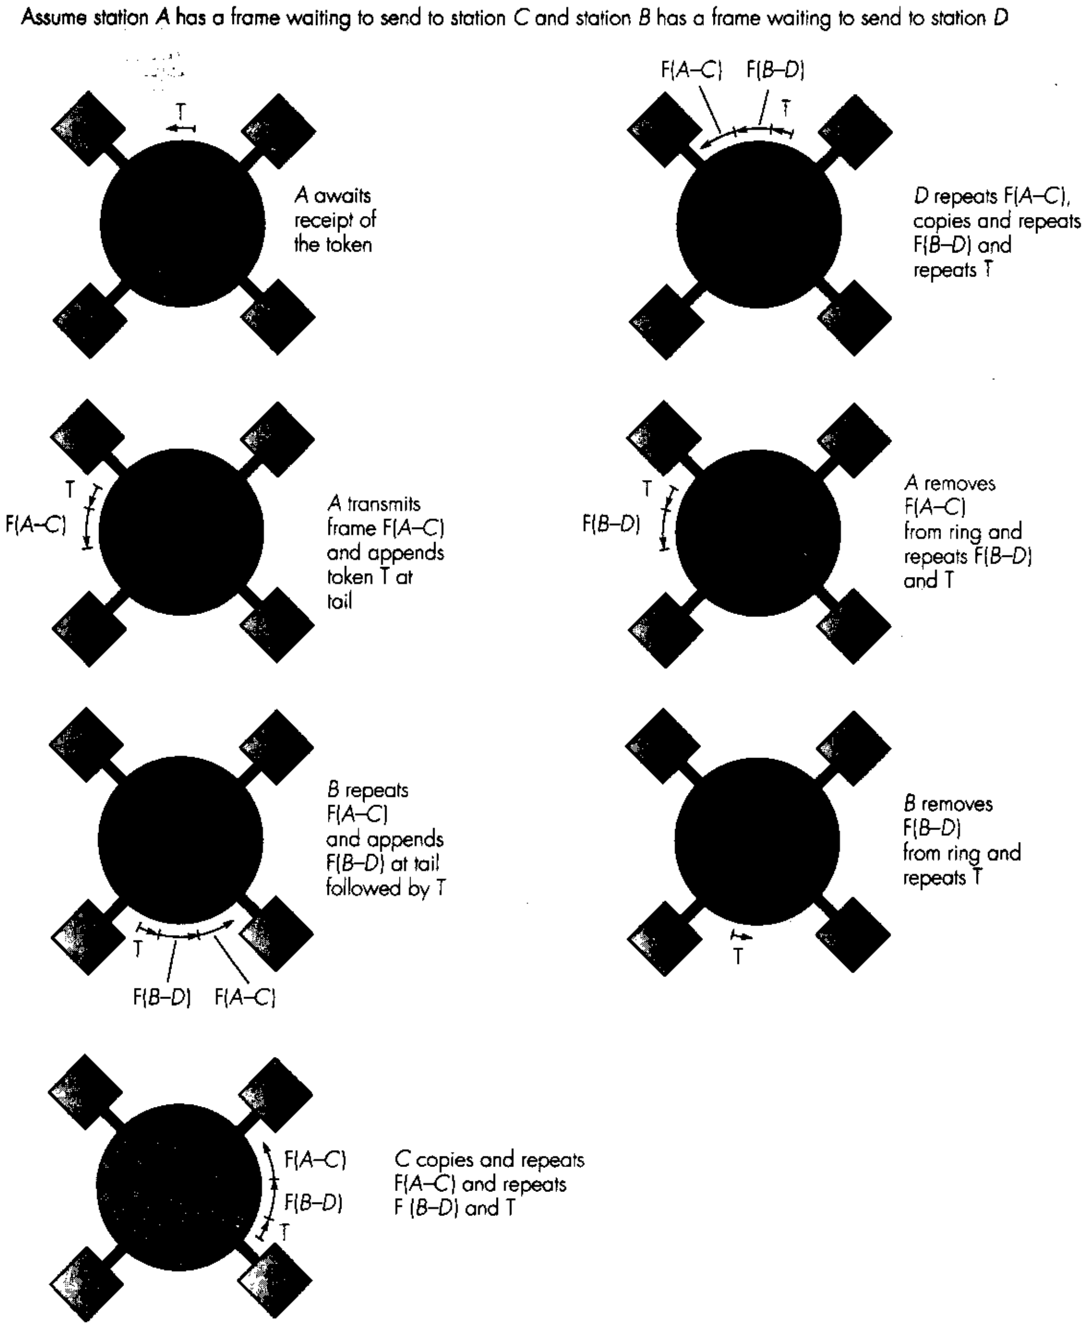

Figure 8.24 FDDI transmission example.

waiting frames before releasing the token. If the THT is positive, then the station can transmit for up to this interval. If it is negative, then the station must forego transmitting any waiting frames for this rotation of the token. A positive THT is thus known as an **early token** while a negative THT is known as a **late token**.

Example 8.5

[redacted]

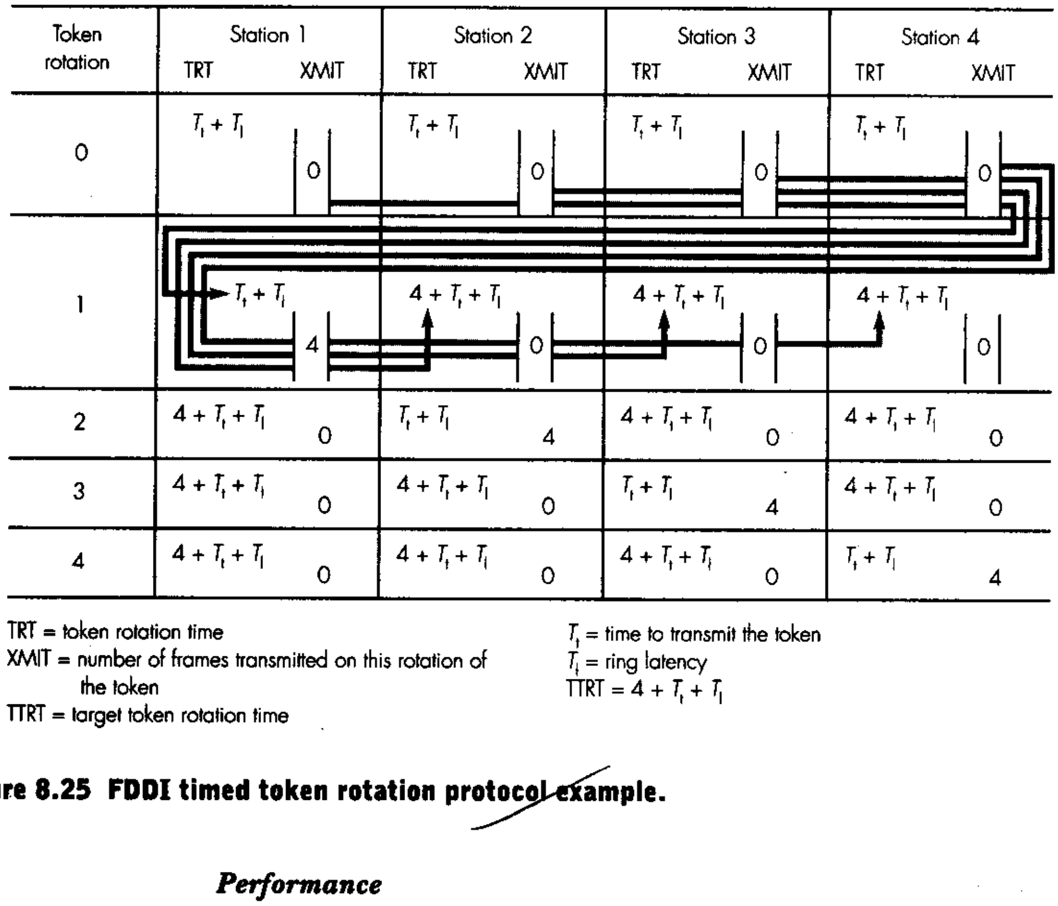

TRT = token rotation time
XMIT = number of frames transmitted on this rotation of the token
TTRT = target token rotation time
T_t = time to transmit the token
T_l = ring latency
TTRT = $4 + T_t + T_l$

Figure 8.25 FDDI timed token rotation protocol example.

Performance

There are two important performance measures associated with shared-medium networks: the **maximum obtainable throughput** and the **maximum access delay**. Both are strongly influenced by the MAC algorithm used to share the available transmission capacity between the attached stations. As we have just seen, in the case of an FDDI ring this is based on the use of a control token and the timed token rotation protocol. The important parameter associated with the protocol is the TTRT and, since this must be preset to the same value in all stations, it is important to quantify its effect on both the obtainable throughput and the access delay.

Although the nominal throughput of an FDDI ring is 100 Mbps, because of the ring latency and access control mechanism, the maximum obtainable throughput is less than this. Maximum obtained throughput (and access delay) implies that there are always frames waiting to be sent on receipt of the token at each ring interface. Under such conditions, we can compute the obtainable maximum throughput by considering the operating scenario in Example 8.5. In this, on receipt of the token, the first station transmits a set of frames up to the point when the TTRT expires before passing on the token. All the other stations in the ring are blocked from transmitting any frames during this rotation of the token. It is not until the token is passed to the next station in the ring on the following rotation that the second transmission burst of frames occurs.

As we can deduce from this, the transmission time lost during successive rotations of the token is made up of two parts: the time lost while the set of frames transmitted by the first station rotate around the ring – the ring latency – and the time taken for the token to be passed to the next station on the ring. We can express the maximum utilization of the nominal ring capacity, U_{max}, as follows:

$$U_{max} = \frac{\mathrm{TTRT} - T_l}{(\mathrm{TTRT} - T_l) + T_l + T_t + (T_l / n)} = \frac{n\,(\mathrm{TTRT} - T_l)}{n\,(\mathrm{TTRT} - T_l) + (n+1)\,T_l + T_t}$$

where TTRT is the target token rotation time, T_l is the ring latency, T_t is the time to transmit the token, and n is the number of stations in the ring.

In practice, the TTRT is very much greater than T_t. Hence the maximum utilization can be approximated to:

$$U_{max} = \frac{n\,(\mathrm{TTRT} - T_l)}{n\mathrm{TTRT} - T_l} \quad \text{or approximately,} \quad \frac{\mathrm{TTRT} - T_l}{\mathrm{TTRT}}$$

We can conclude that, to achieve a high level of ring utilization, we must select a TTRT which is significantly greater than the total ring latency. In addition, the selected TTRT must allow at least one frame of the maximum size to be present on the ring. The maximum size frame is 4500 bytes which requires 0.36 ms to transmit at 100 Mbps. The maximum ring latency was derived in Example 8.4 and, allowing for the secondary ring to be in use and the 500 stations to be all dual-attach stations, this is approximately 2 ms. Therefore the minimum TTRT allowable is 2.36 ms which, allowing for a safety margin, is set to 4 ms in the standard.

The access delay is defined as the time delay between the arrival of a frame at the ring interface of the source station and its delivery by the ring interface at the destination station. Thus access delay includes any time spent waiting in the ring interface queue at the source until a usable token – that is, an early token – arrives. It is meaningful only if the offered load is less than the maximum obtainable throughput of the ring, otherwise the interface queues continuously build up and the access delay gets progressively larger.

Providing the offered load is less than the maximum obtainable throughput, the maximum (worst case) access delay can also be deduced from Example 8.5. Recall that all stations simultaneously receive a set of frames at their ring interface queues such that they will each utilize the full TTRT quota for each rotation of the token. Assuming that station 1 is the first station able to transmit frames, the maximum access delay will be experienced by the last station on the ring – that is, station 4 – since this will receive a usable token only at the start of the fourth rotation of the token. Also, if the four frames in its ring interface queue are all addressed to station 3, then a further delay equal to the ring latency will be experienced while they are all transmitted around the ring to this station. The general expression used to compute the maximum access delay of an FDDI ring, A_{max}, is given by:

$$\begin{aligned} A_{max} &= (n-1)\ (\mathrm{TTRT} - T_l) + nT_l + T_l \\ &= (n-1)\ \mathrm{TTRT} + 2T_l \end{aligned}$$

where A_{max} is the worst-case time to receive a usable token and the other terms have the same meaning as before.

Clearly, since the ring latency is fixed for a particular ring configuration, the larger the ring TTRT, the larger the maximum access delay. As Example 8.6 shows, for all but the largest networks, the maximum obtainable throughput of an FDDI ring can be obtained with a TTRT equal to the minimum value of 4 ms. Hence although in the standard the maximum TTRT can be as high as 165 ms, it is common to operate the ring with a TTRT significantly less than this.

Example 8.6

Derive the maximum obtainable [illegible]

(i) 2 km ring with 20 stations,

(ii) [illegible]

(iii) 100 km ring with 500 stations.

The three ring configurations [illegible] and hence the same computed ring latency [illegible]

(a) Maximum available throughput [illegible]

(b) Maximum access delay [illegible]

(i) From Example 8.4, [illegible] Hence:

[illegible]

8.6 Continued

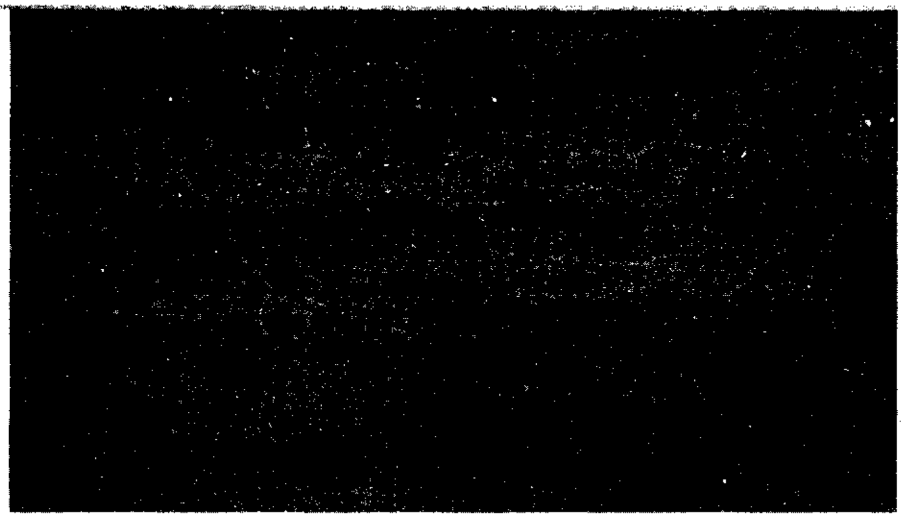

8.7 High-speed LANs

As the application of LANs has become more diverse, so the demands on them in terms of information/data throughput have increased. As we have just described, by using a combination of bridges and a high bit rate backbone, the throughput of the total LAN is determined by the maximum throughput of each LAN segment. As we explained in Sections 8.3 and 8.4, the maximum throughput of the two basic LAN types is only a fraction of the bit rate that is used. Hence in order to meet the higher throughput requirements of the newer multimedia applications, a number of high-speed LAN types have been developed. These include variations of the basic Ethernet LAN and, since this is by far the most widely installed type of LAN, we shall restrict our discussion to three variations of Ethernet: (shared) **Fast Ethernet, Switched Fast Ethernet**, and **Gigabit Ethernet**.

8.7.1 Fast Ethernet

The aim of Fast Ethernet was to use the same shared, half-duplex transmission mode as Ethernet but to obtain a $\times 10$ increase in operational bit rate over 10BaseT while at the same time retaining the same wiring systems, MAC method, and frame format. As we explained in Section 8.3, when using hubs with unshielded twisted-pair (UTP) cable, the maximum length of drop cable from the hub to a station is limited to 100 m by the driver/receiver electronics. Assuming just a single hub, this means that the maximum distance between any two stations is 200 m and the worst-case path length for collision detection purposes is 400 m plus the repeater delay in the hub. Clearly, therefore, a higher bit rate can be used while still retaining the same CSMA/CD MAC method and minimum frame size of 512 bits. In the standard, the bit rate is set at 100 Mbps over existing UTP cable. Hence the standard is also known as **100BaseT**.

Line code

The major technological hurdle to overcome with Fast Ethernet was how to achieve a bit rate of 100 Mbps over 100 m of UTP cable. Category 3 UTP cable – as used for telephony, and the most widely installed – contains four separate twisted-pair wires. To reduce the bit rate used on each pair, all four pairs are used to achieve the required bit rate of 100 Mbps in each direction.

With the CSMA/CD access control method, in the absence of contention for the medium, all transmissions are half-duplex, that is, either station-to-hub or hub-to-station. In a 10BaseT installation, just two of the four wire pairs are used for data transfers, one in each direction. Collisions are detected when the transmitting station (or hub) detects a signal on the receive pair while it is transmitting on the transmit pair. Since the collision detect function must also be performed in 100BaseT, the same two pairs are used for this function. The remaining two pairs are operated in a bidirectional mode, as shown in Figure 8.26(a).

The figure shows that data transfers in each direction utilize three pairs – pairs 1, 3, and 4 for transmissions between a station and the hub and pairs 2, 3, and 4 for transmissions between the hub and a station. Transmissions on pairs 1 and 2 are then used for collision detection and carrier sense purposes as with 10BaseT. This means that the bit rate on each pair of wires need only be 100/3 = 33.33 Mbps.

If we use Manchester encoding, a bit rate of 33.33 Mbps requires a baud rate of 33.33 Mbaud which exceeds the 30Mbaud limit set for use with such cables, as above this, unacceptably high levels of crosstalk are obtained. To reduce the baud rate, a 3-level (**ternary**) **code** is used instead of straight (2-level) binary coding. The code used is known as **8B6T** which means that, prior to transmission, each set of 8 binary bits is first converted into 6 ternary (3-level) symbols. From the example shown in Figure 8.26(b), we can deduce that this yields a symbol rate of:

$$\frac{100 \times 6/8}{3} = 25\,\text{Mbaud}$$

which is well within the set limit.

The three signal levels used are $+V$, 0, $-V$ which are represented simply as +, 0, –. The codewords are selected such that the line is DC balanced, that is, the mean line signal is zero. This maximizes the receiver's discrimination of the three signal levels since these are then always relative to a constant 0 (DC) level. To achieve this, we exploit the inherent redundancy present in the use of 6 ternary symbols. The 6 ternary symbols means that there are 729 (3^6)

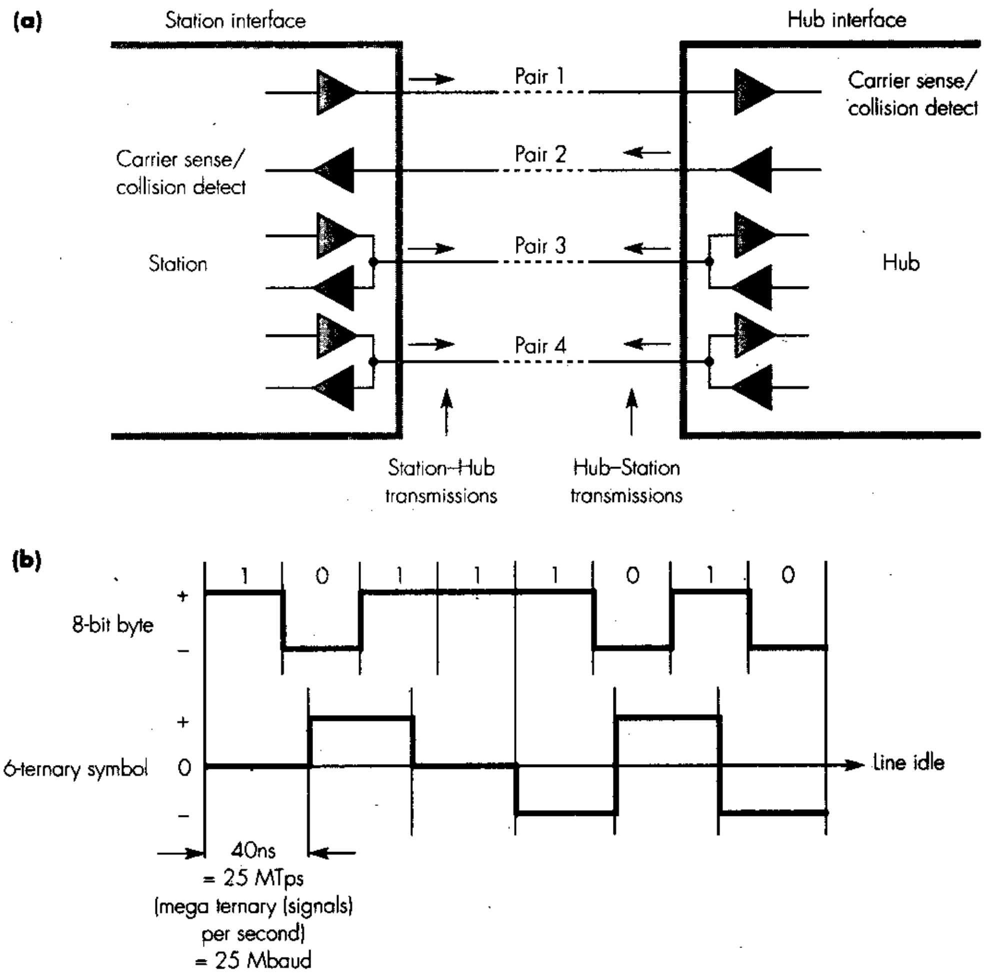

Figure 8.26 100 Base T: (a) use of wire pairs; (b) 8B6T encoding.

possible codewords. Since only 256 codewords are required to represent the complete set of 8-bit byte combinations, the codes used are selected, firstly, to achieve DC balance and secondly, to ensure all codewords have at least two signal transitions within them. This is done so that the receiver DPLL maintains clock synchronization.

To satisfy the first condition, we choose only those codewords with a combined weight of 0 or +1 and 267 codes meet this condition. To satisfy the second condition, we eliminate those codes with fewer than two transitions – five codes – and also those starting or ending with four consecutive zeros – six codes. This leaves the required 256 codewords, the first 128 of which are listed in Table 8.2.

Table 8.2 First 128 codewords of 8B6T codeword set.

Data byte	Codeword	Data byte	Codeword	Data byte	Codeword	Data byte	Codeword
00	-+00-+	20	-++-00	40	-00+0+	60	0++0-0
01	0-+-+0	21	+00+--	41	0-00++	61	+0+-00
02	[illegible]	22	-+0-++	42	0-0+0+	62	+0+0-0
03	0-++0-	23	+-0-++	43	0-0++0	63	+0+00-
04	-+0+0-	24	+-0+00	44	-00++0	64	0++00-
05	+0-+0-	25	-+0+00	45	00-0++	65	++0-00
06	+0-0+-	26	+00-00	46	00-+0+	66	++00-0
07	[illegible]	27	-+++--	47	00-++0	67	++000-
08	-+00+-	28	0++-0-	48	00+000	68	0++-+-
09	0-++-0	29	+0+0--	49	++-000	69	+0++--
0A	0-+0+-	2A	+0+-0-	4A	+-+000	6A	+0+-+-
0B	0-+-0+	2B	+0+--0	4B	-++000	6B	+0+--+
0C	-+0-0+	2C	0++--0	4C	0+-000	6C	0++--+
0D	+0-+-0	2D	++00--	4D	+0-000	6D	++0+--
0E	+0-0+-	2E	++0-0-	4E	0-+000	6E	++0-+-
0F	+0--0+	2F	++0--0	4F	-0+000	6F	++0--+
10	0--+0+	30	+-00-+	50	+--+0+	70	000++-
11	-0-0++	31	0+--+0	51	-+-0++	71	000+-+
12	-0-+0+	32	0+-0-+	52	-+-+0+	72	000-++
13	-0-++0	33	0+-+0-	53	-+-++0	73	000+00
14	0--++0	34	+-0+0-	54	+--++0	74	000+0-
15	--00++	35	-0+-+0	55	--+0++	75	000+-0
16	--0+0+	36	-0+0-+	56	--++0+	76	000-0+
17	--0++0	37	-0++0-	57	--+++0	77	000-+0
18	-+0-+0	38	+-00+-	58	--0+++	78	+++--0
19	+-0-+0	39	0+-+-0	59	-0-+++	79	+++-0-
1A	-++-+0	3A	0+-0+-	5A	0--+++	7A	+++0--
1B	+00-+0	3B	0+--0+	5B	0--0++	7B	0++0--
1C	+00+-0	3C	+-0-0+	5C	+--0++	7C	-00-++
1D	[illegible]	3D	-0++-0	5D	-000++	7D	-00+00
1E	[illegible]	3E	-0+0+-	5E	0+++--	7E	+---++
1F	[illegible]	3F	-0+-0+	5F	0++-00	7F	+--+00

DC balance

As we have just indicated, all the codewords selected have a combined weight of either 0 or +1. For example, the codeword + – – +00 has a combined weight of 0 while the codeword 0+++– – has a weight of +1. Clearly, if a string of codewords each of weight +1 is transmitted, then the mean signal level at the receiver will move away rapidly from the zero level, causing the signal to be misinterpreted. This is known as **DC wander** and is caused by the use of transformers at each end of the line. The presence of transformers means there is no path for direct current (DC).

To overcome this, whenever a string of codewords with a weight of +1 is to be sent, the symbols in alternate codewords are inverted prior to transmission. For example, if a string comprising the codeword 0 + + + – – is to be sent, then the actual codewords transmitted will be 0+ + + – – , 0 – – – + + , 0 + + + – – , 0 – – – + + , and so on, yielding a mean signal level of 0. At the receiver, the same rules are applied and the alternative codewords will be reinverted into their original form prior to decoding. The procedure used for transmission is shown in the state transition diagram in Figure 8.27(a).

To reduce the latency during the decoding process, the 6 ternary symbols corresponding to each encoded byte are transmitted on the appropriate three wire pairs in the sequence shown in Figure 8.27(b). This means that the sequence of symbols received on each pair can be decoded independently. Also, the frame can be processed immediately after the last symbol is received.

End-of-frame sequence

The transmission procedure adopted enables further error checking to be added to the basic CRC. We can deduce from the state transition diagram in Figure 8.27(a) that the running sum of the weights is always either 0 or +1. At the end of each frame transmission – that is, after the four CRC bytes have been transmitted – one of two different **end-of-stream** (**EOS**) codes is transmitted on each of the three pairs. The code selected effectively forms a checksum for that pair. The principle of the scheme is shown in Figure 8.27(c).

In this figure, we assume the last of the four CRC bytes (CRC-4) is on pair 3. The next codeword transmitted on pair 4 is determined by whether the running sum of the weights on that pair – referred to as the checksum – is 0 or +1. The EOS function is complete at the end of this codeword and the length of the other two EOS codes are reduced by two or one times the latency, that is, 4T or 2T. This means the receiver can detect reliably the end of a frame since all signals should cease within a short time of one another. This allows for very small variations in propagation delay on each pair of wires.

Collision detection

An example station–hub transmission without contention is shown in Figure 8.28(a). Recall that a station detects a collision by detecting a signal on pair 2 while it is transmitting and, similarly, the hub detects a collision by the

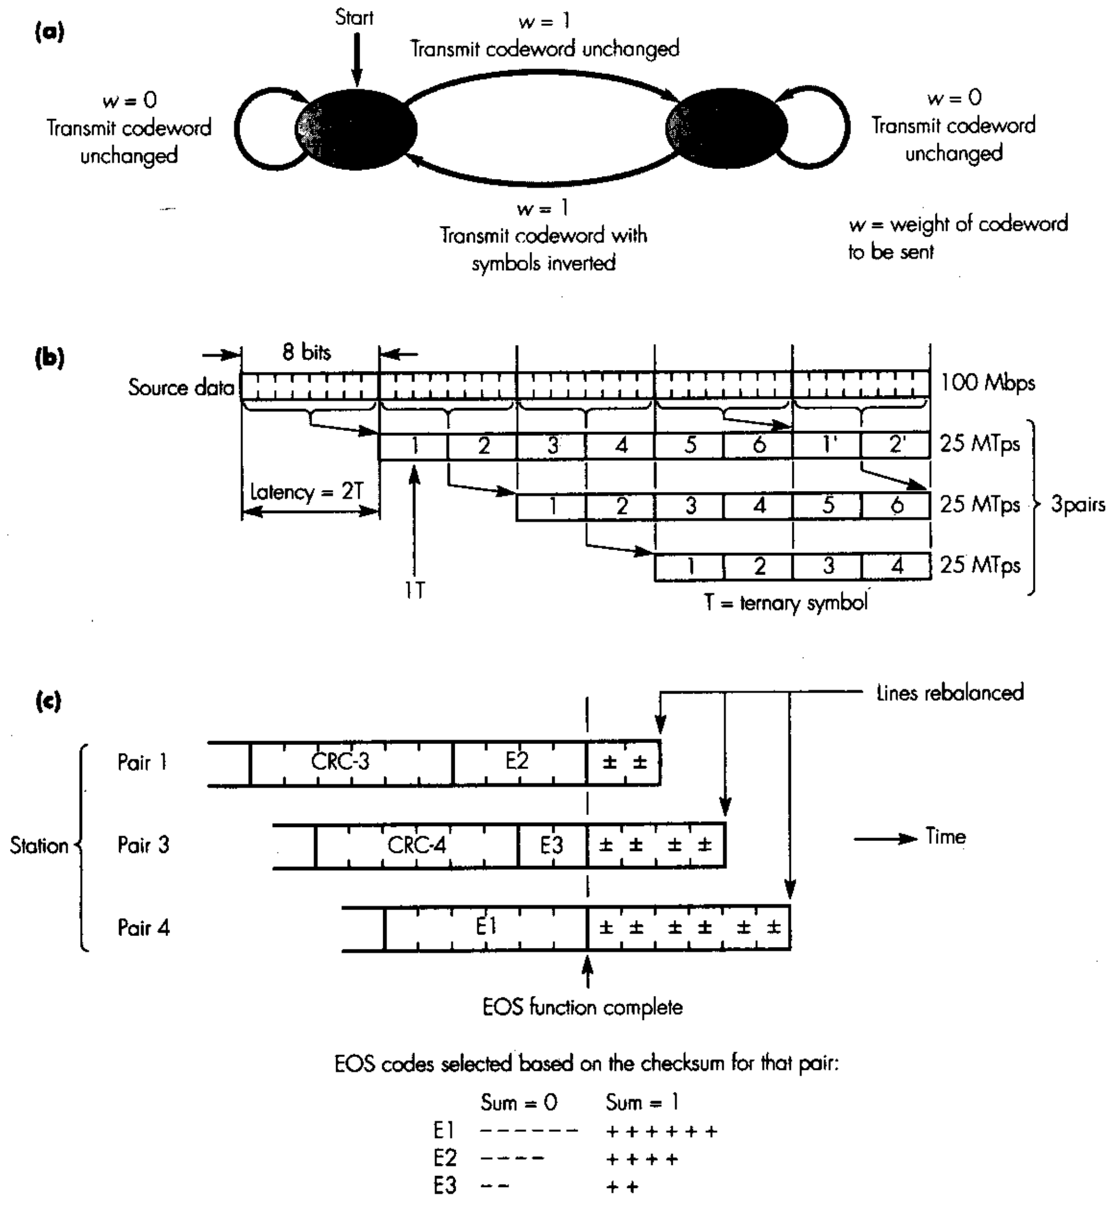

Figure 8.27 100BaseT transmission detail: (a) DC balance transmission rules; (b) 8B6T encoding sequence; (c) end of stream encoding.

presence of a signal on pair 1. However, as Figure 8.28(a) shows, the strong (unattenuated) signals transmitted on pairs 1, 3, and 4 from the station side each induce a signal into the collision detect – pair 2 – wire. This is near-end crosstalk (NEXT) and, in the limit, is interpreted by the station as a (collision) signal being received from the hub. The same applies for transmissions in the reverse direction from hub to station.

To minimize any uncertainty the preamble at the start of each frame is encoded as a string of 2-level (as opposed to 3-level) symbols, that is, only positive and negative signal levels are present in each encoded symbol. This increases the signal-level amplitude variations which, in turn, helps the station/hub to discriminate between an induced NEXT signal and the preamble of a colliding frame.

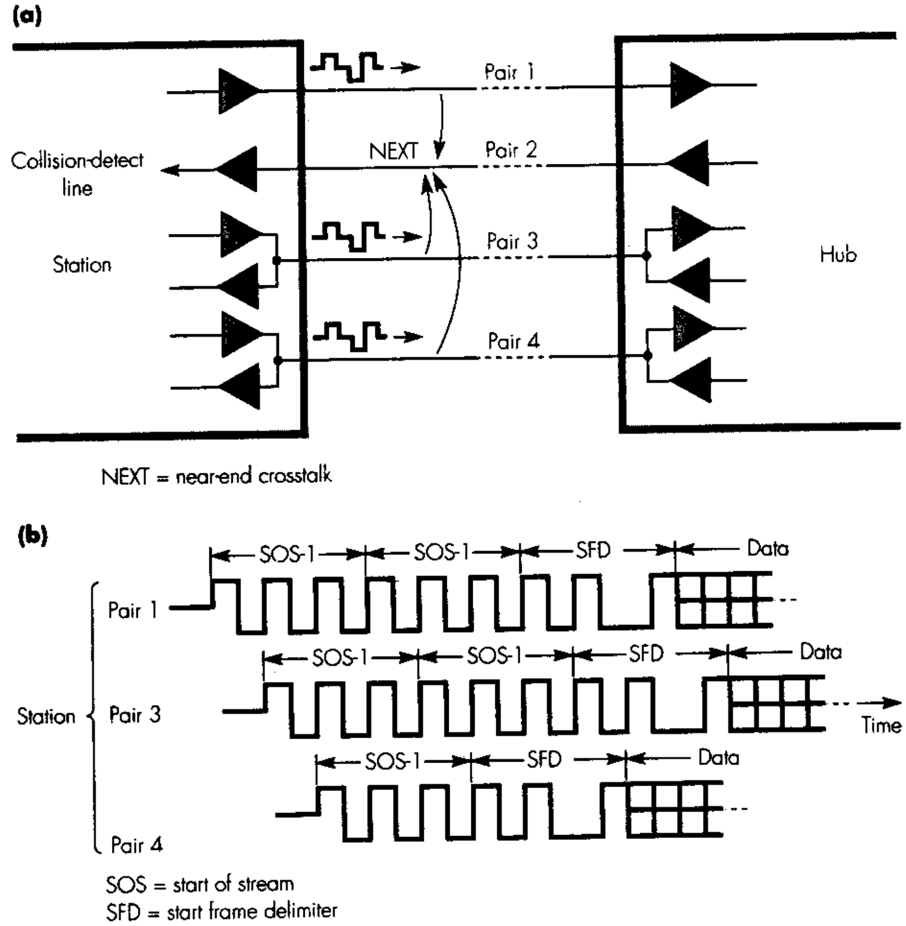

Figure 8.28 Start-of-frame detail: (a) effect of NEXT; (b) preamble sequence.

The preamble pattern on each pair is known as the **start of stream** (**SOS**) and is made up of two 2-level codewords, SOS-1 and SFD. The complete pattern transmitted on each of the three pairs is shown in Figure 8.28(b) and, as we can see, the SFD codeword on each pair is staggered by sending only a single SOS-1 on pair 4. This means that the first byte of the frame is transmitted on pair 4, the next on pair 1, the next on pair 3, and so on. An acceptable start of frame requires all three SFD codes to be detected, and the staggering of them means that it takes at least four symbol errors to cause an undetectable start-of-frame error.

On detecting a collision, a station transmits the jam sequence and then stops transmitting. At this point, the station must be able to determine when the other station(s) involved in the collision cease transmitting in order to start the retry process. In practice, this is relatively easy since, in the nontransmitting (idle) state with 8B6T encoding, a zero signal level is present on the three data wires. This means that there is no induced NEXT signal in the collision detect wire which, in turn, enables the completion of the jam sequence from the hub side to be readily determined. Also, to improve the utilization of the cable, the interframe gap time is reduced from 9.6 μs to 960 ns.

100BaseX

In addition to the 100BaseT standard, a second Fast Ethernet standard is available, which is known as **100BaseX**. Unlike 100BaseT which was designed for use with existing category 3 UTP cable, 100BaseX was designed for use with the higher quality category 5 cable now being used in most new installations. In addition, it is intended for use with STP and optical fiber cables. The use of various types of transmission media is the origin of the "X" in the name.

Each different type of transmission medium requires a different physical sublayer. The first to be developed is that for use with multimode optical fiber cable as used in FDDI networks. Recall from the last section, the FDDI LAN is intended primarily as a backbone subnetwork since, unlike 100BaseT, it can span a distance up to several kilometers. Transmissions over an FDDI network use a bit encoding scheme known as 4B5B (sometimes written as 4B/5B) and this has also been adopted for use with 100BaseX, this version being known as **100BaseFX**.

The cable comprises two fibers, one of which is used for transmissions between the station and hub and the other for transmissions between the hub and the station. As with 10BaseT, collisions are detected if a (colliding) signal is present on the receive fiber during the time the station is transmitting a frame. However, because of the additional cost of both the electrical-to-optical conversion circuits and the associated optical plugs and sockets that are required per port, the cost of the MAC unit associated with the NIC is higher than that used with 100BaseT. Hence the most popular type of Fast Ethernet is 100BaseT and 100BaseFX is used primarily when longer drop cables are required.

8.7.2 Switched Fast Ethernet

As we explained at the start of Section 8.7.1, Fast Ethernet uses the same shared, half-duplex transmission mode as Ethernet. Hence in applications that involve access to, say, large enterprise Web servers, even though the server can handle multiple transfers concurrently, the overall access time and throughput experienced by the various stations using the server is limited by the shared access circuit connecting the server (station) to the hub.

In order to allow multiple access/transfers to be in progress concurrently, two developments have been made: the first, the introduction of a switched hub architecture, and the second, duplex working over the circuits that connect the stations to the hub. The resulting type of hub is known as a **Fast Ethernet switch**.

Switch architecture

The general architecture of a switching hub is shown in Figure 8.29. As we can see, each station is connected to the hub by means of a pair of (duplex) lines which, typically, are implemented as dual UTP (or STP) cables or dual multimode fiber cables. Recall from the last section, each UTP (and STP) cable contains four separate twisted pairs. In the case of 100BaseT, three pairs are used to transmit the 100 Mbps bitstream – in a half-duplex mode – and

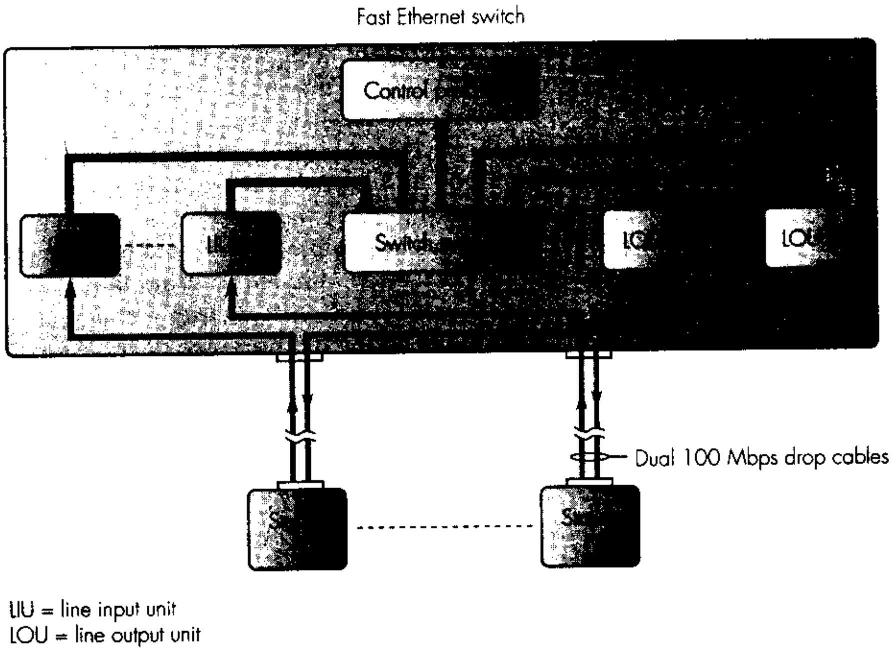

Figure 8.29 Fast Ethernet switch schematic.

the fourth pair is used to perform the carrier sense and collision detection functions. With a switching hub, however, CSMA/CD is not used and instead all stations can transmit and receive frames concurrently. Hence, as with 100BaseT, three pairs in each cable are used collectively to transmit frames (at 100 Mbps) in each direction.

In the case of dual multimode fiber cables, each fiber is used to transmit at 100 Mbps over several kilometers, one in each direction of transmission. Since the 4B5B coding scheme is used, the line signaling rate is 125 Mbaud. In addition, an active signal is maintained on each fiber continuously by transmitting an idle symbol during the idle period between frames. This ensures the receiver DPLL can maintain clock synchronism between successive frame transmissions.

Because each station can transmit frames simultaneously, a frame may be received at multiple input ports of the hub – and hence require processing – simultaneously. Similarly, two or more frames may require the same output line simultaneously. Hence associated with each input and output line is a memory buffer that can hold several frames waiting to be either processed (input) or transmitted (output). The frames – memory pointers to the start of the frame in practice – are stored in a FIFO queue. The control processor then reads the pointer to the frame at the head of each input queue in turn, obtains the destination MAC address from its head, and transfers the frame pointer to the tail of the required output queue to await transmission.

In order to retain the same connectionless mode of operation of the other LAN types, when the switch is first brought into service – and subsequently at periodic intervals – the control processor enters a learning state. This is similar to that used in transparent bridges which we described in Section 8.5.1. Hence when in the learning state, the switch simply initiates the onward transmission of a copy of each frame received from an input line onto all output lines. Prior to doing this, however, the control processor reads the source address from the head of the frame and keeps a record of this, together with the input port number on which the frame was received, in a routing table. The contents of the routing table are then subsequently used to route each received frame to a specific output port. As we can deduce from this, there is a store-and-forward delay associated with a switch. Also, as with a bridge, the FCS at the tail of each frame is used to check for the presence of transmission errors prior to the frame being forwarded and corrupted frames are discarded.

Flow control

As we can see from the above, under heavy load conditions it is possible for all the frame buffers within the switch to become full. At this point, therefore, the control processor must discard any new frame(s). Alternatively, an optional feature associated with switched hubs is to incorporate flow control into the switch. When using flow control, should the control processor find that the level of memory in use reaches a defined threshold, it initiates the transmission of what is called a **Pause** frame on all of its input ports.

On receipt of a Pause frame, the attached station must then stop sending any further frames to the switch until either a defined time has expired or it receives a notification from the switch that the overload condition has passed. Having sent a Pause frame, the control processor monitors the level of memory in use and, when this falls below a second level, it sends out a **Continue** frame on all input por's to inform the attached stations that they can now resume sending new frames.

Network configurations

An example (small) network configuration that includes a Fast Ethernet switching hub is shown in Figure 8.30. As we can see, in order to obtain a high level of throughput, the two servers are connected directly to the switch by means of duplex 100 Mbps lines. All the end-user stations then gain access to the servers through either a 10BaseT or a 100BaseT hub. As we explained in the last section, both these types of hub operate in the half-duplex repeating mode. Hence the duplex uplink port connecting each hub to the switch has bridging circuitry within the hub to temporarily buffer all frame transfers to and from the switch and to perform the CSMA/CD MAC protocol associated with the shared medium hub ports.

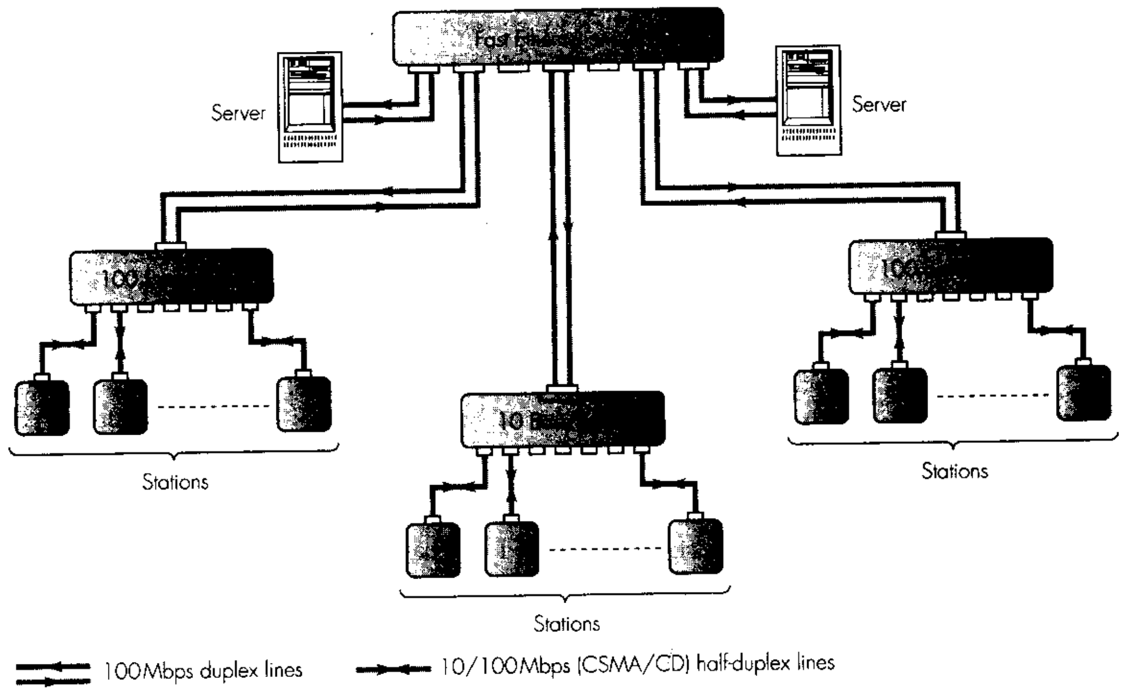

Figure 8.30 Example network configuration with a Fast Ethernet switch and 10/100BaseT hubs.

8.7.3 Gigabit Ethernet

As the name implies, the drop cables associated with Gigabit Ethernet hubs operate at 1000 Mbps (1 Gbps). They have been introduced to meet the throughput demands of an increasing number of servers that hold files containing multimedia information; examples include web pages comprising very high resolution graphics, motion video, and general audio. The hub can be either a simple repeater hub – that is, one that has no memory associated with it – or a switched hub.

An example application of a repeater hub is to distribute the output of a powerful supercomputer (performing, say, 3D scientific visualizations) to a localized set of workstations. The main issue when operating in the repeater mode is to ensure that the round-trip delay between any two stations – the slot time – exceeds the time required to transmit the smaller allowable frame of 512 bits. The time to transmit a 512-bit frame at 1 Gbps is 0.512 microseconds whereas the slot time is in the order of 2 microseconds. To overcome this, a technique known as **carrier extension** is used. This, as the name implies, ensures that a known signal is present on the line for the duration of the slot time and therefore a collision can still be detected. Alternatively, a technique called **frame bursting** can be used which allows multiple short frames to be transmitted one after the other up to the slot time.

In the case of a switched hub, these are used to perform a similar role to the Fast Ethernet switch we showed in Figure 8.30. As we can deduce from the figure, if all the attached hubs are 100BaseT – or Fast Ethernet switches – then the switch providing the backbone function must also be able to handle the increased load generated by the various hubs. Normally, however, the lines connecting each 100BaseT hub/switch to the Gigabit Ethernet hub operate at 100 Mbps and only the servers that are connected directly to the Gigabit hub operate at the full speed of 1 Gbps. However, since in this case both hubs operate in the duplex mode, the slot time is not an issue.

In terms of cabling, repeating hubs can use either category 5 UTP cable with a drop cable length of up to 100 m (1000BaseT) or STP cable providing the length of the drop cable is limited to 25 m (1000BaseCX). In the case of a switched hub used as a backbone, however, optical fiber cable is used that supports drop cable lengths of up to 200 m – using multimode fiber (1000BaseSX) – or up to 10 km if monomode fiber (1000BaseLX) is used. In the case of multimode fiber, an 8B/10B coding scheme is used; that is, each 8-bit group is encoded into a 10-bit symbol. Hence, since the line bit rate is 1 Gbps, the line signaling rate is 1.25 Gbaud. Apart from this, the internal architecture of the hubs is similar to those used in Fast Ethernet hubs and switches.

8.8 LAN protocols

As we have learnt, there is a range of different types of LAN each of which uses a different MAC method, transmission mode, and transmission medium. In the context of our reference model, however, these differences only manifest themselves at the MAC sublayer (of the link layer) and the physical layer. The link control sublayer – known as the **logical link control (LLC) sublayer** in the context of LANs – then offers a standard link layer service to the network layer above it.

Apart from FDDI, which is defined in ISO9314, the various protocol standards associated with LANs are all part of the **IEEE 802 series.** The framework used for defining the various standards is shown in Figure 8.31(a) and a selection of the protocols which we have described in the previous sections are listed in Figure 8.31(b).

(a)

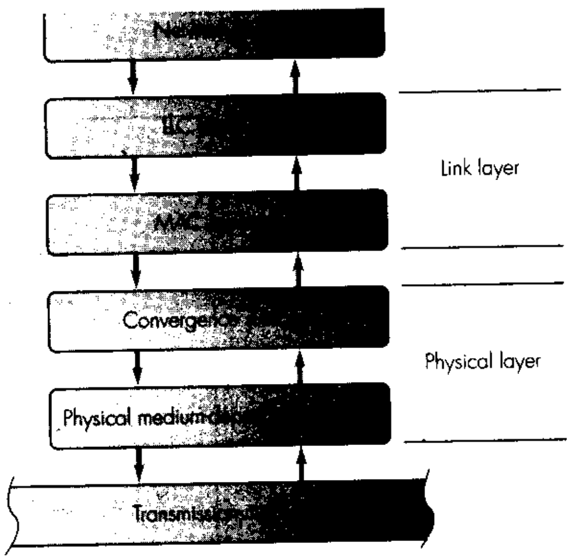

(b)

IEEE	802.1	Station management
	802.1d	Transparent bridges
	802.2	Logical link control (LLC)
	802.3	CSMA/CD (Ethernet) bus
	802.3u	Fast Ethernet
	802.3x	Hop-by-hop switch flow control
	802.3z	Gigabit Ethernet
	802.5	Token ring

Figure 8.31 LAN protocols: (a) protocol framework; (b) examples.

8.8.1 Physical layer

To cater for the different types of media and transmission bit rates, the physical layer has been divided into two sublayers: the **physical medium-dependent (PMD) sublayer** and the (physical) **convergence sublayer (CS)**. To facilitate the use of different media types, a **media-independent interface (MMI)** has been defined for use between the convergence and PMD sublayers. The role of the convergence sublayer is then to make the use of different media types and bit rates transparent to the MAC sublayer.

The set of signals associated with both interfaces vary slightly for the different types of Ethernet (802.3) and token ring (802.5). As an example, the set of signals associated with the different types of Ethernet are shown in Figure 8.32.

As we explained in Section 8.7.1, at bit rates in excess of 10 Mbps it is not possible to use clock encoding – for example Manchester – because the resulting high line signal transition (baud) rate would violate the limit set for use over UTP cable. Instead, bit encoding and a DPLL are used and, to

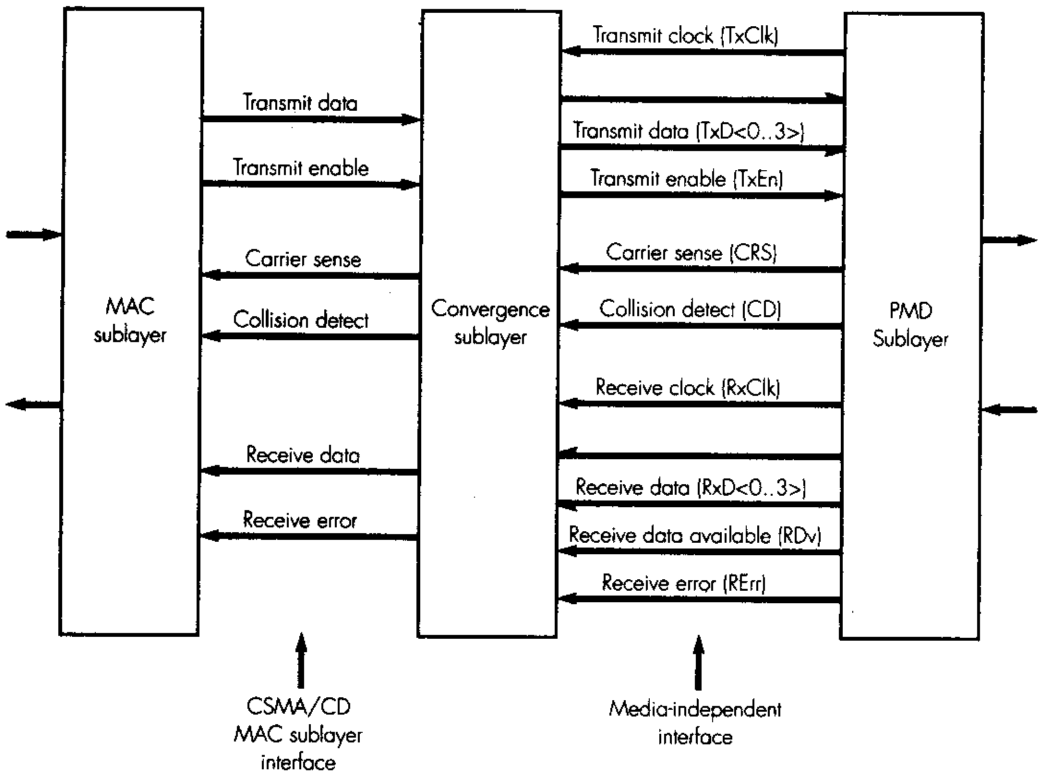

Figure 8.32 Fast Ethernet media-independent interface.

ensure the transmitted signal has sufficient transitions within it, the encoding schemes use one or more groups of 4 bits in each transmitted symbol. For example, one 4-bit group when 4B5B coding is used (100BaseSX), and two 4-bit groups when 8B/10B coding is used (1000BaseSX). Hence, all transfers over the MII are in 4-bit nibbles. The other control lines are concerned with the reliable transfer of these nibbles over the interface. The major functions of the convergence sublayer, therefore, are to convert the transmit and receive serial bitstream at the MAC sublayer interface into and from 4-bit nibbles for transfer across the MII, and, when half-duplex transmission is being used, to relay the carrier sense and collision detect signals generated by the PMD sublayer to the MAC sublayer.

8.8.2 MAC sublayer

Irrespective of the mode of operation of the MAC sublayer, a standard set of user service primitives is defined for use by the LLC sublayer. These are:

MA_UNITDATA.request
MA_UNITDATA.indication
MA_UNITDATA.confirm

A time sequence diagram illustrating their use is shown in Figure 8.33. For a CSMA/CD LAN, the confirm primitive indicates that the block of data associated with the request has been successfully (or not) transmitted – part (a) – while for a token ring LAN it indicates that the request has been successfully (or not) delivered, part (b).

Each service primitive has parameters associated with it. The MA_UNITDATA.request primitive includes: the required destination address (this may be an individual, group, or broadcast address), a service data unit (containing the data to be transferred – that is, the LLC PDU), and the required class of service associated with the PDU. This is used with token ring networks, for example, when a prioritized MAC protocol is being used.

The MA_UNITDATA.confirm primitive includes a parameter that specifies the success or failure of the associated MA_UNITDATA.request primitive. However, as we show in Figure 8.33, the confirm primitive is not generated as a result of a response from the remote LLC sublayer but rather by the local MAC entity. If the parameter indicates success, this simply shows that the MAC protocol entity (layer) was successful in transmitting the service data unit onto the network medium. If unsuccessful, the parameter indicates why the transmission attempt failed. As an example, if the network is a CSMA/CD bus, "excessive collisions" may be a typical failure parameter.

8.8.3 LLC sublayer

The LLC protocol is based on the high-level data link control (HDLC) protocol which we described earlier in Section 6.8. It therefore supports both a

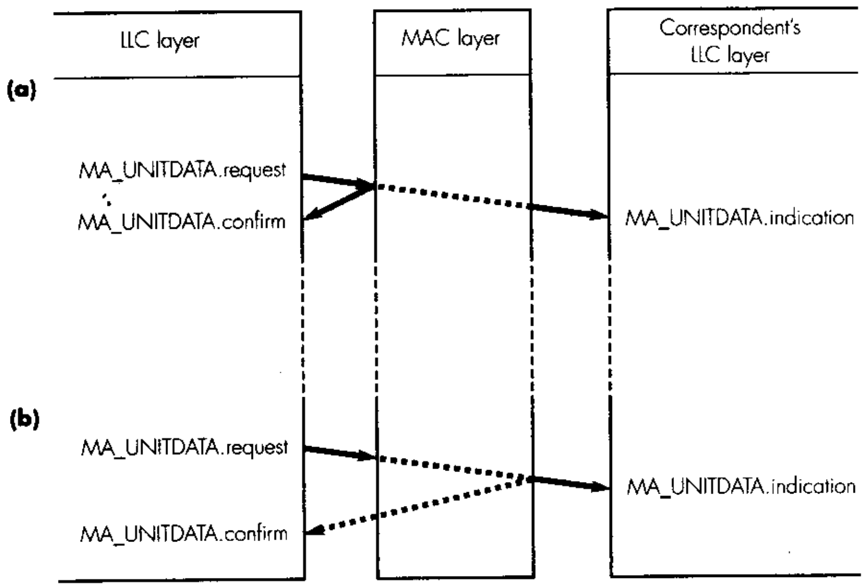

Figure 8.33 MAC user service primitives: (a) CSMA/CD; (b) token ring.

supports both

connectionless (best-effort) and a connection-oriented (reliable) mode. In almost all LAN networks, however, only the connectionless mode is used. Hence, since this operational mode adds only minimal functionality, when the older Ethernet MAC standard is being used – see the discussion on frame formats in section 8.3 – the LLC sublayer is often not present. Instead, the network layer – for example the Internet protocol (IP) – uses the services provided by the MAC sublayer directly.

When the newer IEEE 802.3 MAC standard is being used, the LLC sublayer is present. However, since it operates in the connectionless mode, the only user service primitive is L_DATA.request and all data is transferred in an unnumbered information (UI) frame. The interactions between the LLC and MAC sublayers are as shown in Figure 8.34.

The L_DATA.request primitive has parameters associated with it. These are: a specification of the source (local) and destination (remote) addresses and the user data (service data unit). The latter is the network layer protocol data unit (NPDU). The source and destination addresses are each a concatenation of the MAC sublayer address of the station and an additional service access point (SAP) address. In theory, this can be used for interlayer routing purposes within the protocol stack of the station. In applications such as the Internet, however, this feature is not used and both the destination SAP (DSAP) and the source SAP (SSAP) are set to AA (hex). In addition, two further fields are added. Collectively, the two fields form what is called the **subnet**

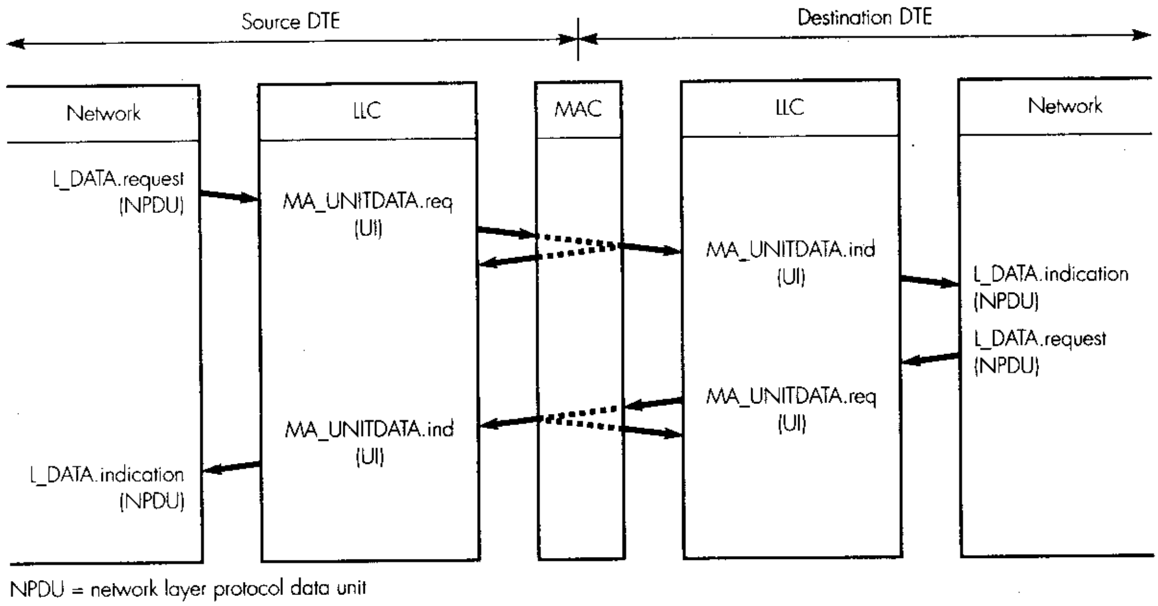

Figure 8.34 LLC/MAC sublayer interactions.

access protocol (**SNAP**) header. The first is a 3-byte field known as the *organisation* (*org*) *code* – which, with the Internet, all three bytes are set to zero – and a two-byte *type* field. This is the same as that used in the original Ethernet standard and indicates the network layer protocol that created the NPDU.

A more detailed illustration of the interactions between the LLC and MAC sublayers is shown in Figure 8.35. The LLC sublayer reads the destination and source LLC service access point addresses (DSAP and SSAP) – from the two address parameters in the event control block (ECB) associated with the L_DATA.request service primitive – and inserts these at the head of an LLC PDU. It then adds an 8-bit *control* (*CTL*) field – set to 03 (hex) to indicate it is an unnumbered information (UI) frame – the 3-byte org code, the type field, followed by the network layer protocol data unit in the user data field. The resulting LLC PDU is then passed to the MAC sublayer as the user data parameter of an MA_UNITDATA.request primitive in a MAC ECB. Other parameters include the MAC sublayer destination and source addresses (DA and SA), the desired service class, and the number of bytes (length indicator) in the user data field. Typically, the service class is used by the MAC sublayer protocol entity to determine the priority to be associated with the frame if a token network is being used.

On receipt of the request, the MAC protocol entity creates a frame ready for transmission on the link. In the case of a CSMA/CD bus network, it creates

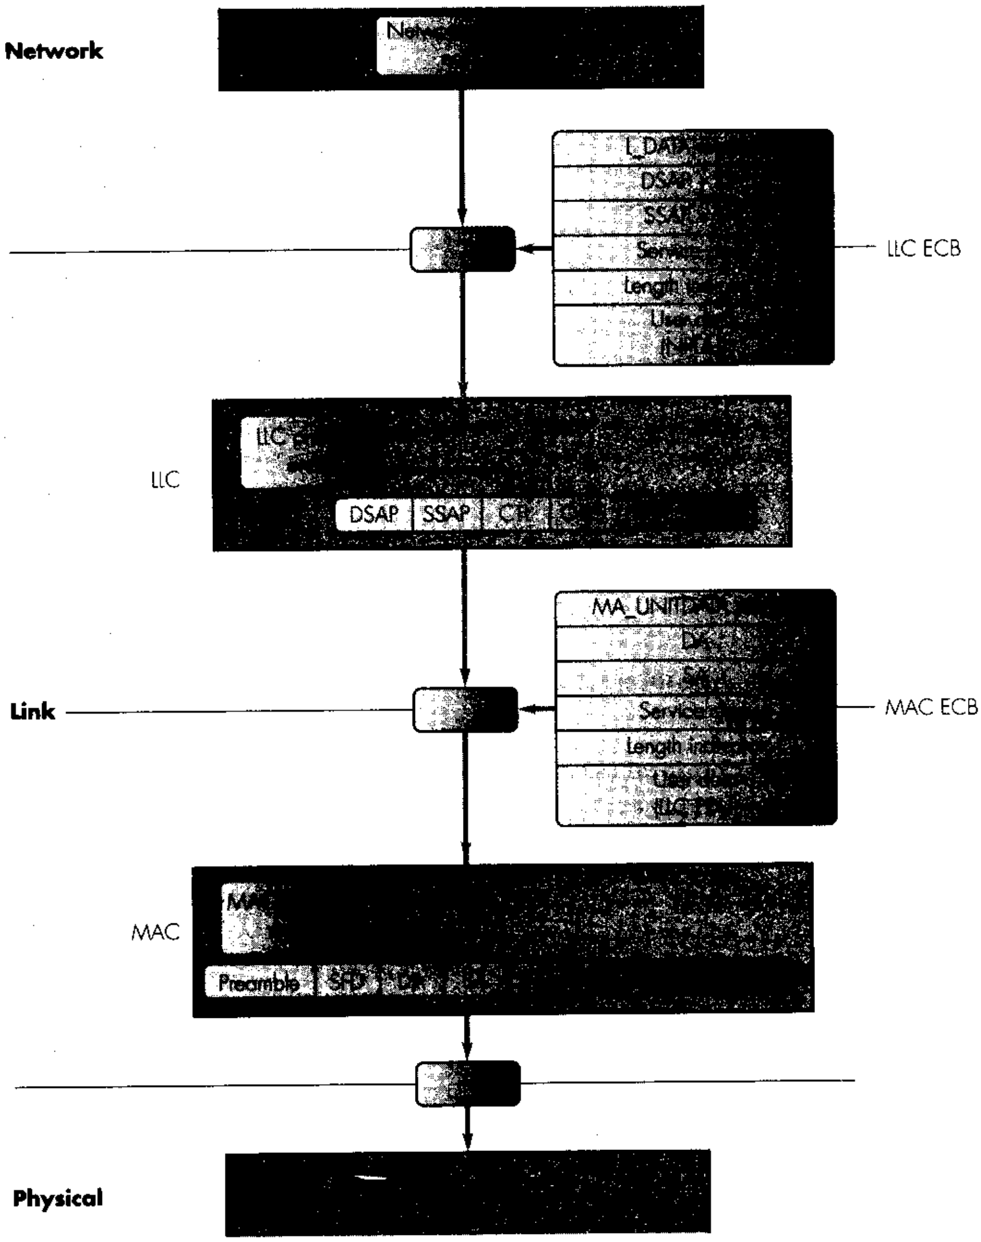

ECB = event control block

Figure 8.35 Interlayer primitives and parameters.

a frame containing the preamble and SFD fields, the DA and SA fields, an I-field, and the computed FCS field. The complete frame is then transmitted bit serially onto the cable medium using the appropriate MAC method.

A similar procedure is followed in the NIC of the destination station except that the corresponding fields in each PDU are read and interpreted by each layer. The user data field in each PDU is then passed up to the layer/sublayer immediately above together with the appropriate address parameters.

8.8.4 Network layer

The most popular protocol stack used within LANs is **Novell NetWare**. Hence it is common for all the stations connected to a site LAN to communicate using this stack. The network layer protocol associated with this is a connectionless protocol known as the **internet packet exchange (IPX)** protocol and, as its name implies, it can route and relay packets over the total LAN. There is no LLC sublayer associated with the stack and hence the IPX protocol communicates directly with the MAC sublayer.

The protocol stack used within the Internet is TCP/IP and the network layer protocol associated with this stack is a connectionless protocol known as the Internet protocol (IP). Hence, as we showed earlier in Figure 5.11, it is common for server machines such as email servers that need to communicate across the Internet, to support both protocol stacks, IPX to communicate with client stations over the site LAN and IP to communicate with other servers over the Internet. Since both IPX and IP are connectionless protocols, they use a single packet to transfer all information over the related network with the full network address of both the source and destination stations in the packet header. We shall defer further discussion of network layer protocols until Chapter 9 when we discuss the operation of the Internet.

8.9 Multisite LAN interconnection technologies

As we saw in the introduction to Section 8.2, in multisite enterprise networks, the LANs associated with the different sites are interconnected together to form an enterprisewide network. As we shall explain, various technologies are available to do this and the choice of technology is determined by the volume of intersite traffic. If this is low, then a low bit rate – up to 56 kbps – leased line with modems is used or, if this is not sufficient, one or more leased 64 kbps ISDN channels. Alternatively, if the volume of traffic is high, then high bit rate leased lines are used, typical bit rates being 1.5 or 2 Mbps or multiples of these. In addition, a number of high bit rate switched services have become available from some telecom providers.

The aim of an enterprise network is to provide an integrated structure so that all communications across the total enterprise ARE transparent to the site/location of the employee's PCs/workstations and the various enterprisewide servers. Hence providing a physical link between sites is only the first step in connecting the LANs at the various sites together. Associated with each LAN is a gateway that provides the additional higher-level protocol support needed to ensure that the physical location of the application-level services is transparent to the end users.

In this section, we first describe the operation of the typical intersite gateways in Section 8.9.1. The first is based on remote bridges and the second on IP/IPX routers. We then describe two examples of switched services: switched ISDN (in Section 8.9.2) and frame relay in Section 8.9.3. A typical enterprise network architecture based on high bit rate leased circuits is then described in Section 8.9.4.

8.9.1 Intersite gateways

The simplest way of providing a transparent connection between sites is to use a **remote bridge** as the intersite gateway. Normally, such bridges are connected directly to the LAN backbone on one port and to the intersite line/network termination on the other. A similar bridge is then used at each site. An example network architecture is shown in Figure 8.36.

As we explained in Section 8.3, the 48-bit MAC address used in LANs to route frames are built into the MAC chipset at the time of its manufacture and each contains a unique address. In practice, therefore, it is possible to treat the interconnected set of site LANs as a single LAN and use just the MAC addresses at the head of each frame for routing.

As we have just explained, various schemes are used to set up connections between sites. Irrespective of the scheme used, however, there is a separate connection set up between each site and hence each pair of LANs. Thus even though each remote bridge has only two physical ports, logically, they behave like multiport bridges with a separate link to each of the other remote bridges. Normally, the spanning tree algorithm we described in Section 8.5.1 spans only a single site/LAN and the remote bridge at the site does not take part in the derivation of the active spanning tree. It does, however, perform the basic learning and forwarding/filtering functions. Hence after the learning phase, it can route all subsequent frames over their appropriate intersite link.

As we show in Figure 8.36, an alternative (and more flexible) way of providing a transparent connection between sites is to use an IP/IPX router as the gateway. Routers are used when either the site LANs are of different types and/or the intersite communications facility is itself a network and requires routing. In both these cases, MAC addresses cannot be used for routing purposes as they only have local significance to a particular site. Hence routing must be carried out at the network layer using either IP or IPX.

As we explained earlier in Section 8.8.4, both IP (as used with the Internet) and IPX (as used with Novell NetWare) are connectionless protocols and all information is transferred over the intersite communications facility in self-contained packets. At the head of each packet is the (enterprisewide) network address of both the source station and the destination

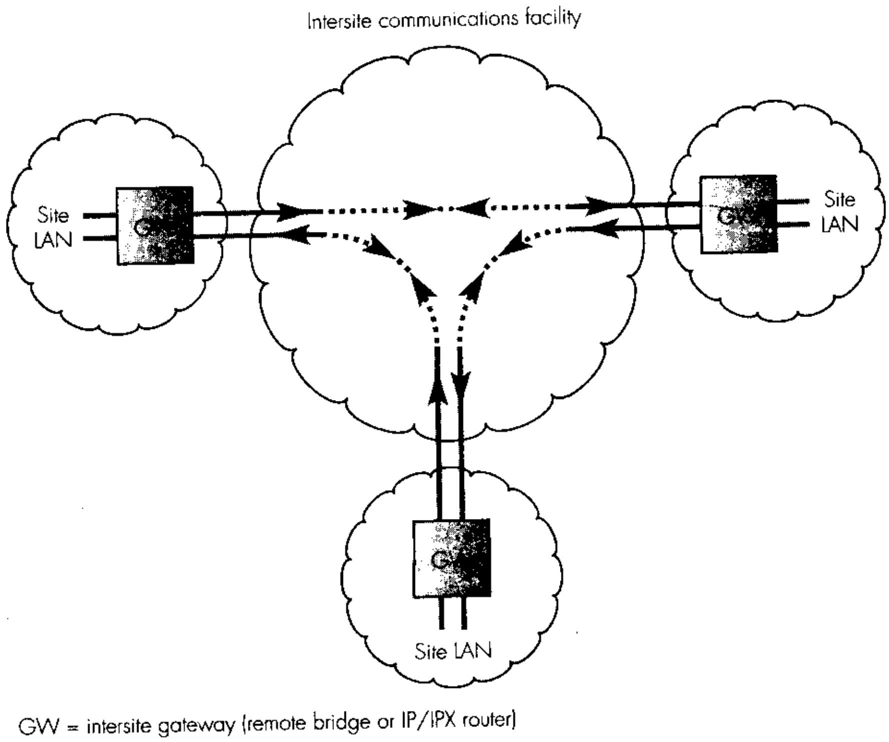

Figure 8.36 Example enterprise network architecture.

station. As we shall expand upon in the next chapter when we discuss the operation of the Internet, each network address comprises a network identifier, which identifies uniquely each site network (LAN) within the context of the total enterprise, and a station identifier which identifies each station at that site. All routing over the intersite network is then carried out using these network addresses. It is for this reason that, when the IP protocol is used, the enterprise network is known also as an intranet. We shall defer further discussion of how the routing is carried out until Section 9.6 when we describe the operation of the IP protocol in some detail.

8.9.2 ISDN switched connections

As we explained in Section 7.2.3, a basic rate interface to an ISDN offers two, separately-switched 64 kbps digital channels. Also, with an ISDN, the connection setup time of a channel is only a fraction of a second. Hence for small enterprises, one solution is to set up (and close down) a channel on demand. For medium-sized enterprises, although the same approach is often acceptable, multiple 64 kbps channels are required to meet the intersite bandwidth requirements. As we explained in Section 7.2.3, to meet this type of requirement a service known as **ISDN multirate** is available from some telecom providers that allows the user to request the bandwidth of a call to be any multiple of 64 kbps. The service is also known as switched $p \times 64$ where p can be up to either 23 or 30 depending on the provider. In order to use the multiple channels as a single high bit rate channel, however, it is necessary to use a device called an **inverse multiplexer**. The general scheme is shown in Figure 8.37(a).

We introduced the concept of multiplexing in Section 7.2.3 as a means of transmitting multiple low bit rate – 64 kbps – digitized voice signals over a single high bit rate line. In contrast, inverse multiplexing is used to derive a single high bit rate channel using multiple lower bit rate channels. For example, an inverse multiplexer can be used to derive a single 384 kbps channel from six independent 64 kbps channels.

At the sending side, the inverse multiplexer first sets up the appropriate number of 64 kbps channels to the required destination. It then proceeds to segment the high bit rate digital stream output by the user equipment – for example a remote bridge – ready for transmission over the multiple low bit rate channels. At the receiving side, once the multiple channels have been set up, the inverse multiplexer accepts the bitstreams received from these channels and reassembles them back into a single high bit rate channel for onward transmission to the receiving terminal equipment.

Hence the function of the inverse multiplexers, in addition to call setup, is to make the segmentation and reassembly operations transparent to the user equipment. In practice, because each channel is set up independently, they may all traverse the ISDN trunk network across different routes/paths. This results in small time differences in the signals received from each channel. To compensate for this, the inverse multiplexer at the receiving side must perform delay compensation and resynchronization of the reassembled bitstream. The general scheme is shown in Figure 8.37(b).

To obtain a similar service to that provided by switched $p \times 64$, inverse multiplexers are avaliable that enable the user terminal equipment to set up and clear multiple 64 kbps circuits on demand. The technique is known as **bonding** and the principles of the scheme are shown in Figure 8.37(c).

In order to set up the high bit rate channel – in response to a user request – the inverse multiplexer at the calling side requests a single 64kbps channel to be set up to the remote site. Once this is in place, the inverse multiplexer at the called side sends back a list of available local (extension) numbers. The calling multiplexer sets up the required number of additional channels – one at a time – and the two user equipments can then exchange data. During the call, a user terminal equipment can request that the aggregated bandwidth available is either increased or decreased by setting up or clearing channels dynamically. For example, if the terminal equipments are remote bridges then the aggregate bandwidth can be regulated to match the data traffic being exchanged at any point in time. An international standard has now been defined in **ITU-T recommendation H.221**, which is concerned with the operation of this type of equipment.

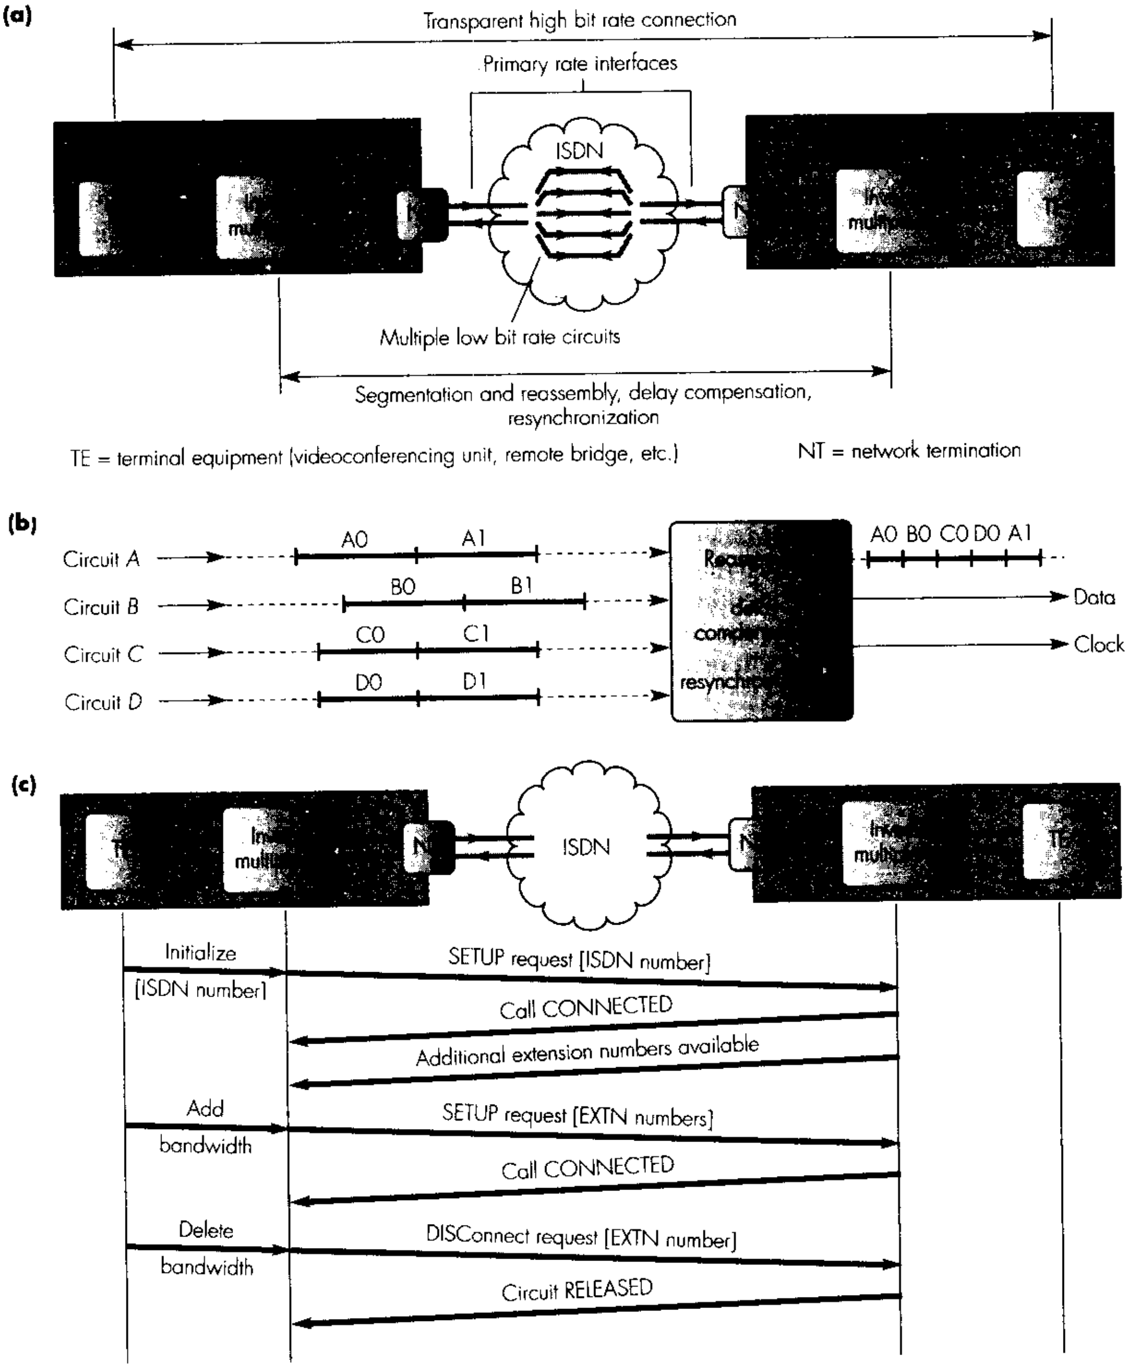

Figure 8.37 Inverse multiplexing: (a) principle of operation; (b) reassembly schematic; (c) bonding protocol.

8.9.3. Frame relay

Frame relay was initially defined as a service provided through an ISDN. Subsequently a number of telecom providers have set up networks to offer just a frame relay service. As the name implies, with frame relay, the multiplexing and routing of frames is performed at the link layer. Moreover, the routing of frames is very straightforward so the channel bit rate can be high, typical rates being up to 34 or 45 Mbps. A schematic diagram of a typical public frame relay network is shown in Figure 8.38.

The customer first informs the service provider of the sites that need to be interconnected. The provider then creates a set of permanent virtual connections that interconnect all the sites by setting up appropriate routing table entries in each frame relay node. The provider then informs the network manager at each site of the identifier that should be put into the header of a frame to reach each of the other sites.

All frames are multiplexed together onto the link connecting the **customer interface equipment** (**CIE**) to its nearest node. Logically, this appears to the customer equipment like a set of point-to-point connections between itself and all the other sites, each identified by the corresponding identifier.

The identifier is known as the **data link connection identifier** (**DLCI**) and is put into the header of each frame. It is then used by the network to route (relay) the frame to its intended destination.

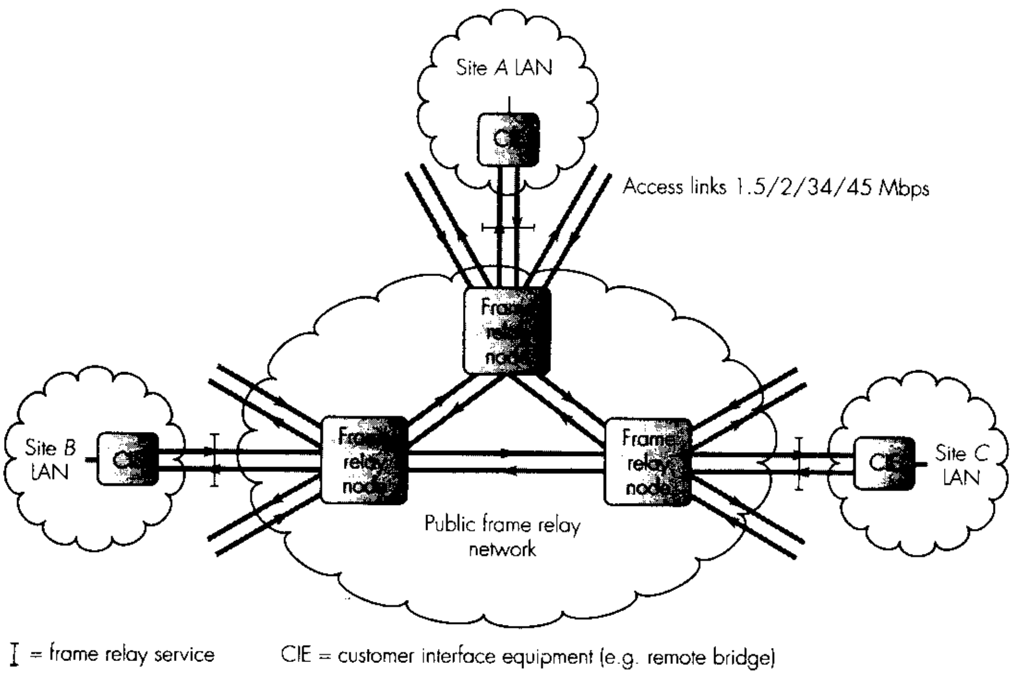

Figure 8.38 Public frame-relay network schematic.

The format of each frame is based on that used in the HDLC protocol and is shown in Figure 8.39(a). It comprises a 2-byte (extended) address header with no control field owing to the lack of any error control. In addition to the DLCI, the header contains the **forward explicit congestion notification (FECN) bit**, the **backward explicit congestion notification (BECN) bit** and the **discard eligibility (DE) bit**. These are used for controlling congestion within the network.

The DLCI, like the VCI in packet-switched networks, has only local significance on a specific network link and therefore changes as a frame traverses the links associated with a virtual path. When the virtual path is first set up, an entry is made in the routing table of each frame relay node along the route of the incoming link/DLCI and the corresponding outgoing link/DLCI to reach the intended destination of the frame. An example set of entries is shown in Figure 8.39(b) and the related routing of each frame is illustrated in Figure 8.39(c).

When a frame is received, the frame handler within each node simply reads the DLCI from within the frame and combines this with the incoming link number to determine the corresponding outgoing link and DLCI. The new DLCI is written into the frame header and the frame is queued for forwarding on the appropriate link. The order of relayed frames is thus preserved and their routing is very fast.

Since multiple frame transfers can be in progress concurrently over each link within the network, during periods of heavy traffic an outgoing link may become temporarily overloaded resulting in its queue starting to build up. This is known as **congestion** and the additional congestion control bits in each frame can be used to alleviate this condition.

Whenever the frame handler relays a frame to a link output queue, it checks the size of the queue. If this exceeds a defined limit, the frame handler signals this condition to the two end users involved in the transfer. This is done in the forward direction by setting the FECN bit in the frame header. In the backward direction, it is done by setting the BECN bit in the header of all frames which are received on this link. If the condition persists, the frame handler also returns a special frame called a **consolidated link layer management (CLLM)** frame to all CIEs that have routes (paths) involving the affected link. Such frames are simply relayed by each intermediate frame relay node in the normal way.

When the frame handler in a CIE receives an indication of network congestion, it temporarily reduces its frame forwarding rate until there are no further indications of congestion. If the overload increases, however, the frame relay node must start to discard frames. In an attempt to achieve fairness, the DE bit in the frame header is used since this is set by the frame handler in each CIE whenever the rate of entry of frames into the network exceeds the peak rate agreed at subscription time.

To minimize the possibility of wrongly delivered frames, the CRC at the tail of each frame is used to detect bit errors in the frame header (and information) fields. Then, if an error is detected, the frame is discarded. With the

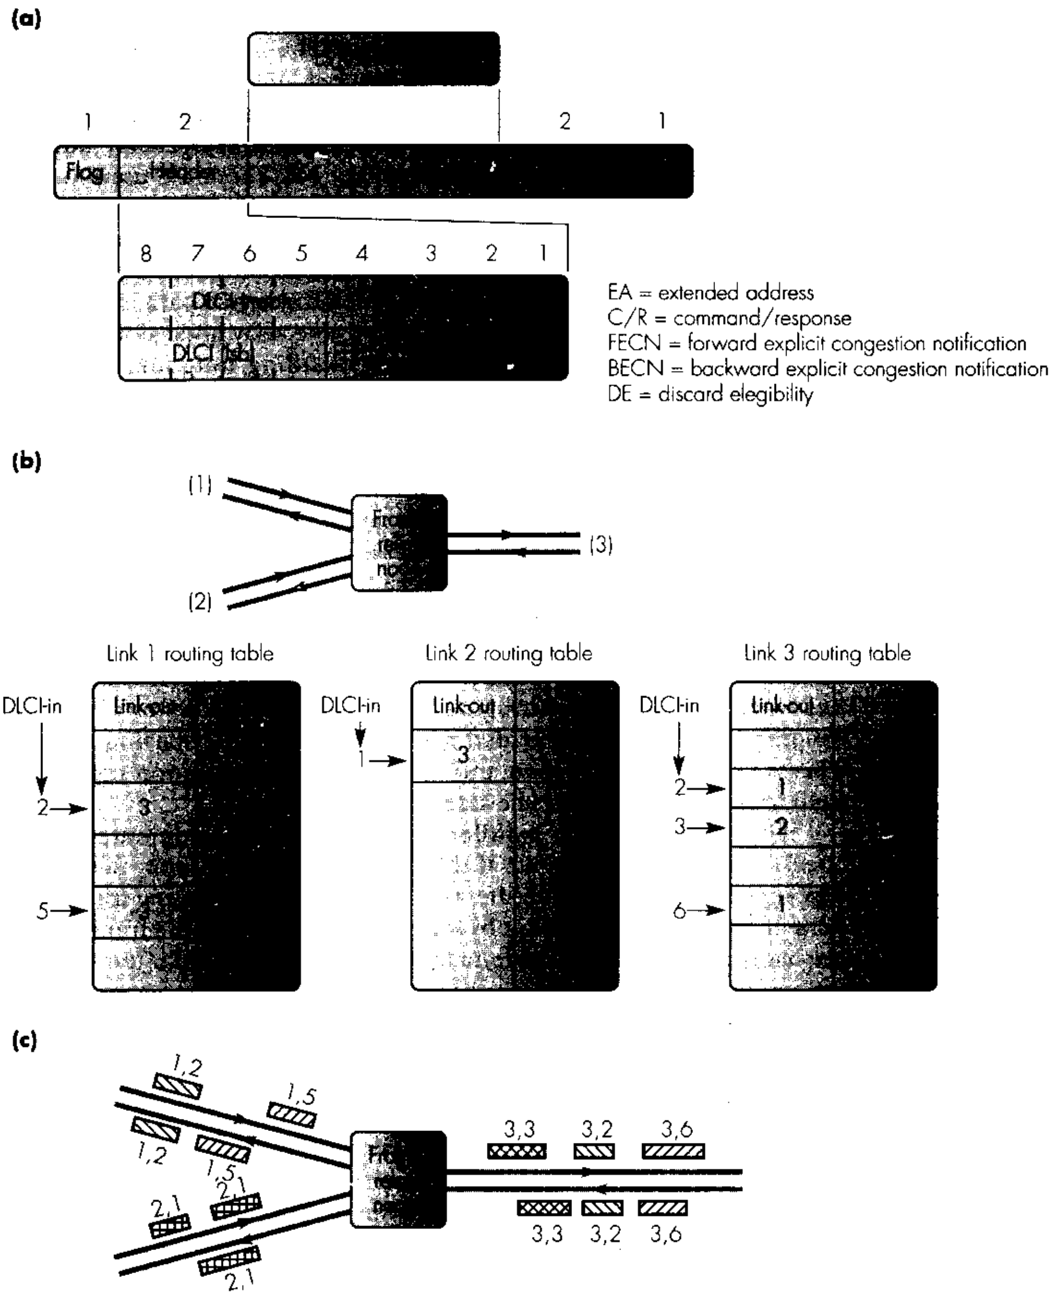

Figure 8.39 Frame relay principles: (a) frame format; (b) frame routing; (c) frame relay schematic.

frame relay service, error recovery is left to the higher protocol layers in the end-user stations. The cost of the service is based on the actual number of frames transferred.

8.9.4 High bit rate leased lines

A typical large multisite enterprise network based on high bit rate leased lines is illustrated in Figure 8.40. As we explained in Section 8.1, in addition to LAN (data) traffic, the network must also support intersite telephony traffic. Typically, the leased lines are either DS1/T1 (1.5 Mbps) or E1 (2 Mbps) or higher such as DS3 (44 Mbps) and E3 (34 Mbps). The multiple 64 kbps channels these contain are divided between telephony – for PBX interconnection – and data – for LAN interconnection – on a semipermanent basis using network management within each site multiplexer.

In the case of telephony, although a 64 kbps channel is used by the PBX for each call, it is now common to use each 64 kbps channel of the intersite leased circuit for more than one call. As we explained earlier in Section 4.2, there is now a range of compression algorithms/circuits available that provide good quality voice communication using 32, 16, or 8 kbps. This means that 2, 4, or 8 calls can be multiplexed into a single 64 kbps channel giving a substantial saving in the number of channels required between sites. The technique used to do this is known as **subrate multiplexing** and is particularly worthwhile over costly international leased lines.

In the case of data traffic, a common approach is to use a (private) **frame relay adapter** (**FRA**). These operate in a similar way to the frame relay nodes in a public frame relay network. The traffic between each pair of sites is allocated a portion of the bandwidth – number of 64 kbps channels – of the related leased line by network management. A set of DLCIs are then assigned to each path/route and these are loaded by network management into the routing tables of the interconnected set of FRAs. A related DLCI for each path is also passed to the remote bridge (or router) at each site and, once the set of MAC addresses associated with each path have been learnt by the bridge, it writes the appropriate DLCI in the header of each frame prior to passing the frame to its local FRA. The role of the FRA is then to relay each frame received from the remote bridge using the preassigned entries in its routing table and the corresponding set of aggregated channels.

Although not shown, there is a single network management station for the total network and all the devices shown in the figure have network management (agent) software within them. Normally, the management station is connected to one of the site multiplexers and a single 64 kbps channel of the intersite leased circuits is then used to transfer management-related messages to/from the management station and all the other network devices. The messages include configuration information to all devices – routing table entries, bandwidth allocations, and so on – and fault reports from the devices. In this way, should a fault develop or a reconfiguration be necessary, this can be done from a central site in a secure way.

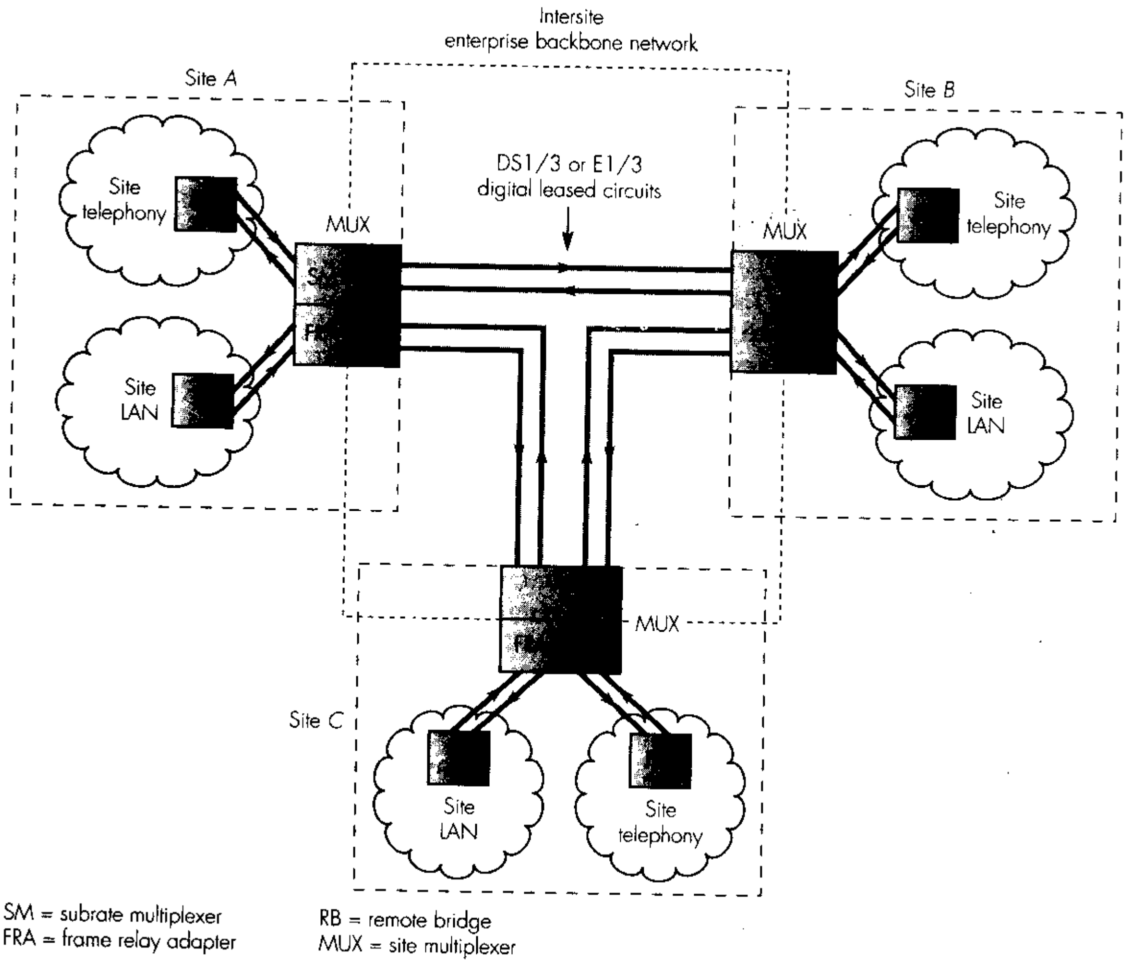

Figure 8.40 Schematic of large multisite enterprise network based on multiplexers and high bit rate leased circuits.

Finally, although most private networks are run and managed by the enterprise to which they belong, a number of telecom providers now offer the option for an equivalent private network to be set up within the provider's network. This is then known as a **virtual private network (VPN)**. These offer the same set of services to the private network but are managed and operated by the telecom provider. In general, VPNs are more expensive than private networks but they have the advantage that the enterprise need not then be involved in the recruitment of staff who are not concerned with the core business of the enterprise.

Summary

A summary of the various topics discussed in this chapter is given in Figure 8.41.

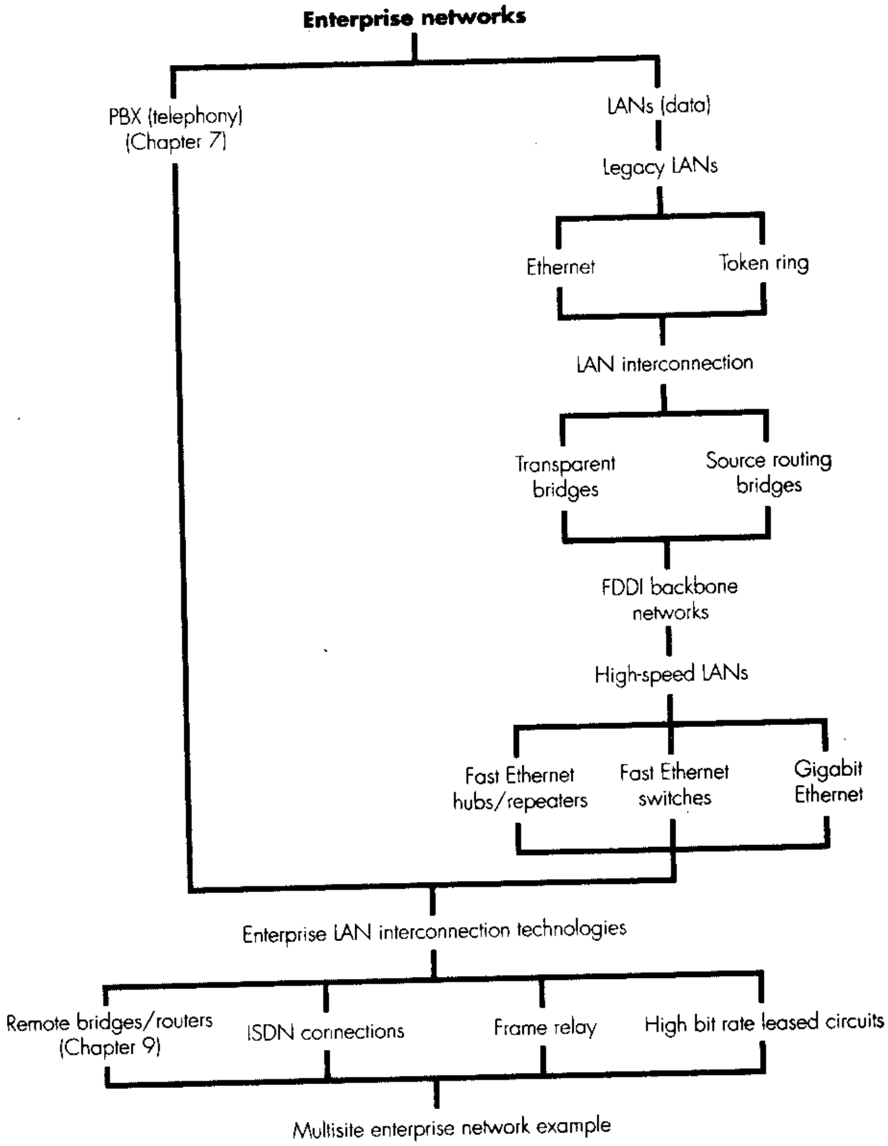

Figure 8.41 Summary of the topics discussed in this chapter relating to enterprise networks.

Exercises

Section 8.1

8.1 What is the meaning of the term "enterprise network"? Describe the factors that determine when such networks are created.

Section 8.3

8.2 Explain the meaning of the following terms relating to the CSMA/CD medium access control method:
- (i) multiple access,
- (ii) broadcast mode,
- (iii) collision,
- (iv) carrier sense.

8.3 With the aid of Figure 8.1, explain the meaning of the term "slot time" and how this is derived. State how the slot time determines the maximum throughput of the LAN.

8.4 State the use of a jam sequence with the CSMA/CD MAC method and explain why a truncated binary exponential process is used.

8.5 Explain the origin of the hub configuration as is now used for Ethernet LANs. Also, with the aid of a diagram, explain how the broadcast mode of operation is achieved.

8.6 Two UTP hubs to which user stations are attached are each connected to a third hub by optical fiber cable in order to gain access to a server that is attached to the third hub. Derive the maximum length of optical fiber cable that can be used.

8.7 With the aid of the frame format shown in Figure 8.3(a), explain:
- (i) the clock encoding method used and how the start of a new frame is detected,
- (ii) how each station that receives a frame determines from the destination address whether the frame contents are intended for it,
- (iii) the use of the type/length field.

8.8 Explain the meaning and use of:
- (i) an interface gap,.
- (ii) a backoff limit,
- (iii) an attempt limit.

Hence, with the aid of flow diagrams, describe the transmit and receive procedures followed by the CSMA/CD MAC sublayer.

Section 8.4

8.9 With the aid of Figure 8.5, explain how the transmission of frames over the ring is controlled using a token. Include in your explanation the two alternatives that are used by a station to release the token.

8.10 With the aid of the two wiring configurations used with a token ring shown in Figure 8.6, explain the meaning/use of the following:
- (i) hub/concentrator,
- (ii) station coupling unit,
- (iii) trunk coupling unit.

8.11 Explain the meaning/use of the following relating to a token ring:
- (i) active ring monitor,
- (ii) minimum ring latency,
- (iii) elastic buffer.

8.12 With the aid of the diagrams in Figure 8.7, explain:
- (i) how the start and end of a frame is detected,
- (ii) the role of the token, monitor, priority, and reservation bits in the access control field,
- (iii) the role of the A and C bits in the frame status field.

8.13 With the aid of flow diagrams, explain the transmission and reception procedures of a frame with a token ring LAN. Include the meaning/use of the token hold timer.

8.14 Assume the same operating conditions as used in Example 8.1. After a period of inactivity, stations 1, 2, and 3 generates a frame to send. If the priority of the frames is 2, 0, and 2 respectively, derive and show in table form the transmissions made by each station for five further rotations of the token.

Section 8.5

8.15 How is a bridge different from a repeater? What are the advantages and disadvantages of each?

8.16 With the aid of the bridge architecture shown in Figure 8.12(a), explain the operation of a bridge including the meaning of the terms:
(i) promiscuous mode,
(ii) forwarding database,
(iii) bridge learning.

8.17 With the aid of Figure 8.12(b), explain how the entries in the two forwarding databases would be built up assuming the following message exchanges: 1 to/from 3, 1 to/from 5, 2 to/from 4, 2 to/from 6.

8.18 With the aid of the LAN topology shown in Figure 8.13, explain why the learning process described in your solution to Exercise 8.16 would not work. Hence state how the topology needs to be changed.

8.19 In relation to the spanning tree algorithm, explain the meaning of the terms:
(i) spanning tree,
(ii) root bridge,
(iii) designated cost,
(iv) root path cost,
(v) root port,
(vi) designated port.

8.20 Assume the same bridged LAN topology as shown in Figure 8.14. Determine the active (spanning tree) topology for the following conditions:
(i) bridge B1 fails
(ii) all bridges are in service but segments S1, 2, 4, and 5 have three times the designated cost of segments S3 and 6 (that is, they have a higher bit rate)
(iii) the same designated costs as in (ii) but bridge B5 has a higher priority than the other bridges.

8.21 With the aid of the LAN topology shown in Figure 8.15, state the routing entry held by station A to reach station D. Hence describe how a frame sent by A is routed to D. Include in your description the structure of the routing information carried in the frame header.

8.22 In relation to the route discovery algorithm used with a source routing bridged LAN, explain the meaning/use of
(i) a single-route broadcast frame,
(ii) an all-routes broadcast frame.

Hence, with the aid of the LAN topology shown in Figure 8.17, explain how the indicated designation of each port is derived.

Section 8.6

8.23 With the aid of the LAN topology shown in Figure 8.18, explain the meaning/use of:
(i) multiport bridge,
(ii) a building backbone,
(iii) a site-wide backbone.

8.24 With the aid of Figure 8.19, explain the meaning/use of the following relating to an FDDI backbone:
(i) primary and secondary rings,
(ii) single attach and dual attach stations,
(iii) optical coupling unit,
(iv) patch panel wiring concentrator,
(v) polarized duplex connectors.

8.25 State why 4B/5B encoding is used in FDDI LANs rather than differential Manchester encoding.

8.26 With the aid of the interface schematic shown in Figure 8.22, explain the meaning/use of:
(i) the latency buffer,
(ii) 4B5B encoder/decoder,
(iii) clock synchronizer.

8.27 In relation to the timed token rotation protocol used with FDDI, explain the meaning of the terms:
(i) TTRT,
(ii) TRT,
(iii) THT,
(iv) early/late token.

8.28 Define the terms (i) maximum obtainable throughput and (ii) the maximum access delay of a shared-medium network. Hence with the aid of the example shown in Figure 8.25, stating clearly any assumptions you make, derive an approximate formula for each in terms of the TTRT and ring latency, T_r.

Section 8.7

8.29 Derive the slot time in bits for a Fast Ethernet hub operating with drop cables of up to 100 m and a bit rate of 100 Mbps. State the implications of your answer in relation to that of a 10BaseT hub.

8.30 With the aid of the wiring configuration of a 100BaseT hub shown in Figure 8.26(a), explain how the carrier sense and collision detection functions associated with CSMA/CD are carried out.

8.31 State why the 8B6T encoding scheme is used for transmissions over each twisted-pair wire with 100BaseT. Also explain why a DC balancing scheme is required.

8.32 With the aid of the diagrams shown in Figures 8.27 and 8.28, explain:
- (i) how DC balance is maintained,
- (ii) the latency associated with the 8B6T encoding sequence,
- (iii) how additional error detection is obtained by sending one of two different end-of-stream codes,
- (iv) how the start-of-stream code used gives a more reliable collision detection.

8.33 What is the difference between a repeater hub and a switching hub? With the aid of Figure 8.29, explain the operation of a Fast Ethernet switch. Include in your explanation how duplex working is achieved.

8.34 With the aid of the network schematic shown in Figure 8.30, explain how a Fast Ethernet switch can be used to provide faster access to site servers for stations attached to 10/100BaseT hubs.

8.35 In relation to a Gigabit Ethernet repeater hub, explain how the shorter slot time is overcome using
- (i) carrier extension,
- (ii) frame bursting.

8.36 Produce a diagram of a network configuration similar to that shown in Figure 8.30 that shows how a Gigabit Ethernet switch can be used to interconnect several 100BastT hubs/switches to a server. Show clearly the bit rate used on each interconnecting line.

Section 8.8

8.37 With the aid of the protocol framework shown in Figure 8.31, explain the role of the following sublayer protocols:
- (i) PMD,
- (ii) convergence,
- (iii) MAC, (iv) LLC.

8.38 With the aid of the interlayer parameters shown in Figure 8.35, explain the meaning/use of:
- (i) the DSAP and SSAP,
- (ii) SNP including the ORG and type fields,
- (iii) CTL.

Section 8.9

8.39 With the aid of the diagram of an enterprise network shown in Figure 8.36, explain how intersite communications are carried out using the following types of intersite gateway:
- (i) remote bridge,
- (ii) IP/IPX router.

8.40 With the aid of the diagrams shown in Figure 8.37, state the role and describe the operation of an inverse multiplexer. Include in your description the principles of the delay compensation and resynchronization procedures.

8.41 With the aid of the network schematic shown in Figure 8.38, explain how a public frame relay network can be used to create an enterprise network. Hence use the frame format and routing principles shown in Figure 8.39 to explain how frames are relayed/routed between all the network sites. Include in your explanation the role of the following fields in each frame header:

(i) DLCI,
(ii) FECN, BECN and DE bits.

Explain the use of a CCLM frame and the actions carried out if a frame is corrupted.

8.42 With the aid of the network schematic shown in Figure 8.40, explain how both telephony (via a PBX) and data (via a remote bridge or router) are extended to cover an entire enterprise. Include in your explanation the role and operation of:

(i) a subrate multiplexer,
(ii) a frame relay adapter.

Can the latter also be used for telephony?

The Internet

9.1 Introduction

As we saw in Chapter 1, the Internet is a global network that supports a variety of interpersonal and interactive multimedia applications. A user gains access to these applications by means of an end system – normally referred to as a host – which, typically, is a multimedia PC, a network computer, or a workstation. As we showed in Figure 1.2, the Internet comprises a large number of different access networks which are interconnected together by means of a global internetwork. Associated with each access network – ISP network, intranet, enterprise network, site/campus LAN, and so on – is a **gateway** and the global internetwork consists of an interconnected set of regional, national, and international, networks all of which are interconnected together using high bit rate leased lines and devices known as **routing gateways** or simply **routers**.

The Internet operates in a packet-switched mode and Figure 9.1 shows the protocol stack associated with it. In the figure, we assume the network interface card in all hosts that are attached to an access network communicate with other hosts using the TCP/IP protocol stack. As we showed in Figure 5.11 (and explained in the accompanying text) this is not always the