11

Entertainment networks and high-speed modems

11.1 Introduction

As we saw in Section 1.4.3, entertainment applications include movie/video-on-demand, near video-on-demand, broadcast television, and interactive television. Also, as we saw in Section 4.3.4, the bit rate requirements for the (compressed) audio and video associated with these applications are determined by the bandwidth of the transmission channel to be used. Typical bit rates are:

VCR-quality video with sound:	1.5 Mbps (MPEG-1)
Broadcast-quality video with sound:	4/6/8 Mbps (MPEG-2, Main)
Studio-quality television with sound:	9/15/18 Mbps (MPEG-2, Main)
High-definition television with sound:	60/80 Mbps (MPEG-2, High)

Normally, the user interaction channel need only be a low bit rate channel of hundreds of kilobits per second. Currently, only VCR and broadcast quality video are used and hence the communications requirement to provide these services is an asymmetric channel consisting of a high bit rate (1.5/4/6/8 Mbps) channel from the service provider to the subscriber and a low bit rate return channel for interaction purposes.

In terms of the networks that are used, as we saw in Figure 1.16, interactions with a video server – for movie/video-on-demand – can be through either a cable network or the access network of a PSTN or ISDN. Alternatively, near video-on-demand can be obtained through either a cable network or a satellite/terrestrial broadcast network. Each can also provide interactive television with, in the case of satellite and terrestrial broadcast networks, the user interaction channel provided through a separate network such as a PSTN.

With a PSTN, the modems we described in Section 7.2.2 provide only a relatively low bit rate switched channel of up to 56 kbps. In addition to these, however, as we saw in Figure 1.1(c), high-speed modems are available which provide a non-switched asymmetric channel that can support applications such as VCR or broadcast quality video-on-demand and high-speed access to the Internet. In the case of cable networks, in addition to the modems used to support broadcast television, high-speed (cable) modems are also available that provide high-speed access to the Internet and other packet-based services. In this chapter we describe the principle of operation of the different types of broadcast television networks and how interactive services are provided with these networks. We also describe selected aspects of the technology associated with the high-speed modems used in cable and PSTN access networks for Internet access.

11.2 Cable TV networks

Essentially, a basic cable TV network is a television distribution facility. A set of TV programs is received from satellite and/or terrestrial broadcasts and the cable network is then used to distribute this set of programs to subscriber premises. Early cable television networks – of which there are still many in existence – are based entirely on coaxial cable. These were designed to distribute broadcast television (and radio) programs that are received at a central site to customer premises geographically distributed around an area such as a town or city. Such networks are known as **community antenna television (CATV) networks** and the basic elements of this type of network are shown in Figure 11.1.

A shared-medium, tree-and-branch architecture is used with, at the root, the various antennas and associated equipment located in a unit called the **cable headend (HE)**. Each branch of the tree is implemented using a single coaxial cable. With this type of network, the electrical signals associated with each television (and radio) program are analog. Hence in order to transmit multiple TV programs concurrently, modulated transmission is used with

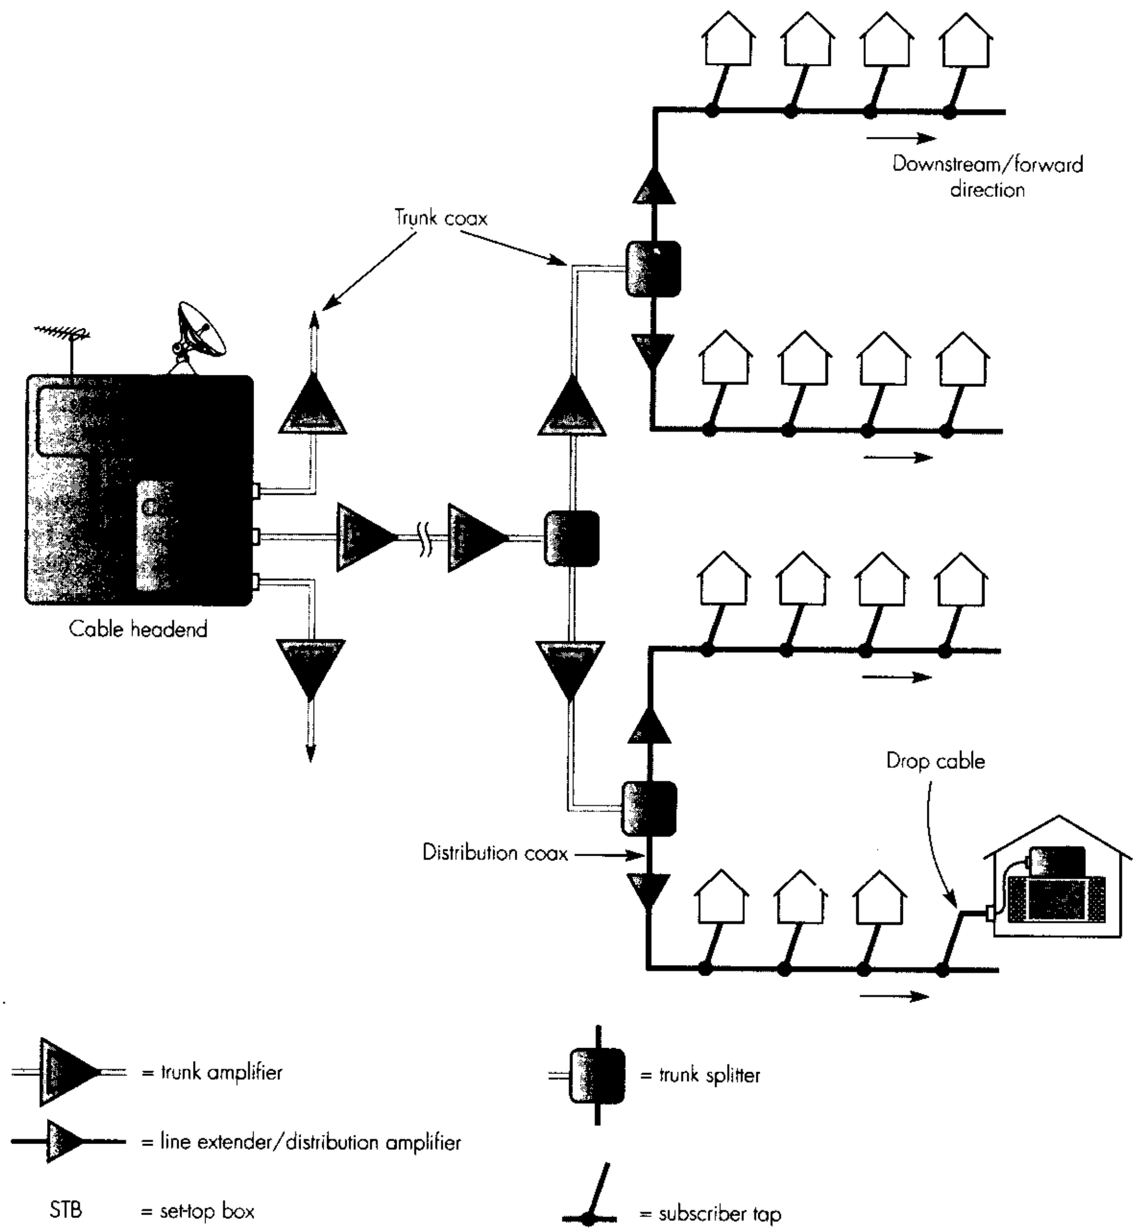

Figure 11.1 Early analog CATV distribution network and components.

each program/channel allocated a fixed portion of the cable bandwidth. The electrical signals of the complete set of TV programs are then multiplexed together onto the cable for distribution to all the network termination points. This technique is known as **frequency division multiplexing (fdm)**.

At each branch node in the cable distribution network, a **trunk splitter** is used to ensure the same set of signals propagate on each branch. Each home passed has a **subscriber tap** attached to the cable which forms the connection point to the distribution cable. A separate **drop cable** is used to connect the tap to the subscriber network termination unit located in the customer premises. The distance from the cable headend to the remotest customer/subscriber premises can be up to tens of miles/kilometers. So to compensate for signal attenuation, a large number of **trunk** and **distribution amplifiers** are required to ensure that the signal received at each subscriber termination is of an acceptable quality with a defined minimum signal-to-noise ratio.

As we saw in Section 2.6.1, a broadcast analog TV signal comprises the video signal and the sound/audio signal. The video signal is made up of the luminance and two chrominance signals and collectively these are transmitted using a separate (single-frequency) carrier for each signal. A typical broadcast analog TV signal is shown in Figure 11.2(a).

The bandwidth used for each TV program/channel in the cable systems of different countries varies. Nevertheless, the same principles apply. Each channel is allocated an appropriate amount of the cable bandwidth – 6 MHz in North America and 8 MHz in Europe – which ensures the signals relating to the two neighboring channels are cleanly separated (by filters) and do not interfere. Associated with each channel is a separate (radio frequency) carrier signal which is separated by 6/8 MHz from the carrier signals of its neighbors. The analog signal of the TV program allocated to that channel is then used to amplitude-modulate the channel carrier signal. The resulting signal is first filtered – to limit the signal bandwidth to 6/8 MHz – and then combined with the signals of all the other channels. The combined signal is then transmitted onto th distribution cable at the cable headend and hence is received by the set-top box of each subscriber. As the subscriber selects a particular program, so the related channel frequency band – and hence carrier signal – is selected. The corresponding carrier signal is then used to demodulate the received signal to obtain the original analog TV signal for that channel.

As we show in Figure 11.2(b), when modulated transmission is used, the bandwidth of a coaxial cable is in the order of 5 MHz through to 900 MHz. In practice, however, the number of channels supported – and hence bandwidth used – is significantly less than this. As we indicated earlier, because of the relatively large area covered by each trunk cable, there may be up to 20 or 30 trunk amplifiers used in each cable run. Each amplifier causes a level of deterioration of the signal as it also amplifies any noise signals introduced by crosstalk and other (external) transmissions. Normally, therefore, the bandwidth used in early cable systems is limited to 88/110 MHz at the low end through to 300, 450, or 550 MHz at the upper end, the actual upper frequency determined by the size of – and hence number of amplifiers used in – the distribution network.

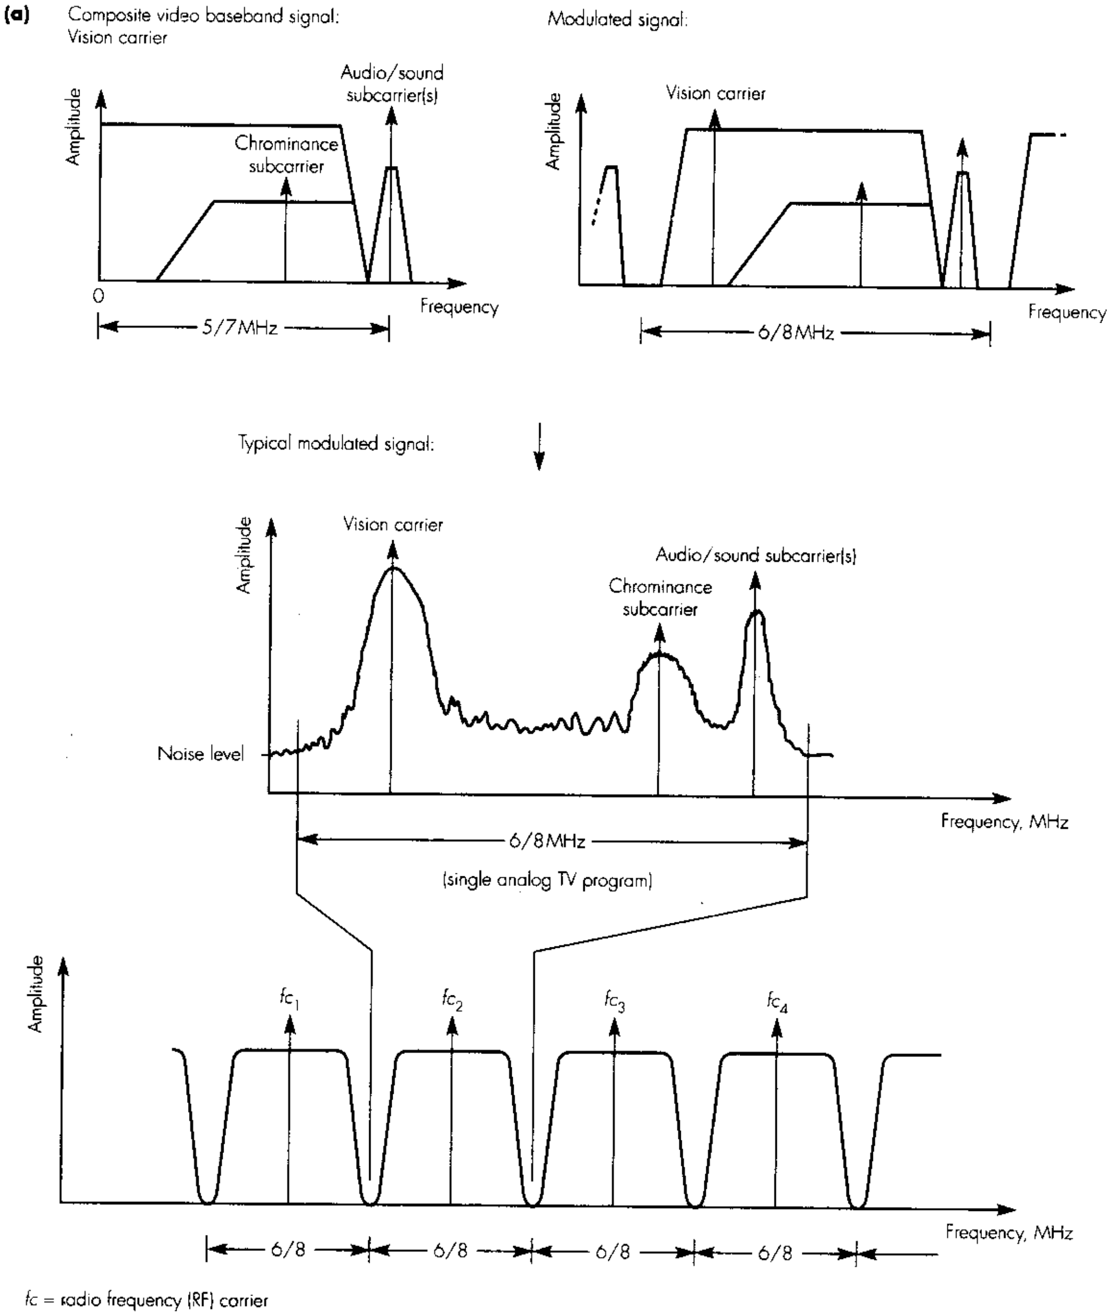

Figure 11.2 Analog CATV principles: (a) analog TV bandwidth requirements; (b) distribution cable bandwidth utilization.

(b)

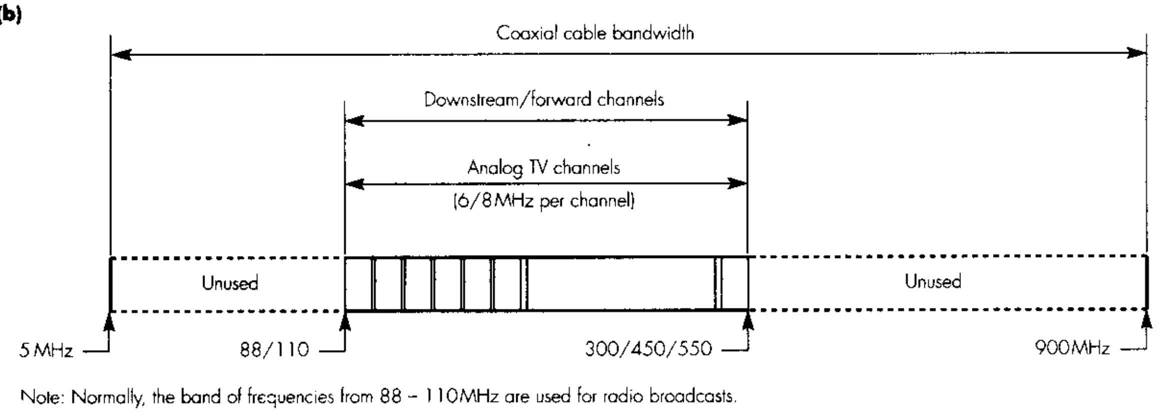

Figure 11.2 Continued

11.2.1 HFC networks

The advent of digital television and the increasing demand for interactive multimedia services – for interactive TV and high-speed access to the Internet for example – has resulted in many operators of cable networks upgrading their networks to support these newer applications. As we indicated in the previous section, even though there is plenty of unused bandwidth in the existing coaxial cable networks, the amount of usable bandwidth is limited by the number of amplifiers required. This is especially the case in the trunk network and hence in the upgraded networks, the main trunk (coaxial) cables have been replaced with optical fiber cables. The resulting network is called **a hybrid fiber coax (HFC) network** and is illustrated in Figure 11.3(a).

As we can see, optical fiber is used in the main trunk network and coaxial cable is limited to the local distribution network. The retention of coaxial cable in this part of the network is done for cost reasons. The number of lines in the distribution network is much larger than the number in the trunk network and hence the cost of upgrading the distribution network to fiber is considerably higher. The fiber used in the trunk network comprises multiple (fiber) trunk cables each terminated in a unit known as a **fiber node (FN)**. Each performs the optical-to-electrical conversion and serves a number of separate coaxial distribution cables. Each distribution cable provides cable services to between 125 and 500 homes distributed over an area of approximately 3 miles/5 kilometers. Typically, therefore, a single FN serves between 500 and 2000 homes and the total network may comprise many FNs.

The use of fiber cable in the trunk network means that the signal attenuation is much reduced so removing the necessity of having a large number of amplifiers in each trunk line. As a result, a much wider portion of the available bandwidth of the coaxial distribution network can be used. A typical utilization is shown in Figure 11.3(b).

(a)

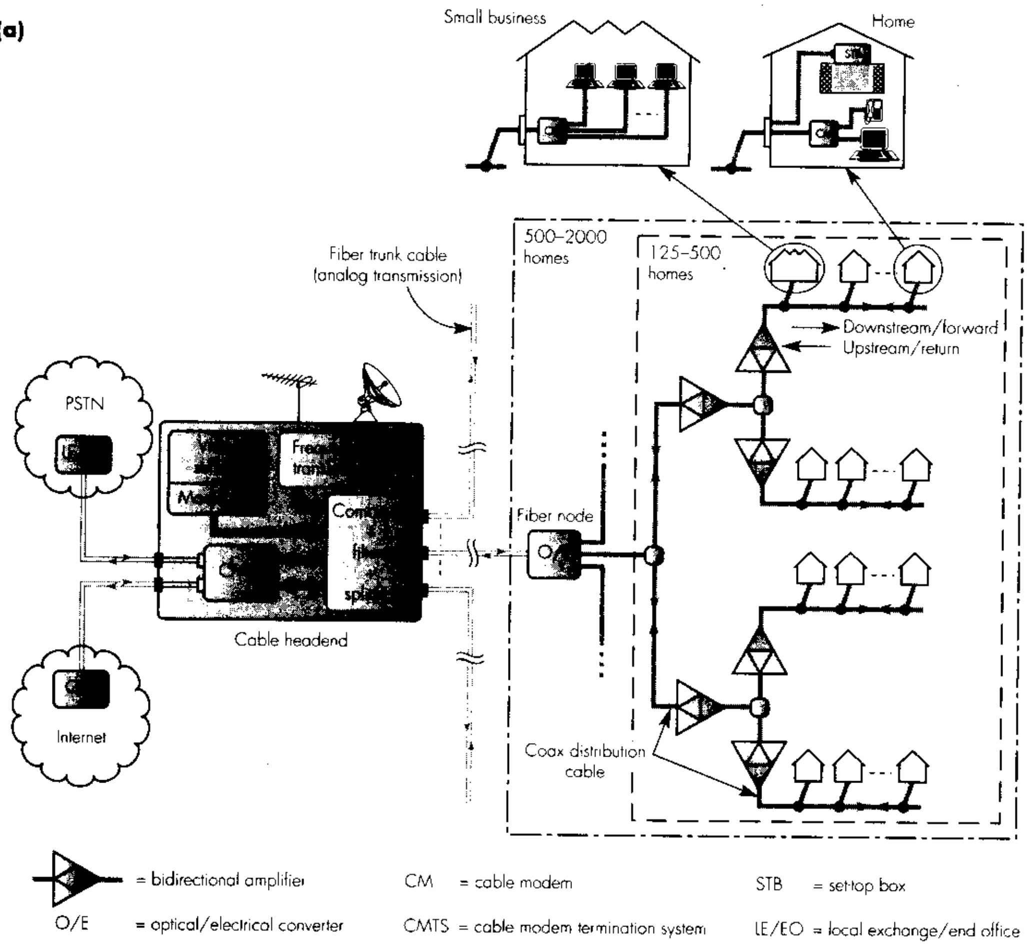

(b)

Downstream frequences:

North America	88–860MHz
Europe	110–862MHz
Japan	88–860MHz

Upstream frequencies:

North America	5–30/42MHz
Europe	5–65MHz
Japan	5–48MHz

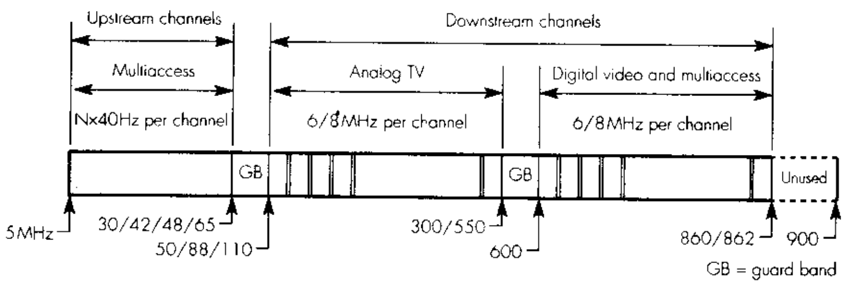

Figure 11.3 Hybrid fiber coax network principles: (a) distribution network and components; (b) cable bandwidth utilization.

As we can see, the bandwidth utilization of the earlier all-coax networks has been extended both at the low and high ends of the available frequency spectrum. To enable existing analog broadcast services to continue to be provided unchanged, the bandwidth from 50/88/100 MHz through to 550 MHz is reserved for such services. Either side of this band is a **guard band (GB)** of 50 MHz at the upper end and 8/10 MHz at the lower end. This is necessary to ensure that the new services are cleanly separated – by means of filters – from the existing services. The additional frequency band in the **downstream/forward direction** from 600 MHz to 860/862 MHz is used for the newer digital services. These include digital TV, video-on-demand and near video-on-demand, and the forward path of a range of two-way communication channels for high-speed Internet access, dedicated high bit rate digital circuits – for LAN interconnection for example – and, in some instances, packet-based telephony.

Although the newer services are all digital, modulated (analog) transmission must be used over the cable distribution network in order to be compatible with the existing analog TV transmissions. All transmissions over the coaxial cable distribution network, therefore, are in an analog form and analog transmission is also used over the fiber cable. Hence, since the signals associated with the newer services are all digital, modems similar to those used with a PSTN are needed to convert the source digital bitstream into (and from) an analog signal for transmission over the distribution network. Because the frequencies used are in the radio frequency spectrum, they are known as **RF modems** and even though radio frequencies are involved, they use a similar modulation scheme to that used in higher bit rate PSTN modems. As we saw earlier in Figure 7.6(c) (and explained in the accompanying text), a multilevel modulation scheme known as QAM is used. In the case of RF modems, however, 64 or even 256 different amplitude and phase combinations are used. These produce usable bit rates in the order of:

64-QAM: 6 MHz band = 27 Mbps, 8 MHz band = 38 Mbps
256-QAM: 6 MHz band = 36 Mbps, 8 MHz band = 50 Mbps

As we can deduce from these values, each additional 6/8 MHz band in the downstream direction can be used to transmit either multiple MPEG-2 (6 Mbps) digital TV programs – or proportionately more MPEG-1 (1.5 Mbps) video programs – or one or more high bit rate data channels. The specific use of the channels is determined by the services to be supported and hence is decided by the cable operator. However, if two-way services are to be supported – interactive TV, video-on-demand, Internet access, and so on – then a related channel in the **upstream/return direction** is required from the subscriber to the cable headend. As we show in Figure 11.3(b), this is achieved by using the lower band of frequencies from 5 MHz through to 30/42/48/65 MHz, the specific upper frequency being determined by the current usage of the cable.

In order to obtain an upstream/return channel, the existing amplifiers in the distribution network are changed to **dual/bidirectional amplifiers**, one to amplify the signals in the downstream frequency bands and the other to amplify the signals in the upstream band. In the case of the fiber trunk network, dual-fiber cable is used, one fiber for the transmission of the downstream signals and the other for the upstream signals. In addition, new subscriber network termination units are required containing **bandsplitting filters** to separate the upstream from the downstream channels.

In some HFC networks, a portion of the upstream bandwidth is already used for the transmission of TV signals (from, say, a camera at a local event to the cable headend for distribution) and hence not all of the bandwidth is available for the newer digital services. Also, in the coaxial cable part of the network, the lower band of frequencies can be of relatively poor quality owing to the ingress of noise from external transmissions and crosstalk from the signals in the various downstream channels. So in order to obtain reliable operation, a more robust modulation scheme is used. This is either QPSK or 16-QAM. Normally, the return channels are not restricted to 6/8 MHz bands and typical bands and example bit rates are:

160 kHz:	320 kbps	1.28 MHz – 2.56 Mbps
320 kHz:	640 kbps	2.56 MHz – 5.12 Mbps
640 kHz:	1.28 Mbps	5.12 MHz – 10.24 Mbps

To support applications such as LAN interconnection and telephony, the forward and return channels are of the same bit rate. Hence symmetric bidirectional communication channels are required. For the various interactive applications, however, the bit rate of the interaction (upstream) channel need only be a fraction of that of the downstream channel. With entertainment applications such as video-on-demand, for example, the traffic in the upstream channel consists mainly of the commands entered by the user at the remote control. Similarly, with Internet access, short commands are used to initiate searching and information retrieval. Hence asymmetry ratios from 10:1 to as high as 100:1 are common.

In some interactive TV applications, the interaction channel is via a PSTN and hence a channel in the upstream direction is required for this. In the case of Internet access, a single high bit rate channel in the downstream direction is used with a lower bit rate channel in the upstream direction for interaction purposes. Each PC/network computer/workstation in all the homes/businesses attached to the cable shares the use of each of these channels in a similar way to the stations attached to a LAN. The channels are known as **multiaccess channels** and we describe aspects of their operation in the next section.

Multiaccess channel operation

As we have indicated, in a cable network all transmissions are either to or from the cable headend rather than between two attached stations as is the case in a LAN. For high-speed Internet access, each pair of channels – one in the downstream direction and the other in the upstream direction – provides a two-way communications path between a station – also referred to as a **customer premises equipment** (**CPE**) – in a home or business and a server attached to the Internet. As we showed in Figure 11.3(a), each station is attached to the cable through a **cable modem** (**CM**). This performs the modulation of the bits in a packet sent by a station for transmission in the upstream channel and the demodulation of the bits in a packet received from the downstream channel. In addition, some CMs support packet-based (IP) telephony. Typically therefore, as we explained in Section 5.3.2, the telephony interface within the CM is based on the H.323 standard. This is concerned with the creation and exchange of call setup and clearing – signaling – messages/packets and the packetization/depacketization of voice samples.

At the cable headend is a unit called a **cable modem termination system** (**CMTS**). This is used to control the transmissions to all the CMs attached to the distribution network and to relay packets to and from the related network – the Internet or a PSTN. Figure 11.4 shows the essential components of a CMTS.

As we shall explain below, in addition to (IP) packets relating to Internet access and telephony, the CMTS uses IP packets to communicate with each CM in order to configure and manage the CM. Typically, therefore, as we see in the figure, a Fast-Ethernet switch is used to relay incoming/outgoing packets to their related interface. This is either a local server (for configuration and management), an access gateway (for Internet access), or an H.323 gateway (for access to a PSTN).

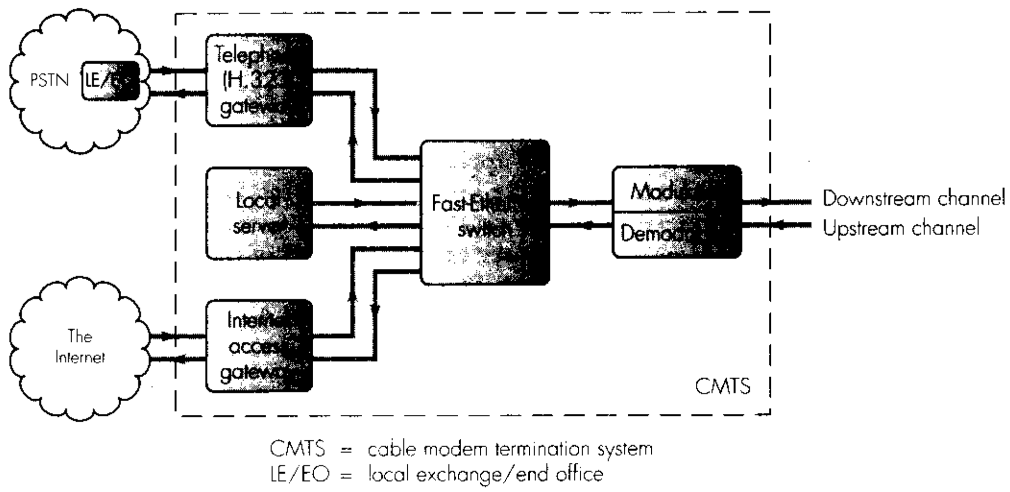

Figure 11.4 Cable modem termination system schematic.

In the downstream direction, each packet is simply broadcast over the assigned downstream channel and hence is received by all the attached CMs. In the upstream direction, however, since the attached CMs must compete for the use of the upstream channel, a suitable medium access control (MAC) protocol must be used to ensure each CM gains access to the channel in a fair way. In a cable network, since a CM cannot receive the transmissions made by a CM that is nearer to the CMTS than itself, MAC protocols such as CSMA/CD and control token cannot be used. So instead, access to the upstream channel is controlled by the CMTS. In practice, there are a number of different schemes that can be used. In the remaining part of this section we describe the principle of operation of just one of these – the **data-over-cable service interface specification (DOCSIS)**, which is very similar to the **IEEE 802.14** standard.

Protocol stack

A schematic diagram showing the protocol stack associated with the CMTS and CM that is used for Internet access is given in Figure 11.5.

Each CM has an integral repeater hub – for example, 10BaseT – within it, and all the stations – there can be more than one at a site – are attached to the hub by, for example, individual twisted-pair drop cables. Hence each CM has its own 48-bit MAC address – which is used for identification purposes by the CMTS during the configuration and management phase – and an additional set of MAC addresses, one for each station attached to the hub. Normally, the latter are acquired by the hub when each station becomes active.

The aim is to make the presence of the cable network – and hence CM and CMTS – transparent to the IP layer in a station that is communicating with a remote server attached to the Internet. To achieve this, forwarding of IP packets through the CM is carried out at the MAC sublayer and hence the CM acts as a bridge. In the case of the CMTS, although bridging can optionally be used, normally, the forwarding operation is performed at the IP layer with the CMTS acting as an Internet access gateway for all the stations – through their related CMs – that are attached to the cable network. Hence, as we explained in Section 9.5.1, the ARP in the IP layer of the CMTS acts as a proxy ARP for all the attached stations. The IP layer must also support the ICMP (see Section 9.7) and IGMP (see Section 9.6.10) protocols. In addition, as we show in the figure, each CM attached to the cable network has additional higher-level protocols above the MAC sublayer. These enable each CM to be configured at start up once the basic MAC-level transmission structure has been established – DHCP and TFTP – and to be remotely managed – SNMP. We described the operation of DHCP earlier in Section 9.10.4 and we describe a selection of the Internet application protocols including TFTP and SNMP in Chapter 14.

Configuration and management

All CMs are diskless and hence need to be configured at start up. This is carried out in two phases. First, since the only address a CM knows at start up is

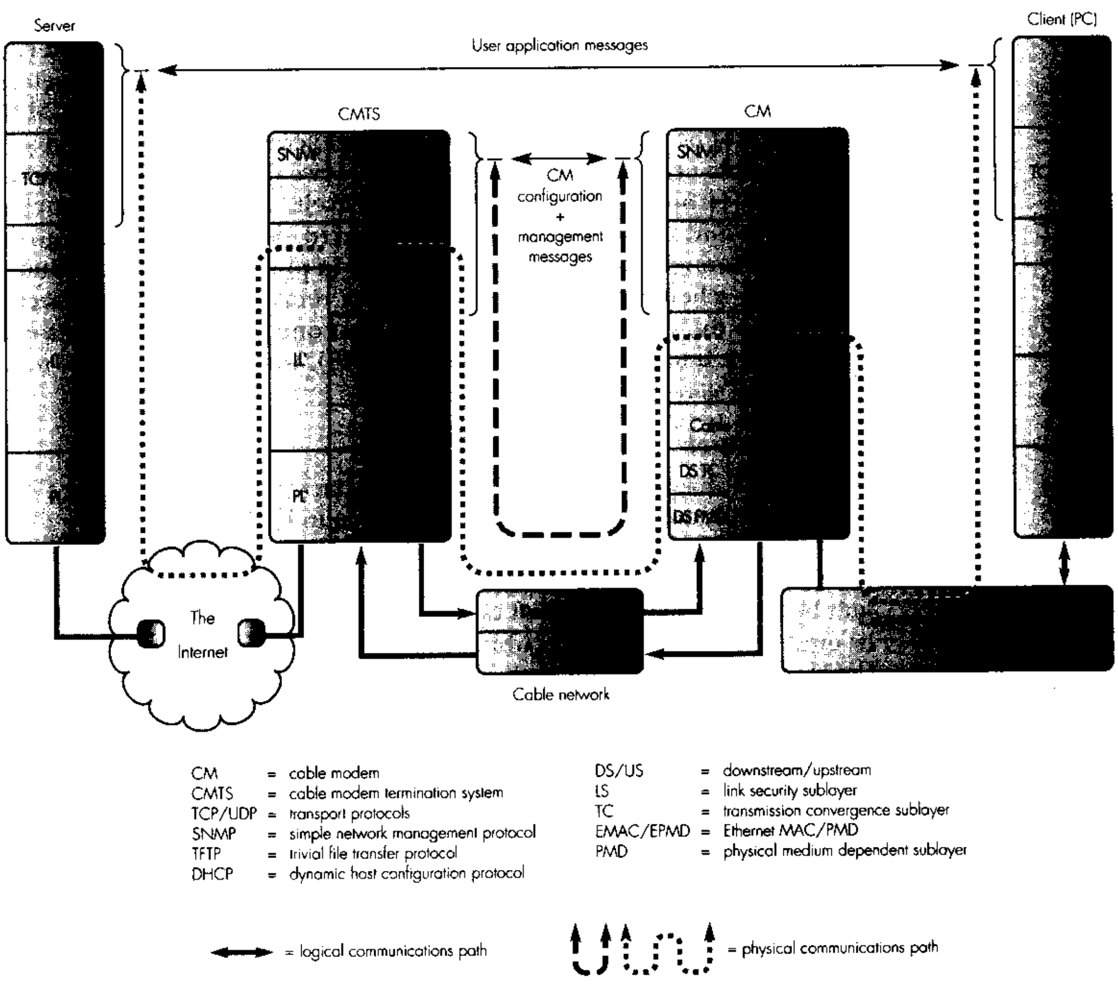

Figure 11.5 CMTS/CM protocol stacks.

its MAC (hardware) address, the **dynamic host configuration protocol (DHCP)** – see Section 9.10.4 – in the CM requests the DHCP in the CMTS to send it the IP address of itself and the CMTS. Also, the name of the servers that supply the time-of-day and hold the modem configuration information. On receipt of this information, the CM proceeds to obtain the time-of-day (ToD) from the ToD server and, using the **trivial file transfer protocol (TFTP)** – see Section 14.5 – the configuration information.

The configuration information includes the public encryption key of the CM and the QoS parameters of the packet flows the CM can support. Examples include best-effort flows associated with Internet access and the constant bit rate flows associated with telephony. We expand upon the subject of encryption later in Section 13.4. Essentially, however, the public key of the CM is used by the CMTS to send the private/secret key that is to be used by the **link security** (**LS**) sublayer of the CM to encrypt/decrypt the payload of each link layer frame sent/received from the cable. Once it has done this, the CM sends a request message to the CMTS to formally register and obtain permission to start to relay frames. The CMTS responds by assigning a unique **service identifier** (**SID**) to the CM and also one or more **service flow identifiers** (**SFID**), one for each type of packet flow the CM can support. The CM can then start to relay frames to/from the stations attached to its **customer premises network** (**CPN**) interface.

As we shall expand upon in Section 14.7, the simple network management protocol (SNMP) is used to enable a management process in the CMTS – or another remote system – to read and change a set of operational parameters that are maintained by the CM. It is also used for a process in the CM to report fault conditions should they arise. Normally, the process in the CMTS displays this information on a computer screen and, should a fault be reported, the network manager can readily determine the location of the CM and the fault that is present.

Cable MAC

As we indicated earlier, all transmissions on the upstream channel – that is, from each CM to the CMTS – are controlled by the cable MAC in the CMTS. Hence when a CM wishes to transmit a frame – containing an IP packet in its payload – the cable MAC in the CM first sends a *Request* (*REQ*) message to the cable MAC in the CMTS which indicates the amount of bandwidth that is required to transmit the frame. On receipt of the REQ message, the CMTS responds by returning a *Grant* message which indicates the amount of bandwidth that has been allocated to the CM.

All transmissions on the upstream .hannel are scheduled to occur in fixed time intervals of 4 ms. Hence, during a transmission interval, the CMTS may receive a number of requests for a portion of the bandwidth in the next transmission interval. A bandwidth allocation algorithm is used to ensure that requests are granted in a fair way. For example, if the waiting frame contains speech samples then it may be allocated a portion of the bandwidth ahead of, say, a frame containing best-effort data traffic. Once this has been carried out, the CMTS divides the subsequent transmission interval into a number of variable-length (time) intervals each with the SID of the requesting CM that can use it. This information is then broadcast on the related downstream channel during the current transmission interval in what is called an *upstream bandwidth allocation MAP* management message.

For each requesting CM, the corresponding entry in the allocation MAP message gives the duration of time it can transmit. If the time duration for a CM has been set to zero, this indicates the request is still pending and the CM must wait for the next transmission interval. Also, since a CM can have only a single outstanding request (per SID) at a time, this blocks the CM from making a new request.

Because of the relatively wide area of coverage of HFC networks, as we saw in Section 6.2.8, the signal propagation delay over the cable can be relatively large. Hence the precise time each requesting CM receives the allocation MAP varies and depends on the physical distance the CM is from the cable headend. To ensure each CM times its frame transmission with those from all of the other requesting CMs, firstly, all CMs must be in time synchronism with the master clock in the CMTS and secondly, in order for the transmitted frames to arrive at the CMTS in their expected time intervals, each CM must know the round-trip propagation delay time between itself and the CMTS. The procedure used to enable a CM to determine this is known as **ranging** and both time synchronization and ranging must be carried out during the initialization phase of the CM prior to the transmission of any frames containing IP packets.

Time synchronization

To ensure all CMs have a common time reference, the CMTS maintains a 32-bit binary counter that is clocked using its master clock of 10.24 MHz. The CMTS then broadcasts a *SYNC management message* that contains the current state of this counter on the downstream channel at periodic (10 ms) intervals. The counter value in the SYNC message is known as the *SYNC time-stamp* and is used by each CM to synchronize its own time reference to this value. The general principle is shown in Figure 11.6.

The upstream channel bitstream is divided in *time* into a stream of numbered **minislots**. The duration of each minislot is a power-of-two multiple of 6.25 microseconds such that the number of bytes per minislot is equal to 16. For example, assuming a bandwidth of 640 kHz is being used for the upstream channel, as we indicated earlier, a typical bit rate is 1280 kbps. Hence each 6.25 μs interval contains $1280 \times 10^3 \times 6.25 \times 10^{-6} = 8$ bits. Therefore the time duration of a minislot, T, is $6.25 \times 16\,\mu s$ and hence the power-of-two multiple, M, is 4.

All frame transmissions in each transmission interval – including frames containing request and timing messages – start at a defined (numbered) minislot and the bandwidth allocation MAP uses units of minislots. The CMTS uses a second 32-bit counter to number each minislot. This is clocked at 6.25 μs intervals which is equal to a clocking rate of 160 KHz. Hence, since this is 1/64th of the CMTS master clock of 10.24 MHz, there is always a fixed relationship between the contents of the minislot counter and the broadcast SYNC time-stamp: the least significant 0 to $(25 - M)$ bits of the minislot counter are the same as the most significant $(6 + M)$ to 31 bits of the SYNC time-stamp.

Ranging

Once a CM is in time synchronism, it invokes the ranging procedure to determine its **round-trip correction (RTC) time**; that is, how much earlier the CM should start transmitting a frame relative to a given minislot number in the

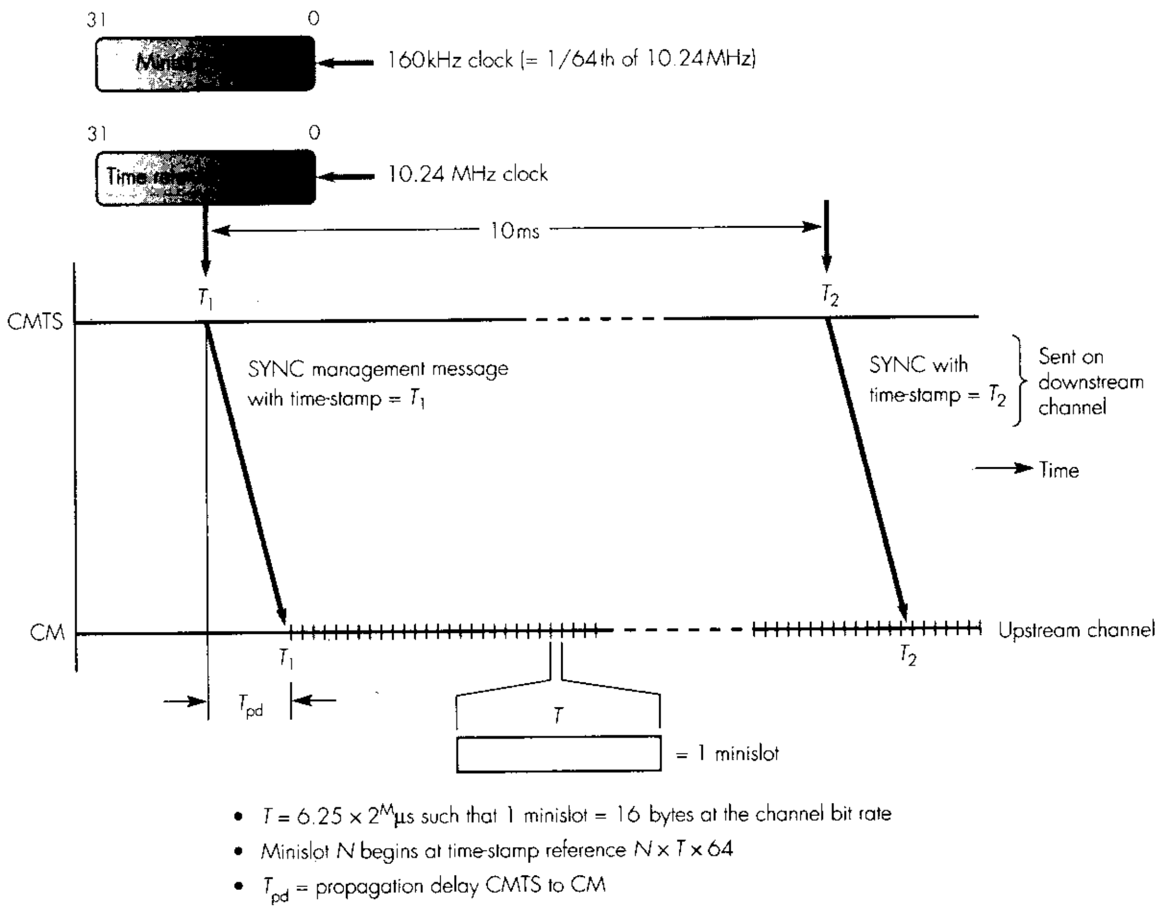

Figure 11.6 Cable MAC time synchronization principles.

bandwidth allocation MAP. As we indicated earlier, this ensures that the transmissions from all the CMs in a transmission interval arrive at the CMTS as if all the CMs were colocated with the CMTS. The principle behind the ranging procedure is shown in the time sequence diagram in Figure 11.7.

At periodic (typically 2 ms) intervals, the CMTS includes in the bandwidth allocation MAP it is transmitting, a reserved time interval – transmission opportunity – known as an *initial maintenance (IM) region.* This is present to enable a new CM to carry out the ranging procedure. Once the CM is in time synchronism, it searches each MAP it receives from the downstream channel for an IM region and, when it finds one, it immediately transmits a *Ranging Request (RNG-REQ)* message with its own (physical layer) fixed delay within it. The duration of this region is such that it can compensate for the worst-case round-trip propagation delay between a CM and the CMTS.

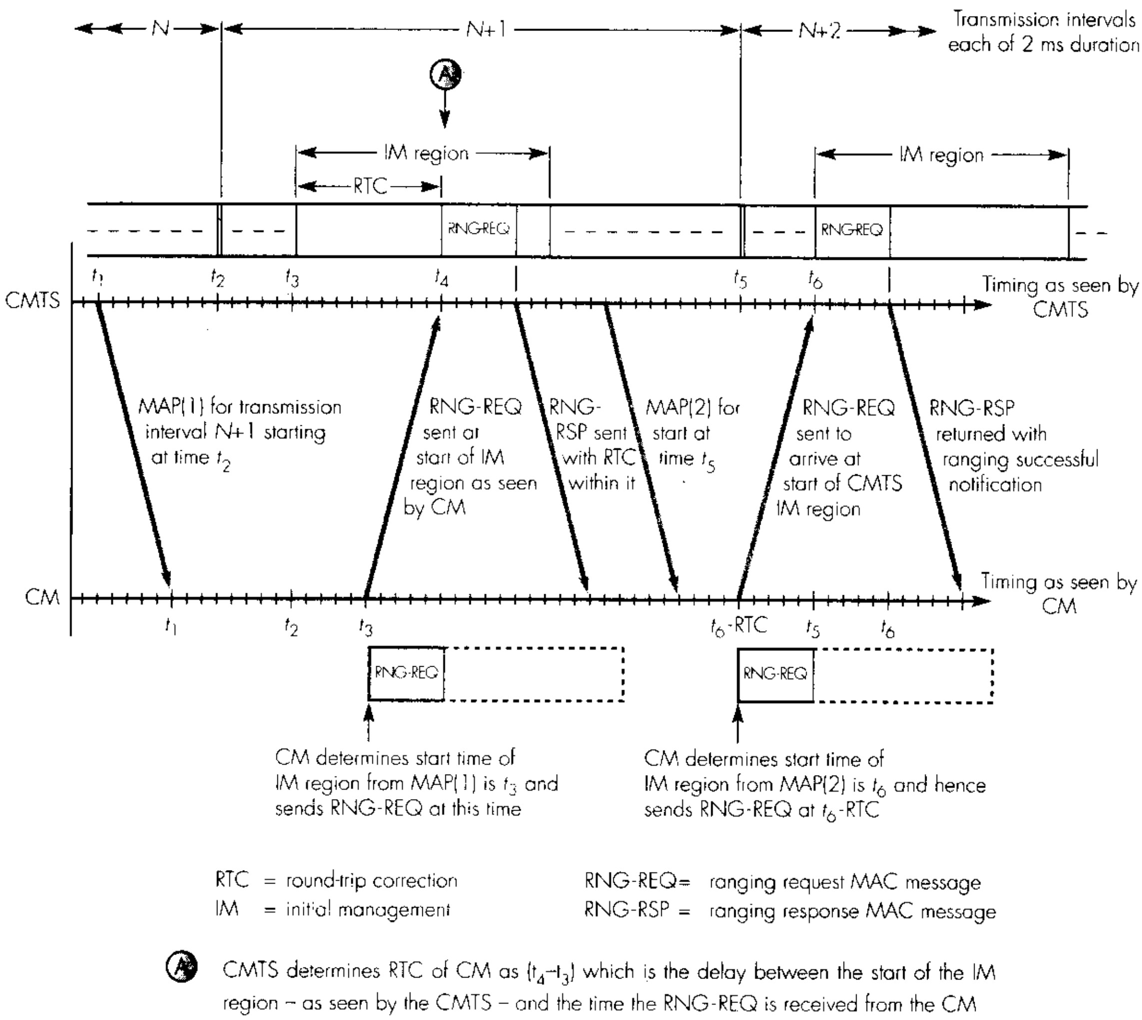

Figure 11.7 Cable MAC ranging procedure principles.

On receipt of the RNG-REQ message, since the CMTS knows the time offset from the start of the current transmission interval to the start of the IM region, the CMTS can estimate the round-trip propagation delay from the CM to the CMTS and back again by the time delay the RNG-REQ message is received relative to the start of the IM region. The CMTS responds by returning a *Ranging Response* (*RNG-RSP*) message in the downstream channel that contains the initial computed timing offset – including the CM fixed delay – to be used by the CM. On receipt of this, the CM first proceeds to locate another IM region in the next MAP it receives. It then sends a second

RNG-REQ message to the CMTS but this time earlier than the start of the IM region – as seen by the CMTS – by a time equal to the initial RTC that it received in the RNG-RSP message. If this arrives at the CMTS at the start of the IM region, then the estimated RTC is correct and the CMTS returns a second RNG-RSP message containing a *Ranging Successful* notification within it. Alternatively, if the arrival of the second RNG-REQ is still later than the start of the IM region, a further REQ/RSP message interchange takes place using a revised RTC time that is indicated in the second RNG-RSP message. This procedure continues until the precise RTC is known and the CM receives a successful notification.

Once a notification success has been received, the CM can then proceed to send frames containing IP packets using the request/grant cycle outlined earlier. First the CM carries out the configuration operation we described in the previous section and this concludes with the CM being registered and allocated a service identifier (SID). It is at this point that the CM can start to relay frames – containing IP packets – to and from the stations that are attached to the hub ports of the CM.

Frame transmission

As we indicated earlier, all transmissions over the upstream channel are controlled by the CMTS using the bandwidth allocation MAP. Prior to sending a frame (containing an IP packet), the CM sends a bandwidth Request (REQ) message to the CMTS indicating the number of minislots required to transmit the frame. The CMTS then responds by including in the next allocation MAP it broadcasts on the downstream channel, a Grant message containing the SID of the CM and the number of minislots within the next transmission interval the CM can use.

Since the Request message must also be sent to the CMTS over the upstream channel, each bandwidth allocation MAP includes a reserved region – similar to the IM region used to send a Ranging Request message – to enable a CM to send a Request message during the current transmission interval. This is called the *REQ region* and, since there may be a number of CMs waiting to send a Request message, a contention resolution procedure must be used to enable each competing CM to share the use of the request region in an equitable way.

Each REQ region can support the transmission of multiple Request messages. Also, as we shall expand upon later, each Request message is 6 bytes long and includes the SID of the CM making the request and the number of minislots requested. The contention algorithm is based on a truncated binary exponential backoff algorithm similar to that used with Ethernet. In the header of each bandwidth allocation MAP is a pair of fields relating to the algorithm, one called the *data backoff start* (*DBS*) and the other the *data backoff end* (*DBE*). Each is a power-of-two value and, in the case of the DBS, a value of 2, for example, indicates a backoff value in the range 0–3, a value of 3, 0–7 and so on. The range is known as the **backoff window** and must always be less than the value derived from the DBE. The latter is known, therefore, as the **maximum backoff window**.

When a CM wishes to make a request, on receipt of the next MAP, it first reads the DBS value from the MAP and proceeds to compute a random number within the derived backoff window. This indicates the number of request opportunities (6-byte slots) from the start of the REQ region the CM must defer – and hence not use – before transmitting its own Request message. For example, if the computed random number from the current backoff window is 3 and the REQ region can support 6 requests, then the CM must defer from using the first 3 request opportunities/slots before transmitting its own request. Alternatively, if the computed random number is 8, the CM must defer from using all 6 request slots and wait for the next REQ region. On determining this from the MAP, the CM must then defer from using a further 2 slots before sending its request.

Once a CM has sent a request, it must then wait for a data *Grant* message in a subsequent MAP containing the SID of the CM and the number of minislots allocated. Providing this is received, then the Request message was received successfully by the CMTS and no collision occurred. If a Grant is not received, then the CM must try again. It first doubles its current backoff window and, providing the size of the new window is less than the maximum backoff window, the CM proceeds to compute a new random number within the limits of the new window. A maximum retry limit of 16 is used and, if this or the maximum backoff window is reached, the CM discards the frame.

This mode of operation is known as the **reservation access mode** and an example transmission sequence illustrating this is given in Figure 11.8. In this example it is assumed that the REQ region contains 4 request slots, the random number generated from the initial backoff window is 2, and the Request message is successfully received by the CMTS. Note that each transmission made by the CM is sent earlier than its allocated time by the round-trip correction time computed during the ranging procedure.

Frame formats

All cable MAC frames have a standard 6-byte header the format of which is shown in Figure 11.9(a). The *frame control* (*FC*) field identifies the type of MAC header and consists of three subfields. The *FC-TYPE* indicates whether the MAC frame contains a user data frame – containing an IP packet – or a cable MAC frame. For the latter, the *FC-PARM* identifies the type of MAC frame; for example, a Request message/frame or a MAC management frame such as a time synchronization frame. The *EHDR-ON* is a single bit and is set to 1 if an extension header is present. We shall give an example of the use of extension headers in the next subsection.

The *MAC-PARM* field in a REQ frame indicates the number of minislots requested or, if an extension header is present, it indicates the number of bytes in the extension header. The *LEN/SID* is either the sum of the number of bytes in the extension header (if present) and the number of bytes following the HCS field or, in a REQ frame, the SID of the requesting

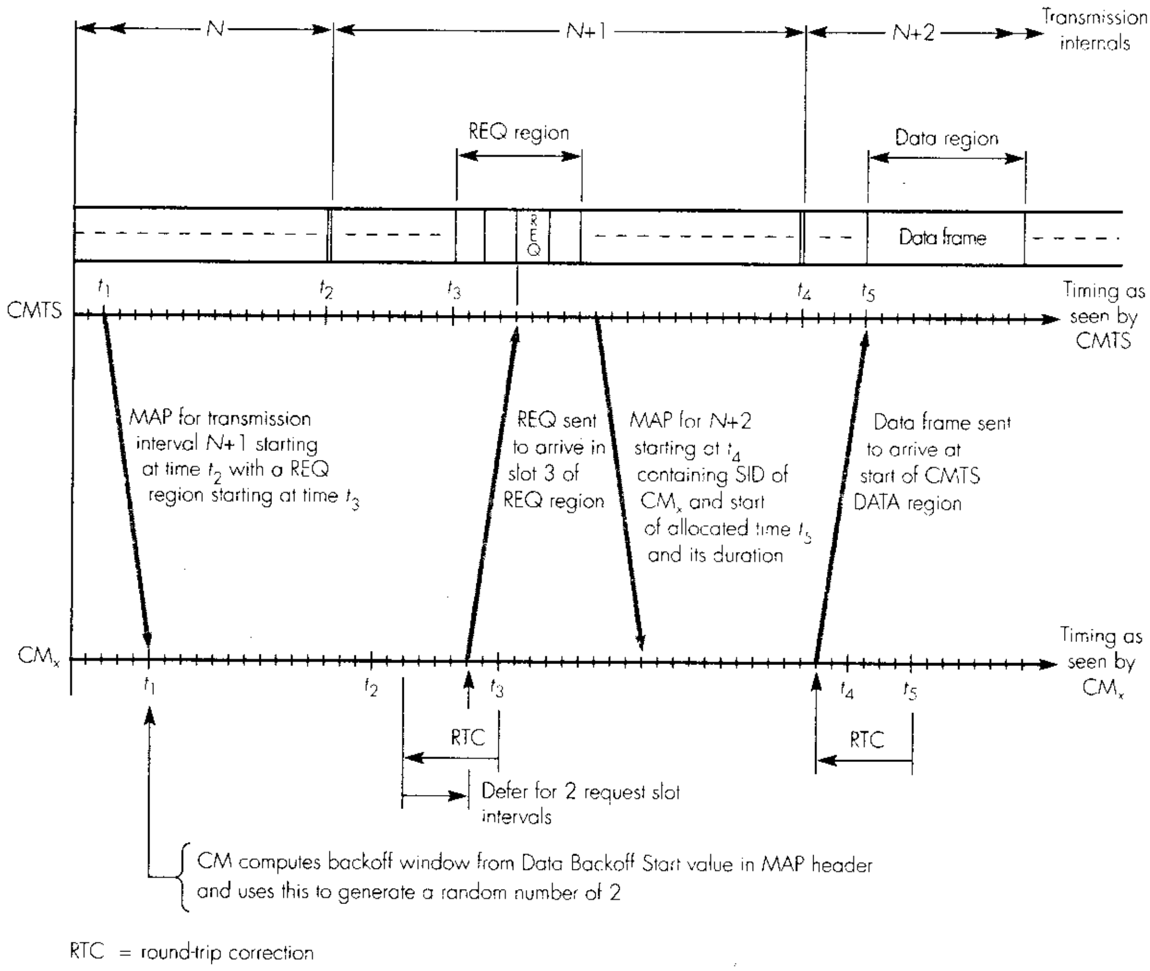

Figure 11.8 Cable MAC reservation access mode of transmission.

CM. The *HCS* is a 16-bit CRC that is used to detect transmission errors in the MAC header. It is used to ensure the integrity of the header fields and, in the contention mode, whether a collision has occurred.

Two example MAC frame types are shown in Figure 11.9(b). The first is the format of a MAC frame containing a user data frame. As we saw in Figure 8.3, for an Ethernet/IEEE 802.3 LAN the length of a frame – including a 4-byte FCS – can be from 64 to 1518 bytes and the *LEN* field in the MAC header indicates the number of bytes in the actual frame being transferred. The second example is the format of a REQ frame and, since this consists of the header only, the *SID* indicates the 14-bit SID of the CM making the request. The *MAC-PARM* contains the number of minislots that are being requested.

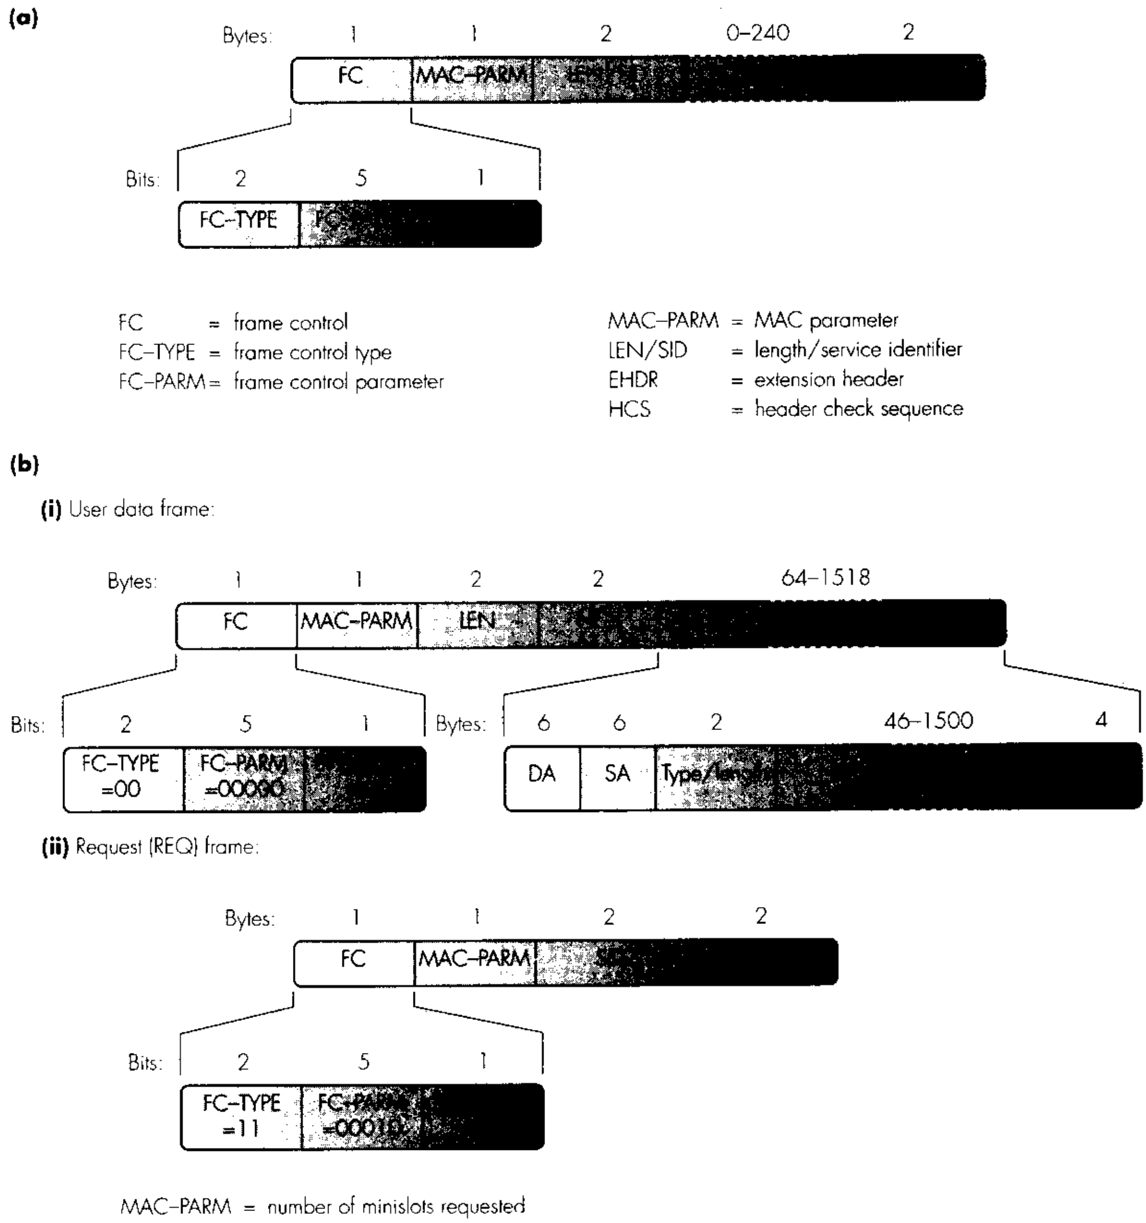

Figure 11.9 Cable MAC frame formats: (a) header field definitions; (b) two example frame types.

Fragmentation

During periods of heavy traffic on the upstream channel, a grant message may specify a smaller number of minislots than have been requested by the CM to transfer the user data frame. When this happens, the MAC layer within the CM automatically fragments/splits the user data frame into a number of smaller fragments for transmission over the cable. At the CMTS, the peer MAC layer then reassembles the fragments to form the initial frame before forwarding this to the Internet gateway.

In order for the cable MAC in the CMTS to determine that a received MAC frame contains a fragment of a larger frame, an extension header is used. The format of each fragment header is shown in Figure 11.10(a) and an example is given in Figure 11.10(b).

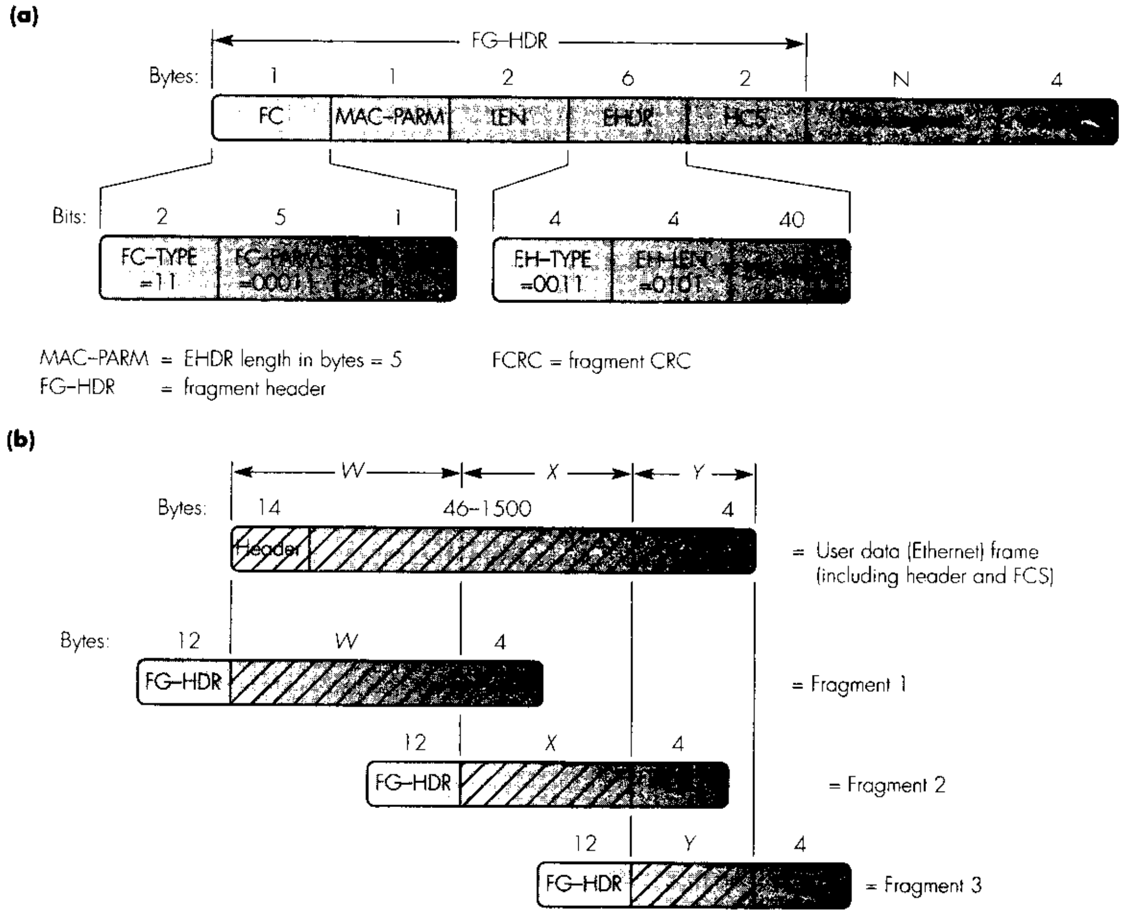

Figure 11.10 User data frame fragmentation principles: (a) fragment header format; (b) example fragmentation overheads.

In general, the EHDR field may contain multiple variable-length fields each of which is defined in a type-length-value (TLV) format. In the case of a fragment EHDR, however, only a single fixed-length field is present. The 4-bit *EH-TYPE* is set to 3 which indicates the MAC frame contains a data fragment. The 4-bit *EH-LEN* is set to 5 and the following 5-byte (40-bit) *EH-VALUE* field then contains a number of fixed-length subfields. These include the 14-bit SID of the CM carrying out the fragmentation, a 1-bit *first fragment* and a 1-bit *last fragment* – each of which is set to 1 for the first and last fragment respectively – and a 4-bit *fragment sequence number*. This is initialized to zero and is incremented by one for each fragment that is sent. Collectively these fields enable the peer cable MAC in the CMTS to reassemble the fragments into the initial user data frame. The 4-byte FCRC is a 32-bit CRC of the fragment. Note that the initial user data frame to be fragmented includes the various (Ethernet) header fields and the original FCS.

In the example shown in Figure 11.10(b), it is assumed that the initial user data frame has to be sent in three cable MAC frames. As we can see, the overheads associated with each of these frames is 16 bytes and hence it is assumed that the first grant message contains sufficient minislots to send $W+16$ bytes, the second $X+16$ bytes and the third $Y+16$ bytes.

Piggyback requests

Also present in the extension header (EHDR) field of each fragment is an 8-bit *REQ* (request) subfield. This plays the same role as the Request message we described earlier. The difference is that there is no contention involved in sending the subsequent request message(s) for the additional minislots that are required to send the remainder of the user data (Ethernet) frame.

When the cable MAC in the CM prepares the first fragment of a larger frame, it computes the number of minislots required to send the remainder of the frame including the 16 bytes of overhead. It then includes this value in the REQ subfield of the frame. This type of request is known as a **piggyback request** since it is being sent in the same frame as the current fragment of data being transferred. As we can deduce from this, this procedure leads to a shorter wait time relative to sending a separate request message. Again, if the number of minislots allocated in the grant message contained in the subsequent bandwidth allocation MAP is less than the number requested, then a further level of fragmentation must be carried out using the same procedure as before. This is repeated until the complete user data frame has been transferred.

Request/data regions

During periods when the upstream channel is lightly loaded, the CMTS may include in each bandwidth allocation MAP it transmits a *Request/Data region*. This is similar to a REQ region except that in addition to Request messages, a CM can also attempt to send a short user data frame. As with a REQ region, this is a contention region and hence the data frame may be corrupted by a simultaneous transmission of a request or data frame from another CM. In order for the CM to know whether a collision has occurred, if the data frame is received without errors, the CMTS returns a *Data Acknowledge* message (containing the SID of the CM) in the next MAP it transmits. In order for the cable MAC in the CM to know the maximum time to wait for an acknowledgment, the CMTS includes in each MAP it transmits an ACK field which indicates the start time of the transmission region that any acknowledgment messages in the MAP relate. If this is later than the time the data message was sent, then a collision is assumed and the CM must try again.

QoS support

The reservation access mode described in the preceding subsections is intended for the transfer of user data frames containing IP packets relating to the best-effort service. Hence all the packets have the same priority value in the type-of-service (ToS) field in the packet header. An additional feature is also included in the cable MAC, however, to enable frames containing packets which have a high priority to be transmitted before frames containing a lower priority packet.

As we saw in Section 1.5.6, the relative priority of a packet is determined by the service class associated with the application/call. There are a number of different service classes each determined by a set of QoS parameters. These may include, for example, a defined worst-case end-to-end delay, jitter, and throughput requirement. In the context of the cable MAC, each service class has a related **service flow** and, for each service flow, the CMTS endeavors to schedule transmissions so that the agreed QoS parameters associated with each service class are met. As we can deduce from this, there are a number of different service flows each of which maps to a related packet priority. Hence packets containing a particular priority value are transmitted according to the rules associated with the related service flow.

Each service flow is characterized by a separate **service flow identifier** and, for applications involving duplex flows – such as telephony – a separate service flow is set up in both the CMTS-to-CM and CM-to-CMTS directions. Also, for applications requiring asymmetric communication channels – that is, the class of service associated with the information flow is different for each direction – the service flow in each direction can be different.

In order to ensure that the agreed QoS parameters associated with each service flow are met, in the upstream direction, a number of additional transmission opportunities are provided by the CMTS for these flows on receipt of the related request from a CM. These include:

- **unsolicited grant**: these are intended for use for service flows involving packets containing real-time information. The CMTS, on receiving the appropriate request message from a CM, reserves in the allocation MAP specific fixed-sized transmission opportunities that are repeated at periodic intervals. The frequency and duration of the intervals is determined by the type of service flow. The CMTS stops generating intervals when it detects they are not being used;

- **real-time polling**: these also involve the CMTS reserving periodic transmission opportunities and are intended for real-time traffic flows such as voice-over-IP (Internet telephony). The CMTS periodically polls/invites each CM with an active service flow to make a bandwidth request at intervals of about 1 ms. The related transmission opportunities cease when the CM fails to respond to a poll.
- **unsolicited grant with activity detection**: the packetization process associated with voice-over-IP sometimes exploits the silence periods between talk spurts by ceasing to send packets during these periods. This service is intended for use with this type of application. All the time packets associated with this flow are being transmitted by a CM, the CMTS continues to provide periodic transmission opportunities for it. Immediately the CMTS detects these are not being used, however, it stops providing them. Then, when the CM detects the flow of packets resumes – and hence the service flow becomes active – it sends a request message to the CMTS asking for the unsolicited periodic transmission opportunities to be resumed. Normally, the CMTS, after it stops providing transmission opportunities, provides a specific request opportunity for this CM/SID at similar periodic intervals;
- **non-real-time polling**: this service is intended for use with non-real-time applications that involve the transfer of large volumes of data; for example large file transfers. With the basic reservation access mode, during periods of heavy loading such transfers may take unacceptably long times to complete. Hence by using non-real-time polling, the CMTS periodically polls/invites CMs to make a bandwidth request at intervals of about 12 s. Such requests always receive some reserved transmission opportunities even during periods of heavy traffic when contention requests are receiving minimal transmission opportunities.

DS TC sublayer

The MAC frames exchanged between the CMTS and a CM have the same basic structure in both the downstream and the upstream channels. In the downstream direction, however, in order to utilize similar receiving hardware to that used for a digital TV channel, the channel bitstream is divided by the transmission control (TC) sublayer into a stream of 188-byte packets. As we showed earlier in Figure 5.20 and explained in the accompanying text, this is the packet format used by the transport multiplexer to transmit multiple digital TV programs over a single high bit rate channel. Each packet is made up of a 4-byte header and a 184-byte payload. The first byte of the header is for synchronization purposes and enables all the receivers to determine the start of each new packet. A second 13-bit field in the header (known as the *payload identifier* (*PID*)) is then used to identify the type of contents in the payload and, if this is data relating to Internet access, it is set to 1FFE (hex). Hence if the total channel bandwidth is being used for Internet access, all packets will contain this value in the packet identifier field. Alternatively, if the channel is being used to transmit both Internet data and digital TV programs, then only those packets containing Internet data will have this PID.

The 184-byte payload field of each packet is used to transfer the cable MAC frames. However, since a packet may contain either multiple short frames or a portion of a long frame, in order to use the available bandwidth efficiently, it is necessary for each receiver to be able to determine the start and end of each frame. Also, since there may be unused bits/bandwidth between successive frames, a means of detecting when these are present is required.

To determine the start of a frame within a packet, the first byte following the 4-byte packet header is used as a *pointer* (offset) to the first byte of a new frame should one start within the current packet. If a new frame does not start within a packet, however, and the packet contains, say, a portion of a packet that straddles multiple packets, then the first byte is not used as a pointer. For each CM to determine if a pointer byte/field is present, a single bit in the 4-byte packet header called the *payload unit start indicator* (*PUSI*) bit is set to 1. Idle periods between frames are indicated by the presence of one or more *stuff-bytes* which have the reserved bit pattern of FF (hex). The *frame control* (*FC*) field at the head of each frame cannot be equal to this. Hence, since the length of each frame can be determined from the fields in the (MAC) frame header, each CM can readily determine the start and end of each frame. Three examples are given in Figure 11.11.

The first assumes the start of the frame immediately follows the pointer byte and just a single frame is present in the packet. The second assumes multiple short frames are packed into a single packet and the third, a single long frame spans multiple packets.

DS/US PMD sublayers

A schematic diagram showing typical packet flows in both the downstream and upstream channels is given in Figure 11.12.

In the downstream direction, the signal is continuous and is the allocated carrier for the channel modulated by the bitstream using either 64 or 256-QAM. As we explained in the previous subsection, the bitstream is divided into a contiguous stream of 188-byte MPEG-2 transport packets each with a 4/5-byte header.

In the upstream direction, the signal is the allocated carrier modulated by the bitstream using either QPSK or 16-QAM. The bitstream in this direction, however, is made up of variable-length bursts of bits/frames each sent by a different CM. Hence in order for the PMD sublayer in the CMTS to be able to receive reliably each frame, a *guard-band* between each frame comprising several symbols at the start and end of each frame is present. Also, in order for the receiving electronics to achieve clock/symbol synchronization and determine the start of each frame, a *preamble* sequence of up to 1024 bits precedes each frame starting at the allocated minislot boundary. This is sent as either 512 (QPSX) or 256 (16-QAM) symbols and terminates with a defined symbol pattern. This is followed by the bits in the (MAC) frame and, since the

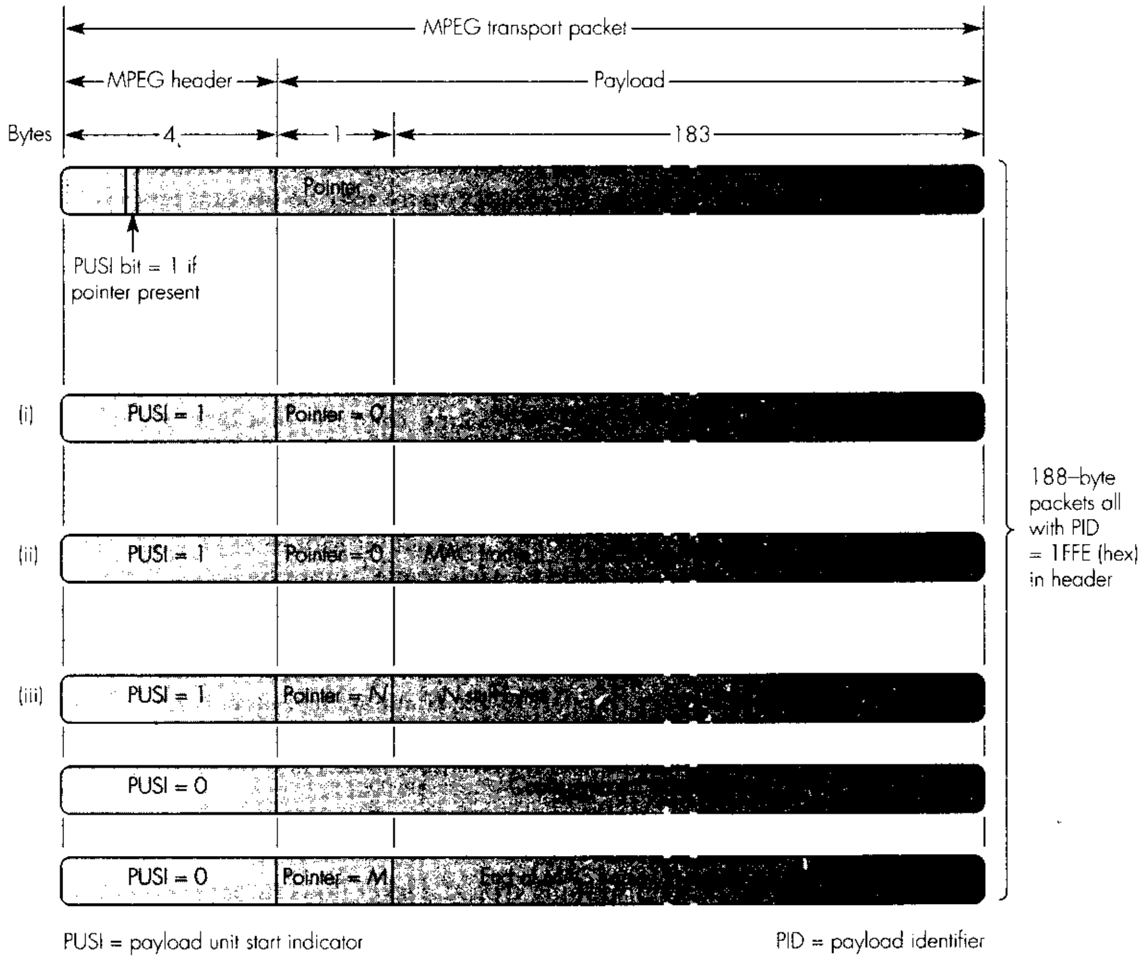

Figure 11.11 Downstream transmission convergence sublayer: examples showing the packing of MAC frames into 188-byte MPEG transport packets.

length of a frame can be determined from the fields in the frame header, no end-of-frame sequence is required.

In addition, as we show in the figure, a forward error control (FEC) scheme is optionally used with both the downstream and upstream channels in order to reduce the probability of the received bitstream containing transmission/bit errors. We discuss the principles of such schemes in Appendix B. The scheme used is based on an (n, k, t) Reed–Solomon (RS) code where k is the number of bytes in the original block of data, n the number of bytes in the block after coding, and t the number of bytes in a block that the code will correct. Normally, in order to exploit the use of the same hardware as is used for digital TV, the block size used in both channels, k, is 188 bytes. This has 16 bytes

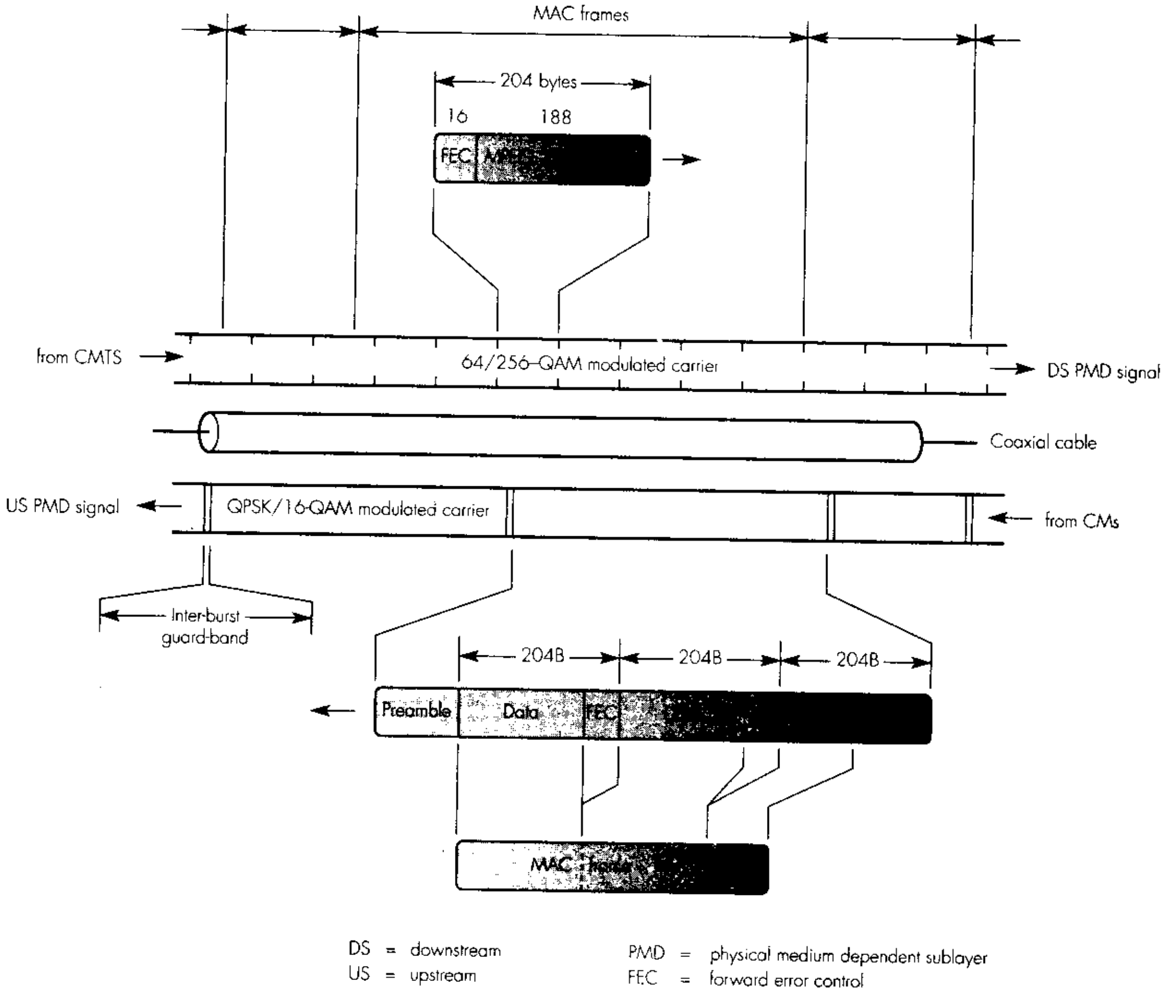

Figure 11.12 DS/US PMD sublayer: example packet flows and their overheads.

of FEC and hence $n = 188 + 16 = 204$ bytes. This code has a minimum *Hamming distance, d,* of 17 bytes and hence will correct up to 8 bytes in each block. This code is written as an RS (204, 188, $t = 8$) code and, in practice, is a shortened version of the RS (255, 239, $t = 8$) code. The FEC bytes are computed by adding 51 zero bytes before the 188 bytes prior to computing the FEC.

Cable intranets

The network architecture we showed earlier in Figure 11.3 is an example of a typical regional HFC cable network. The CMTS in the cable headend provides a single point of access to the Internet and the cable operator acts as an Internet service provider (ISP) for its CM subscribers. Many cable operators, however, have multiple regional networks. Hence, in terms of Internet access, these operators have many thousands of subscribers. In such cases, instead of providing multiple single access points to the Internet, the cable operator creates its own intranet with all the regional headends linked together using high bit rate digital circuits. The general scheme is shown in Figure 11.13.

As we can see, each regional network is similar to a site LAN in a large enterprise intranet. As we saw in Figure 11.4, in the CMTS is a fast-Ethernet switch and, connected to this, are a number of additional local servers that

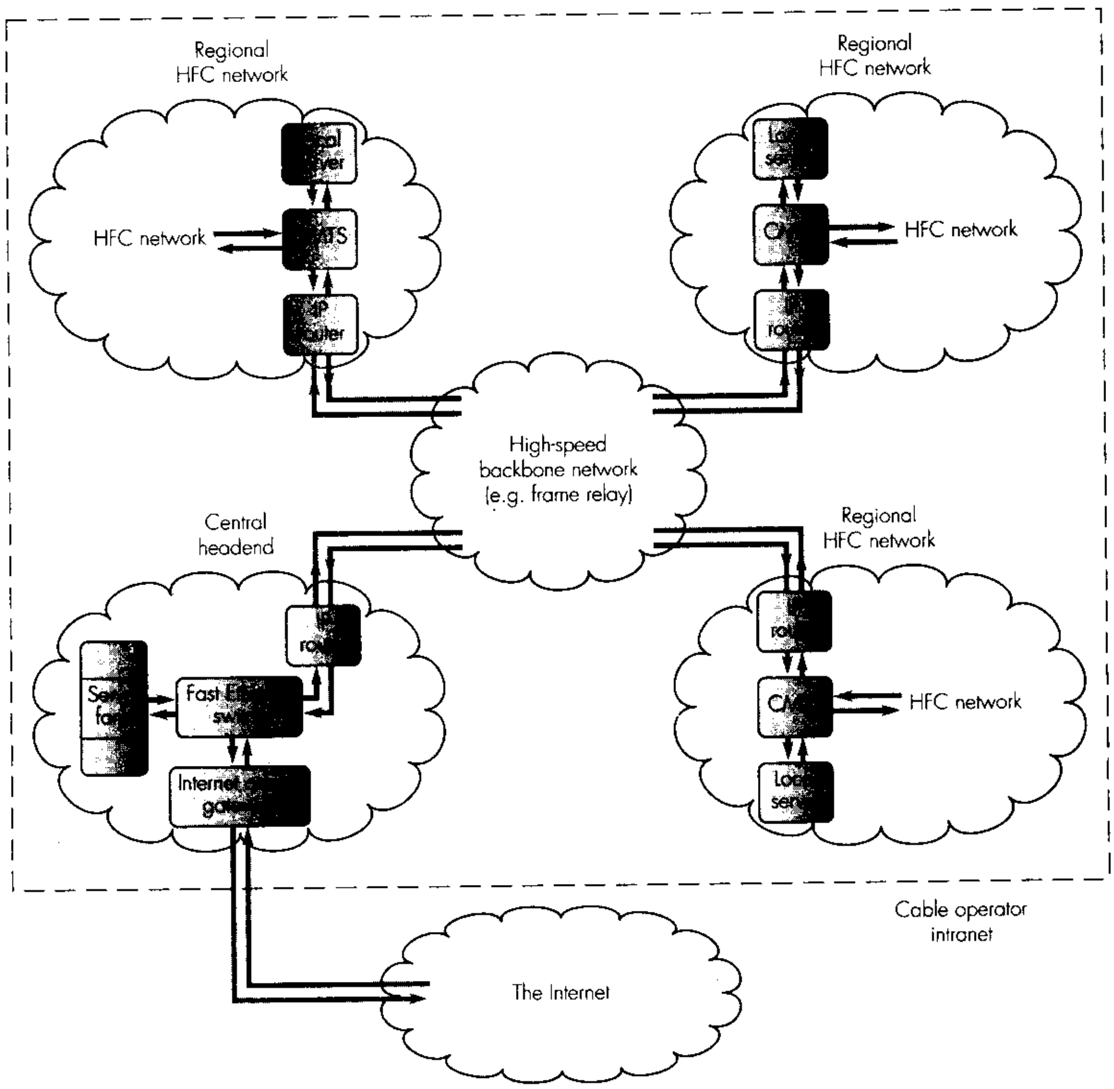

Figure 11.13 Cable operator intranet principles.

provide various services on behalf of its local community of subscribers. For example, an email server, some local Web servers, and one or more cache servers that are used to retain copies of popular Web pages.

When part of an intranet, however, also connected to the switch is a router that forms the interface with other sites. In this way, all communications between subscribers that are connected to the different regional networks of the operator are via the intranet and hence are carried out at high speed. Also, a single high-speed backbone connection with the public Internet can be used so increasing the speed of access to external servers.

MMDS and LMDS

In areas with low subscriber densities or where laying cable is difficult, an alternative to coaxial cable is to use terrestrial microwave broadcast transmissions. As we show in Figure 11.14(a) and (b), two alternatives are available, one called the **multichannel multipoint distribution system** (**MMDS**) and the other the **local multipoint distribution system** (**LMDS**). Both use ommidirectional transmitters and, within their area of coverage, provide a similar range of services to those provided with a coaxial cable distribution network. The main difference between the two systems is the area of coverage of the transmissions and the number of channels supported. The main features of both systems are summarized in Figure 11.14(c).

As we can see, a typical MMDS operates in the 2.15–2.7 GHz frequency band. Normally, it has a relatively wide area of coverage with a direct connection between the transmitter and the cable headend. Multiple 6/8 MHz downstream channels are supported which are used for either analog TV broadcasts or, with suitable modems, digital broadcasts. The latter include digital TV and, together with the upstream (return) channel, interactive services including access to the Internet.

An LMDS operates in a higher frequency band as it is intended for local distribution with an area of coverage of up to 5 km from the transmitter. Normally, the transmitter has a direct connection to a fiber node in an HFC network. Hence an LMDS effectively replaces the coaxial cable part of an HFC network. A similar bandwidth to coax is supported and hence a similar range of services to an HFC network.

11.3 Satellite television networks

Satellite networks have been used for many years to deliver broadcast television direct to the home. Initially, the signals relating to all the TV programs were analog but these have now been complemented with digital TV. Also, as with cable networks, in addition to broadcast television, various interactive data services are now supported. In this section we discuss the operation of, firstly, a direct-to-home broadcast TV system, secondly, the principles behind digital TV over satellite and thirdly, how interactive services are provided.

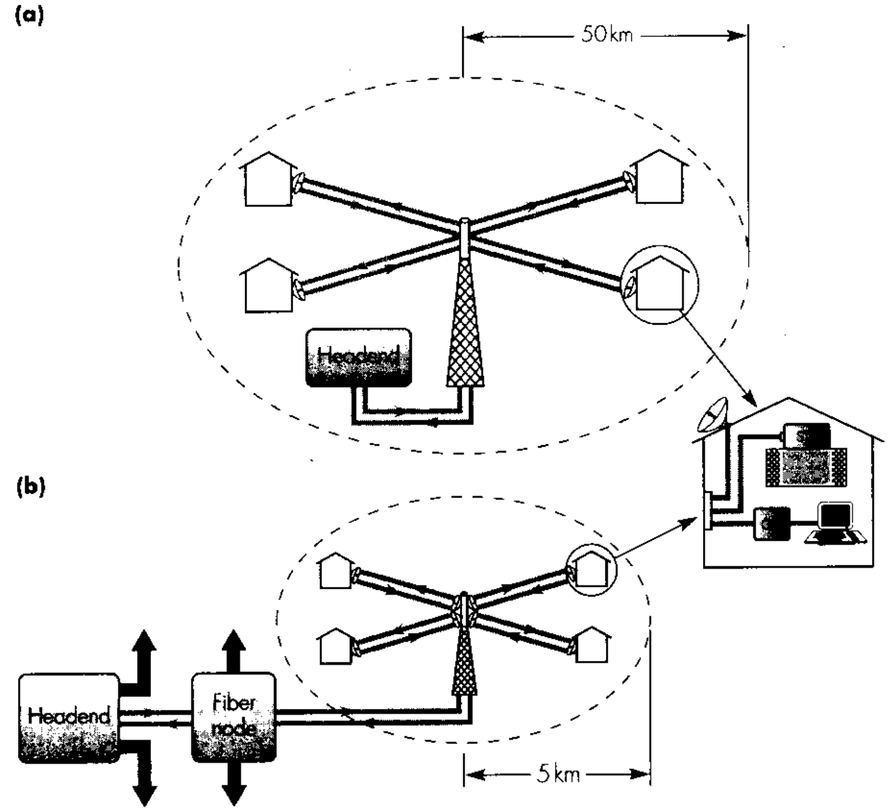

(c)

System	Frequency band (bandwidth) Downstream	Frequency band (bandwidth) Upstream	Area of coverage
MMDS	2.15–2.348 GHz (198MHz)	2.66–2.70 GHz (4MHz)	50km radius with 60cm (24") dish
LMDS	27.5–28.35 GHz (850MHz)	29.2–29.35 GHz (150MHz)	5km radius with 30cm (12") dish

M/LMDS = multichannel/local multipoint distribution system

Figure 11.14 M/LMDS principles: (a) MMDS schematic; (b) LMDS schematic; (c) typical operating parameters.

11.3.1 Broadcast television principles

The basic requirement of an entertainment satellite network is to broadcast a set of TV programs from a program source to the set-top boxes of a large number of subscribers who are physically distributed over a wide geographical area. Figure 11.15(a) shows how this is achieved.

The orbital period of a satellite varies according to its distance above the earth's surface. Satellites used for TV broadcasting, however, are **geosynchronous**

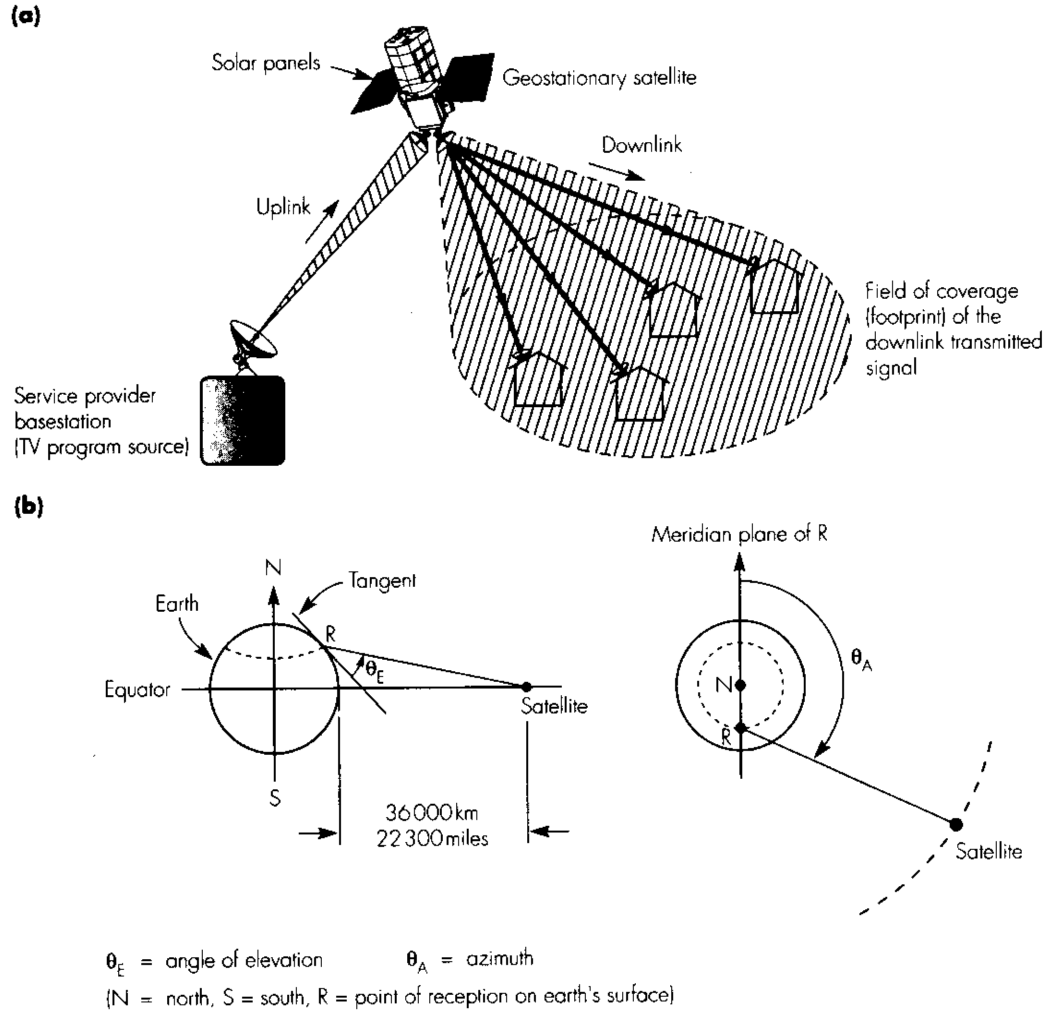

Figure 11.15 Geostationary satellite principles: (a) broadcast network schematic; (b) positioning details.

which means that the satellite orbits the earth once every 24 hours – slightly less than this in practice – in time synchronism with the rotation of the earth. To achieve this, the orbit must be circular and, since the earth rotates around its polar axis, the satellite's orbit must be around the earth's polar axis. A circular orbit is then obtained if the orbit is in the equatorial plane and at an (average) altitude above the equator of approximately 36 000 km/22 300 miles. At this height, the effect of the centrifugal force due to the rotation of the satellite is negated by the terrestrial attraction of the satellite by the earth. Also, from a point on the earth's surface, the satellite appears stationary and hence is known as a **geostationary earth orbit** (**GEO**) – or simply geostationary – satellite.

GEO satellite positioning

The satellite is first launched to an altitude of 36 000 km/22 300 miles above the equator and on-board motors are then used to ensure the orbit is circular and in its assigned position relative to a point on the earth's surface. The on-board motors are also used at periodic intervals during the satellite's operational lifetime to effect small corrections which maintain the satellite's position to within 0.2° of its assigned position. In general, therefore, the useful lifetime of a satellite is determined by the quantity of propellant used to perform these maneuvers. The main power source for the on-board electronic equipment comes from large solar panels but this has to be complemented by batteries to overcome the loss of power from the panels during periods when the satellite is eclipsed by the earth.

Since the latitude of a satellite is 0°, the position of a satellite is defined by its longitude relative to the Greenwich meridian. As we show in Figure 11.15(b), the position relative to a point on the earth is defined by the elevation and azimuth of the satellite. The **elevation** is the angle between the line from the satellite to the point of reception, R, and the tangent to the earth's surface at R. The **azimuth** is the angle between the N–S plane passing through R and the line from the satellite to R measured relative to the north pole.

Frequency allocations

As we saw in Figure 11.2, the bandwidth of an analog TV signal – including guard bands – is either 6 or 8 MHz. As we show in Figure 11.16, with the earliest satellites used for commercial analog television broadcasts, the signal of each TV program is frequency modulated onto a separate carrier which results in a basic channel bandwidth of 36 MHz. A satellite supports multiple channels and these are combined to form the signal that is transmitted from the basestation to the satellite using frequency division multiplexing. Typically, 24 channels are used with a guard band of 4 MHz between adjacent channels.

In order to avoid the signal received from the basestation on the uplink interfering with the signal transmitted by the satellite on the downlink, a separate frequency band is used for transmissions in the uplink and downlink directions. This means that each channel is allocated a separate carrier signal in both the uplink and downlink directions with the same fixed spacing between channels.

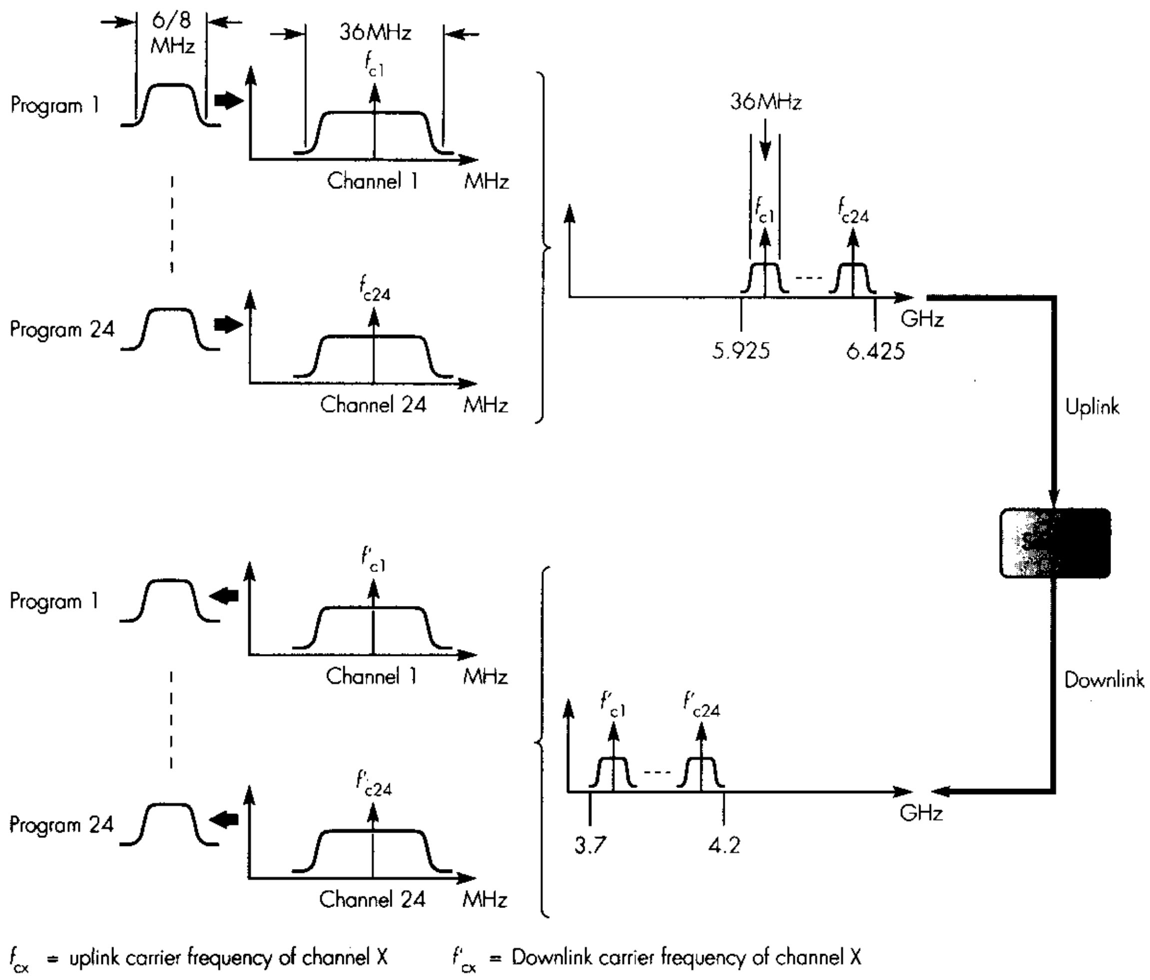

Figure 11.16 Frequency bands used with early satellite systems.

Antenna designs

Owing to the large distance travelled, the signal received from the transmitter is extremely weak. Hence it is essential to receive as much of the transmitted signal as possible. This is achieved by using a parabolic dish made of a reflecting material. Like light, microwaves travel in straight lines and are reflected by metallic surfaces. Hence as we show in part (i) of Figure 11.17(a), the incident waves received from the transmitter are all focussed at the focal point of the dish. Moreover, since the distance travelled by all the waves is the same, they all arrive at the focal point in-phase and their energy, therefore, is additive.

Located at the focal point of the dish is a unit known as a **low-noise block converter (LNB/C)**. This consists of a low-noise amplifier and a frequency

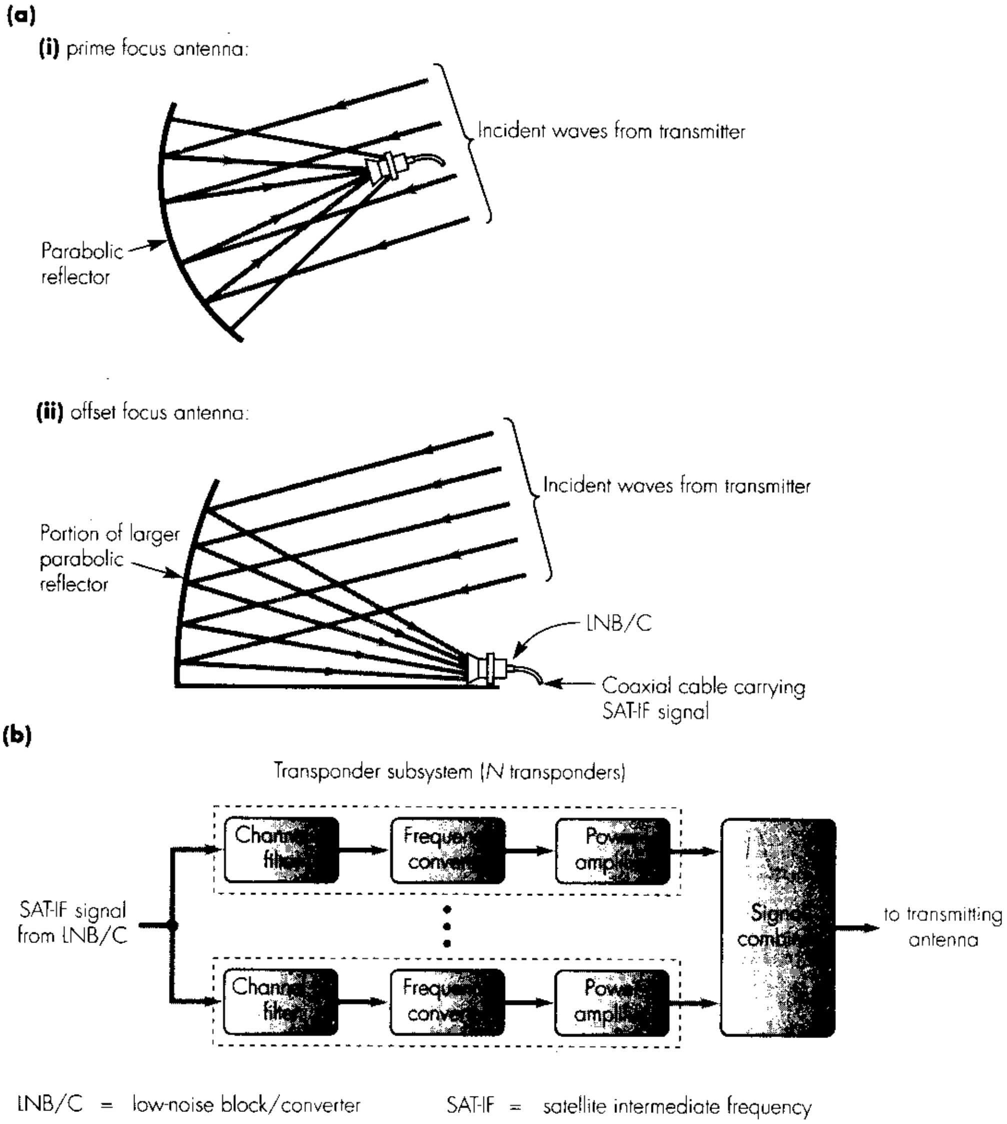

Figure 11.17 Satellite components: (a) two antenna designs; (b) on-board transponder subsystem.

downconverter. As we saw in Figure 11.16, the frequency of the uplink signal with early satellite systems is in the range 5.925 through to 6.425 GHz. Since this exceeds the bandwidth of a coaxial cable, prior to passing the received signal to the on-board electronics for frequency conversion, the received signal is first amplified in the LNB/C and then downconverted into a lower frequency band. This is called the **satellite intermediate frequency** (SAT-IF) and, with early systems, is 950 MHz. It is this signal that is then passed to the on-board electronics by means of a coaxial cable.

The antenna design shown in part (i) of Figure 11.17(a) is called a **prime focus antenna** since the LNB/C is located in the center of the dish at its focal point. The disadvantage of this design is the LNB/C inhibits the direct waves in the center of the dish from being collected. Clearly this reduces the efficiency of the antenna and, for small dishes, the fall in efficiency can be significant. To overcome this effect, the reflector used with more modern antennas is a portion of a larger parabolic mirror. This type of design is known as an **offset focus antenna** and its principle of operation is shown in part (ii) of Figure 11.17(a). As we can see, with this design the LNB/C is attached to the bottom of the dish and hence out of the way of the incoming waves. This has the effect of increasing significantly the antenna's efficiency.

Transponder subsystem

The signal output by the LNB/C is passed to an electronic module within the satellite. This is known as the transponder subsystem and its composition is shown in Figure 11.17(b).

Microwave power amplifiers are only linear over a limited frequency band which dictates that each channel signal is amplified separately prior to its transmission. As we can see in the figure, the individual modulated channel signals are first separated out using filters. Each is then frequency-shifted to its allocated (downlink) frequency band and then amplified. The resulting signals are then combined to form the downlink signal that is broadcast over a defined area known as the satellite's field of coverage or **footprint**. An antenna – satellite dish – on each subscriber premises is then used to receive the broadcast signal.

The signal received by the antenna is first down-shifted in the LNB/C and then passed to the set-top box over a coaxial cable. Electronic circuitry within the STB demodulates the received signal by first filtering out each channel signal. The corresponding carrier signal for each channel is then used to recover the signal of the related TV program that has been selected by the subscriber.

As we saw in Section 6.2.5, satellites use frequencies in the microwave frequency band since these propagate through free space in straight lines and can be focussed into a beam of a defined width. In the uplink direction a narrow beam width is used to ensure the maximum amount of energy in the signal transmitted by the basestation is received by the satellite's receiving antenna. Conversely, a wide beam width is used in the downlink direction to ensure the signal is received by all the antennas within the satellite's footprint. The size of the dish required to receive the signal broadcast by the satellite is determined by the output power of the satellite's transmitter – which is influenced by the number of transponders – and its area of coverage.

Most early systems use a single satellite and, in many instances, cover a wide geographical area with the effect that dishes of between 1 and 4 meters are required. With later systems, however, multiple satellites are used to cover a similar area and hence dishes as small as 45 cm (15 inches) are common. With such dishes, interference-free reception can be obtained if the satellites have orbital positions with a spacing of in the order of 10°.

For most early commercial analog television broadcasts, the frequency bands used are within what is called the **C-band** of the microwave frequency range. As we saw in Figure 11.16, with this band the frequency range of the uplink channels is from 5.925 to 6.425 GHz and that of the downlink channels from 3.7 to 4.2 GHz. Typical channel bandwidths of 40 MHz – including a guard-band – are used. With microwaves, however, it is possible to use the same frequency band twice using both horizontally polarized and vertically polarized transmissions. Hence the total 500 MHz of available bandwidth can support up to 24 active channels and hence transponders.

Normally, in order to improve reliability, most transponder subsystems contain a number of spare transponders (and other units) to replace any that may become defective. A *command and telemetry subsystem* is then used to send commands to the satellite to switch spare units into service should this be necessary. Typically, a satellite has to be replaced after 10 to 12 years of service.

11.3.2 Digital television

Since all satellite transmissions are within an allocated frequency band (within the microwave frequency spectrum), modulated transmission must also be used for the transmission of a digital TV program. As with cable TV, therefore, the bitstream containing the multiplexed set of (digital) TV programs is passed through a (microwave) modem to convert it into an analog signal within the allocated frequency band. However, in order to obtain the same bit error rate probability as that obtained with a cable distribution network, a more robust modulation scheme must be used. Hence instead of 64 or 256-QAM with 6/8 bits per symbol (signal element), the modulation scheme used is QPSK (4-QAM) which, as we saw in Figure 7.6(b), has just 2 bits per symbol.

For most digital TV transmissions, the frequency band used is in what is called the **Ku band** which covers the frequency range from 10.7 through to 14.5 GHz. In the downlink direction, the lowest part of the band from 10.7 through to 11.7 GHz is used mainly for newer analog TV transmissions. For digital TV, example downlink bands are 12.2 through to 12.7 GHz for the North American **digital broadcast satellites (DBS)** and 11.7 through to 12.5 GHz for the European **digital video broadcasting-satellites (DVB-S)**. The DBS system uses three orbital positions with 9° spacing, each allowing full coverage to be obtained. Each position is used by a separate service provider which then uses a number of satellites – normally three – to obtain interference-free reception with small dishes. Typically, the allocated bandwidth is used to provide 32 channels of 24 MHz each supporting a symbol rate of 20 Mbaud. Hence, with QPSK, this gives a typical channel bit rate of 40 Mbps.

The DVB-S system uses two orbital positions with 6.2° spacing, each of which allows full coverage to be obtained. Each position is used by a separate satellite operator which uses two satellites to obtain interference-free reception with small dishes. With this system, the typical allocated bandwidth is used to provide 40 channels of 33 MHz, each supporting a symbol rate of 27.5 Mbaud which, with QPSK, gives a bit rate of 55 Mbps.

Channel interface

A schematic diagram showing the individual blocks associated with each channel interface is given in Figure 11.18(a). Most satellite digital TV transmissions now use the standard, MPEG-2 transport stream multiplex. As we saw in Figure 5.20, the (multiplexed) bitstreams of multiple TV programs are multiplexed together into a single bitstream that is made up of a contiguous stream of 188-byte packets, each comprising a 4-byte header and a 184-byte payload. As we indicated earlier, satellite channels are more susceptible to transmission (bit) errors than cable networks. Hence in addition to using a more robust modulation scheme, a more rigorous forward error control scheme is applied to each 188-byte packet. This involves the addition of check bytes derived using a Reed–Solomon code – as optionally used with cable networks – byte interleaving, and convolutional encoding of the resulting bitstream.

The RS check bytes are added primarily to detect burst errors and the same coding scheme as is optionally used in cable networks is used. As we saw in Section 11.2.1, this is an RS (204, 188) code and the 16 check bytes that are appended to each 188-byte packet enable up to 8 bytes in error in each 204-byte block to be identified and corrected.

As we explain in Appendix B, very long error bursts in a block – that is, greater than 8 bytes – can be broken down into smaller bursts by using a technique known as **interleaving**. Essentially, this involves rearranging the order of transmission of the bytes in each 204-byte block so that an error burst longer than 8 bytes will affect no more than 8 sequential bytes in the original 204-byte block. The principle is shown in Figure 11.18(b).

Prior to transmission each 204-byte block output by the RS coder is fragmented into twelve 17-byte subblocks. The transmitted byte sequence is then the first byte from each of the 12 subblocks, followed by the second byte in each subblock, and so on. The reverse operation is then performed at the receiver to reorder the bytes into their original sequence. In this way, should an error burst of, say, 12 bytes occur within a block during its transmission, this will affect only every seventeenth byte in the original 204-byte block and hence be detected and corrected by the RS decoder.

In addition to burst errors, satellite transmissions are susceptible to randomly distributed single bit errors. To minimize the effect of such errors, the bitstream output by the byte interleaver is passed through a **convolutional encoder**, the principles of which are also described in Appendix B. A typical encoder used with digital TV broadcasts is shown in Figure 11.18(c).

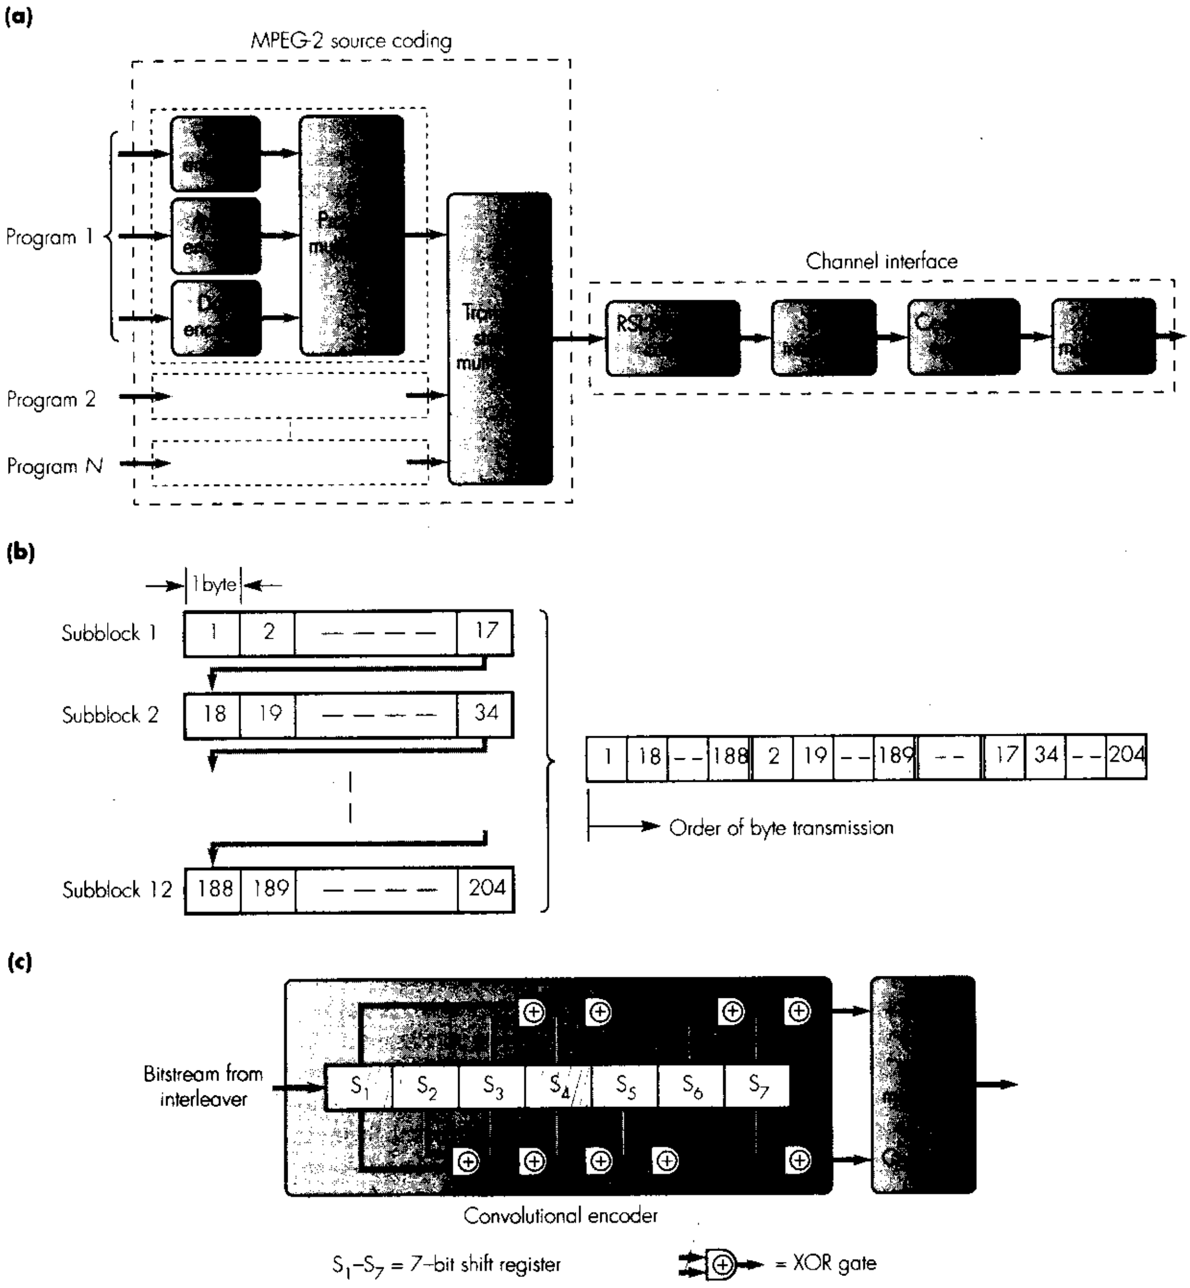

Figure 11.18 Satellite digital television channel interface: (a) schematic; (b) interleaver principle; (c) convolutional encoder.

Essentially, the bitstream output by the byte interleaver is passed through a 7-bit shift register and, for each new bit entering the shift register, two separate XOR operations are performed on the new contents of the register. The pair of bits produced by the two XOR operations are then the two bits that are transmitted for this input bit; that is, for each input bit two bits are transmitted. As we explain in Appendix B, for each type of encoder there is only a limited number of possible pairs of output bits for each new input bit. These are known by the convolutional decoder in the receiver and, for each pair of bits it receives, the decoder computes the *Hamming distance* – see Section 6.6.1 – between the pair of bits it has received and all the known possible pairs. This procedure is repeated for each new pair of bits the decoder receives and, after a predefined number of pairs, the decoder selects the sequence with the smallest Hamming distance as being the most likely sequence of (pairs of) bits that were transmitted. It then replaces each pair of bits in this sequence with the corresponding bitstream that would have produced this sequence.

Since two bits are output for every bit input into the encoder, it is known as a **rate 1/2** convolutional encoder. Hence for every 204-byte block output by the interleaver, 408 bytes are transmitted. It is also possible to operate such encoders at a higher rate by deleting selected bits from the bitstream produced by the encoder. The technique used is called **puncturing** and example rates are 2/3 – 3 output bits for every 2 input bits – 3/4, 5/6, and 7/8. This has the effect, however, of reducing the error correction properties of the code and hence the rate used is a compromise between the level of error correction required and the amount of transmission overheads that are acceptable.

Example 11.1

A digital satellite channel interface uses the MPEG-2 transport stream multiplex, an RS (204, 188) block code, a rate 3/4 convolutional encoder, and QPSK modulation. Determine the number of overhead bits associated with the interface and hence the useful bit rate that is obtained with a channel that operates at (i) 20 Mbaud and (ii) 27.5 Mbuad.

Answer:

MPEG-2 transport stream multiplex comprises a stream of 188-byte packets each of which has a 4-byte header.

With an RS (204, 188) code, for each 188-byte block the RS encoder adds 16 error check bytes.

With a rate 3/4 convolutional encoder, for each 204-byte block the encoder adds 68 bytes of overhead.

Hence total overhead per 184-byte block = 4 + 16 + 68 = 88 bytes.

11.1 Continued

11.3.3 Interactive services

As we saw in Figure 1.16(b) and the accompanying text, in addition to the transmission of direct-to-home analog and digital television broadcasts, satellite networks are also used to support a range of interactive services. As we saw in the figure, the satellite network is used to provide a high bit rate channel from the service provider to the set-top box of each subscriber. As we have just seen, a single satellite channel can be used to broadcast data at rates of up to 6/8 Mbps if a single TV channel is used or up to 27/37 Mbps with a full transponder channel. Hence applications that involve the transmission of the same data to all of the subscribers concurrently can readily be supported. In some instances only local interaction with the STB is required whilst in others a facility is needed to enable the subscriber to interact with the remote information source. Typically, this is a remote server computer attached to the Internet. Hence the most popular way of providing an interaction channel is through either a PSTN (with modems) or an ISDN.

There are many applications that can exploit this type of facility and, for description purposes, they can be divided into a number of categories determined by the level of interactivity involved:

- **local interaction**: as we saw in Figure 11.18(a), each MPEG-2 program multiplexer, in addition to the audio and video associated with the (TV) program, also supports an optional data channel. This is used for a variety of purposes. For example, to transmit information about the players in a football game that is being broadcast. Typically, as this is received, it is stored in the STB and the subscriber can then locally interact with the STB to have selected information displayed on the TV screen while the game is being played.

  In other applications of this type, the data is not associated with a particular TV program but is from, say, a server that is located at the basestation of the service provider and has a direct connection to the

Internet. Examples of this are electronic newspapers and magazines. Typically, therefore, the data that is broadcast occupies a separate TV channel. This again is stored in the STB and subsequently can be accessed interactively by the subscriber and displayed on the TV screen. Normally, such channels are controlled by **conditional access**, that is, the broadcast bitstream has been scrambled (randomized) and requires a key to be unscrambled. Another example in this category is **pay-per-view**. The local interactions in these cases involve the subscriber inserting a *smart card* into a slot in the STB;

- **anonymous response to broadcasts**: with this category, a low bit rate interaction channel is involved which, typically, is a PSTN. Examples are when a subscriber votes in a talent contest that is being broadcast or responds to an opinion poll. The subscriber simply calls one of a given set of telephone numbers and the call is logged;
- **purchase requests**: typically, this is in response to a product or service that is being offered via a TV broadcast. For such services the subscriber must interact with a remote location in order to enter credit card details and address information, for example. Normally, this is through a PSTN or ISDN and a call center.

11.4 Terrestrial television networks

Prior to the introduction of small-dish direct-to-home satellite broadcasts in the early 1990s, most homes received broadcast television by means of either cable TV networks or terrestrial networks; that is, both the transmitting and receiving antennas are at ground – or nearly ground – level. So although a significant percentage of homes now receive broadcast TV by means of cable and satellite networks, terrestrial networks are still widely used.

11.4.1 Broadcast television principles

Most terrestrial TV transmissions are in the very high frequency (VHF) and the ultra-high frequency (UHF) bands of the electromagnetic frequency spectrum from 47 through to 860 MHz. As we indicated in Section 6.2.5, above 100 MHz electromagnetic waves travel in straight lines. They also suffer more attenuation than lower-frequency radio waves and, since both the transmitter and receiver are at ground level, large buildings, hills, mountains, and so on between the transmitter and a receiver all impair the transmitted signal and hence give rise to poor reception. Thus to cover a wide geographical area requires a significant number of transmitters distributed around the total reception area. A typical arrangement is shown in Figure 11.19(a).

The signal from the program source to the distributed set of transmitters can be sent by microwave links or land lines, typically optical fiber. In some countries, each transmitter operates in a different frequency band from its

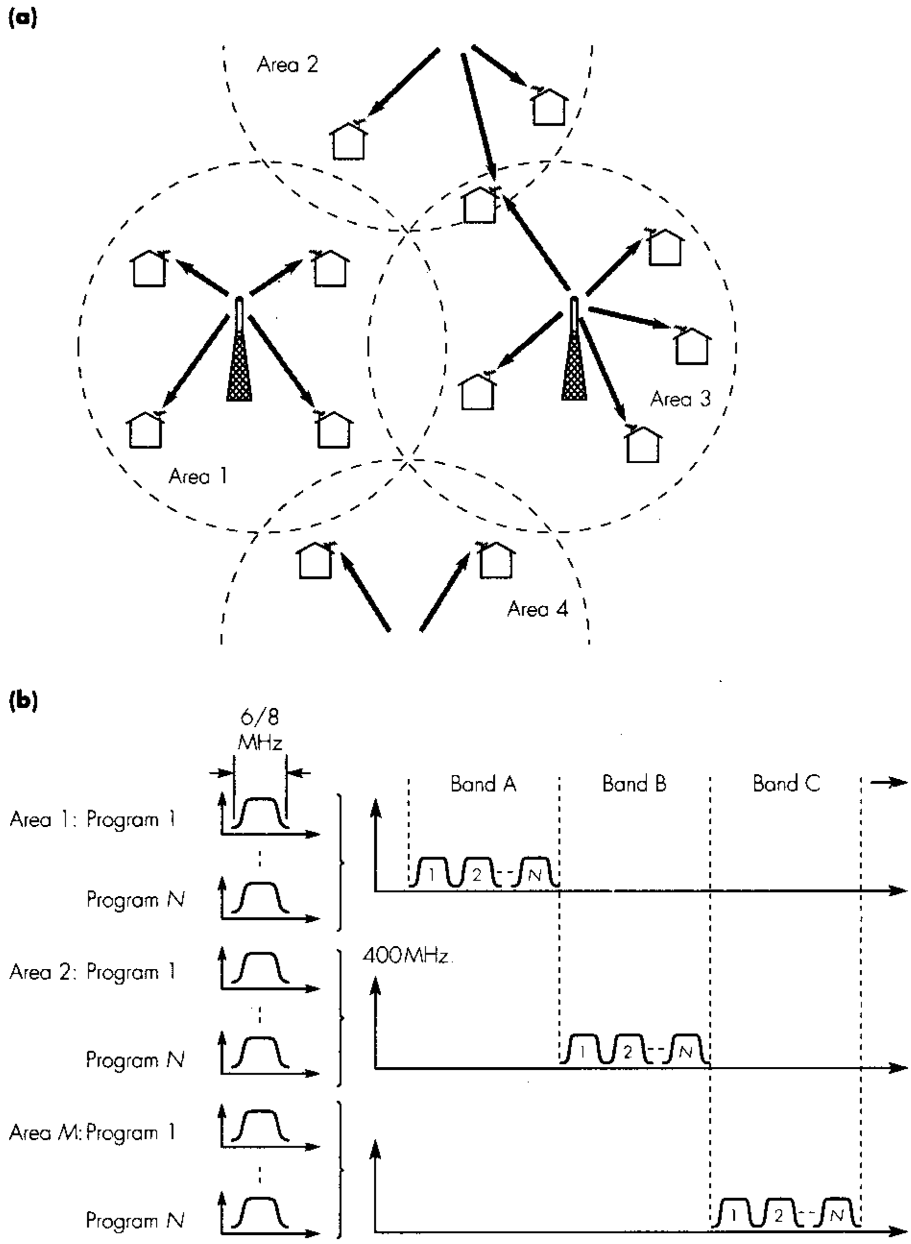

Figure 11.19 Terrestrial television principles: (a) broadcast network schematic; (b) example frequency usage.

neighbors and the network is then known as a **multiple-frequency network** (**MFN**). In other countries, all transmitters operate using the same frequency band and hence the network is called a **single-frequency network** (**SFN**). An example usage of the frequency band in an MFN is shown in Figure 11.19(b) and, as we saw earlier in Figure 6.6, with an MFN it is possible to reuse the frequency bands using a cell structure.

11.4.2 Digital television

The MPEG-2 source coding and channel interface blocks used for digital TV broadcasts over terrestrial networks are similar to those used with satellite networks which we showed in Figure 11.18(a). The main difference is the type of modulation scheme that is used. In order to obtain a wide field of coverage, terrestrial broadcasts are omnidirectional and, as a result, the receiving antenna may receive multiple copies of the same signal from a variety of paths. Although most power is in the direct path (line-of-sight) signal, microwaves are reflected from buildings. Also, various atmospheric conditions can cause the broadcast signal to be refracted back to earth. These waves take a slightly longer time to reach the receiving antenna than the direct wave and lead to an effect known as **multipath**.

All electromagnetic waves travel through free space at the speed of light, about ($3 \times 10^8\ \mathrm{ms}^{-1}$). The wavelength of a signal is defined as the distance travelled by the wave in the time duration of a single cycle of the signal. That is,

Wavelength, $\lambda = c/f$ meters

where f is the frequency of the signal. Hence a signal of 500 MHz has a wavelength of:

$$\lambda = 3 \times 10^8/500 \times 10^6 = 0.6 \text{ meters}$$

Thus, the reflected and refracted waves received by the antenna may have significant phase differences from the direct wave. This is known as **multipath dispersion** or **delay spread** and causes the signals relating to a previous bit/symbol to interfere with the signals relating to the next bit/symbol. This is known as **intersymbol interference** (**ISI**) and, the higher the transmitted bit rate – and hence the shorter each bit cell period – the larger the level of ISI. It is for this reason that a modulation technique called **coded orthogonal frequency division multiplexing** (**COFDM**) is used for the transmission of the high bit rates associated with digital television.

COFDM principles

The main components of a COFDM modulator are shown in Figure 11.20(a) and the principle of operation of the scheme is shown in Figure 11.20(b).

Using COFDM, instead of just a single carrier signal, multiple orthogonal (equally-spaced) carriers are used, each of which is independently modulated

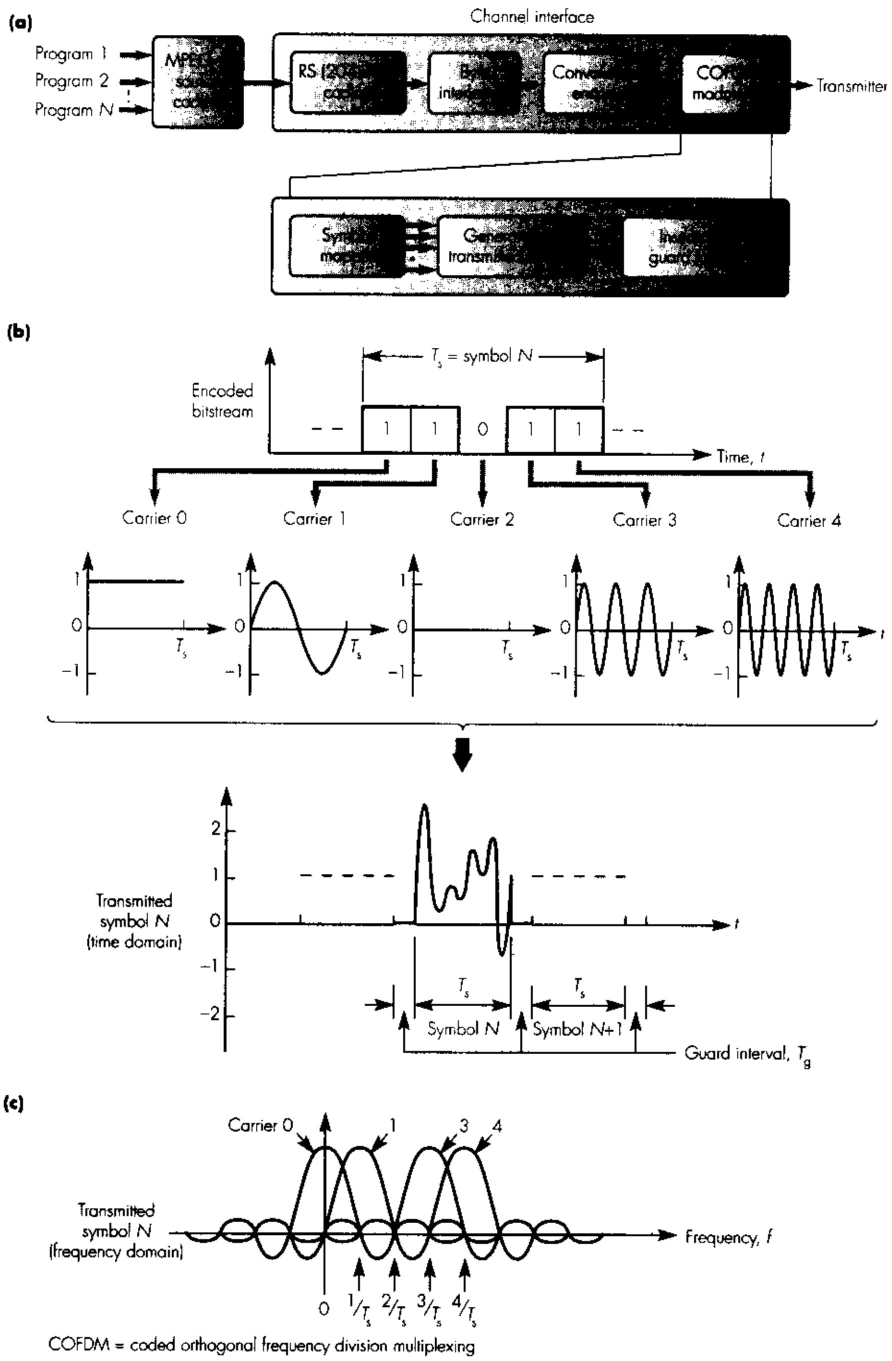

Figure 11.20 COFDM principles: (a) channel interface components; (b) symbol generation (time domain); (c) frequency domain symbol.

by one or more bits from the encoded bitstream to be transmitted. In the simple example shown in Figure 11.20(b), five carriers are used each of which is modulated by a single bit from the transmitted bitstream using on–off keying. The first carrier is a DC level, while the remaining four carriers are all sinusoidal signals of frequencies f_s, $2f_s$, $3f_s$ and $4f_s$ respectively. The period T_s is called the symbol period and, since on–off keying is being used, is equal to the time to transmit 5 bits from the encoded bitstream being transmitted. Alternatively, if N carriers were used each modulated using, say, QPSK, then T_s would be the time to transmit $2 \times N$ bits from the encoded bitstream. During each symbol period, the five modulated signals are added together to form the signal/symbol that is transmitted.

At the receiver, before starting to process each received symbol, the receiver waits a short time interval known as the **guard interval** to ensure that all delayed versions of the direct-path signals that make up the symbol have been received. In this way, instead of the delayed signals interfering with the direct-path symbol, since the same signals make up the symbol, they simply increase the power in the received symbol before it is processed. The processing involves determining which of the five carrier signals are present in the symbol and, based on this (and knowledge of the modulation method that has been used) the receiver can determine the original bitstream that was transmitted. As we show in Figure 11.20(c), by using a set of carriers that are orthogonal, there is a fixed spacing of $1/T_s$ between adjacent carriers which simplifies the processing that is required to determine which of the carriers are present in each received symbol.

In practice, the generation of each symbol is carried out digitally using a mathematical technique called the **inverse discrete Fourier transform (IDFT)**. The digital symbol output by the IDFT is converted into an analog symbol using a digital-to-analog converter prior to transmission. Similarly, the received symbol is first converted into a digital form – using an analog-to-digital converter (ADC) – before the symbol is processed using the DFT to determine the carriers that are present.

Although adding a guard interval avoids intersymbol interference, it also influences the maximum encoded bit rate that can be supported; the longer the guard interval, T_g, the lower the maximum bit rate. Typically, therefore, T_g is limited to a maximum value of $T_s/4$. Hence, in a terrestrial broadcast application, the number of carriers used can be several thousand. For example, with multiple-frequency networks a typical T_g of 50 µs is required to avoid intersymbol interference and hence a typical T_s would be 200 µs. This gives a carrier spacing – and hence fundamental frequency F_s (= $1/Ts$) – of 5 kHz. Hence in a broadcast channel that has a usable bandwidth of, say, 7 MHz, the carriers would be spaced at:

$$0, 5, 10, 15 \ldots 6995, 7000\,\text{kHz}$$

intervals. This means 1401 carriers would be used. In the case of single-frequency networks, however, a typical guard interval of 200 µs is required. This means a typical T_s of 800 µs, an f_s of 1.25 kHz and hence 5601 carriers.

Example 11.2

[redacted]

As we saw earlier, both the IDFT and DFT are performed digitally and hence the processing is carried out in powers of 2. The nearest power of 2 value to 1705 is 2048 (2K) and to 6817 is 8192 (8K) which is the origin of the terms used.

Receiver synchronization

Clearly, it is necessary for the receiver to detect reliably the start of each new symbol that is received; that is, for the receiver to obtain and maintain symbol synchronization. In order to do this, with both network types, the stream of symbols is divided into 68-symbol blocks. Each block of symbols is called a *frame* and the common frame format used is shown in Figure 11.21.

As we can see, along each row are the individual frequency carriers in each symbol. Hence in a 2K system this contains 1705 carriers and in an 8K system 6817 carriers. Although in the simple example we saw in Figure 11.20(b) on–off keying was used, in practical systems each carrier is separately modulated using either QPSK (4-QAM), 16-QAM, or 64-QAM. Since each carrier can transmit 2, 4, or 6 bits per signal transition/element, ignoring the loss of transmission capacity caused by the guard interval, each carrier in an MFN can transmit $4.464 \times 2/4/6$ kbps of data and, in an SFN, $1.116 \times 2/4/6$ kbps of data.

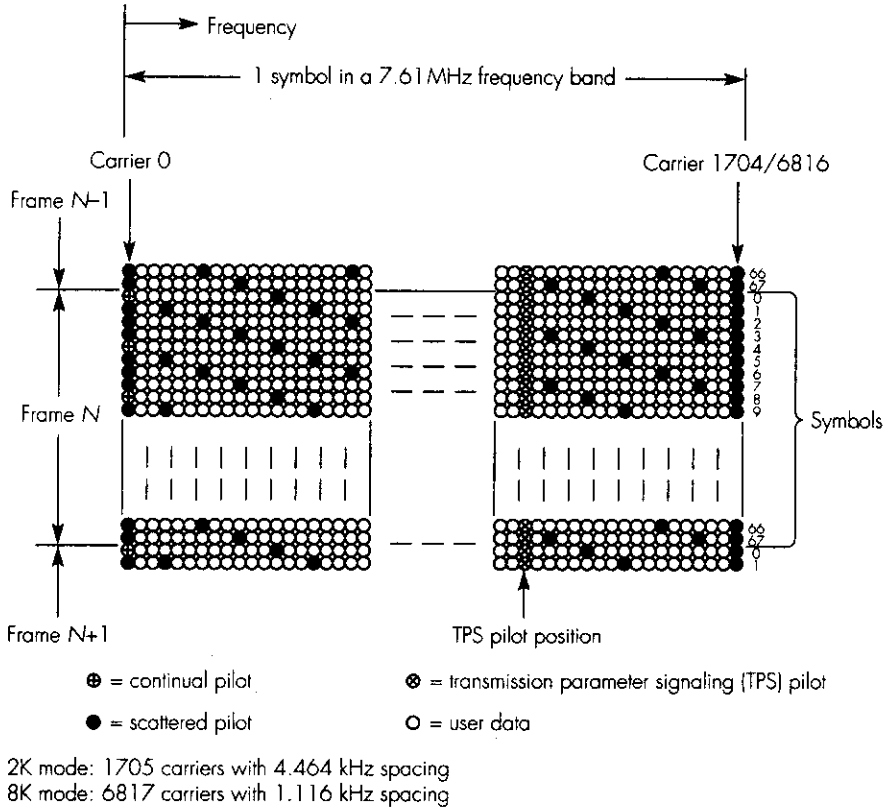

Figure 11.21 DVB-T 2K/8K frame format.

As we show in the figure, the bitstream relating to specific carrier positions in each of the 68 symbols of a frame carries one of two defined bit sequences called **pilots**: one the *continual pilot* and the other the *scattered pilot*. The receiver achieves and maintains symbol synchronization by searching for these two pilot sequences in the known carrier positions. In addition, once in synchronization, the bitstream of the eighteenth carrier from the end of each symbol is used to form a subframe that contains operational parameters to enable the receiver to interpret correctly the received symbol stream. These carriers are called *TPS pilots* and the contents of each subframe include the length of the guard interval, the type of modulation of the other carriers that are being used, and the rate of the convolutional coder.

The bit rate that is available with each 6/8 MHz broadcast channel is determined by the modulation method that is used for each carrier. With the 2K specification and 16-QAM for example, after removing the pilots and allowing for the loss of transmission capacity caused by the guard interval, a typical channel bit rate is 24 Mbps. As we saw earlier in Example 11.1, the overheads associated with the other blocks in the channel interface – the RS (204, 188) coder and a rate 3/4 convolutional coder – reduce this figure to about 16 Mbps (24 × 184/272). This can be used to carry, say, four 4 Mbps digital TV programs in each 8 MHz broadcast channel.

11.4.3 Interactive services

The interactive services available with terrestrial broadcast networks are similar to those we identified in relation to satellite networks. In general, however, the amount of transmission bandwidth/number of channels available with terrestrial networks is much less than with satellite networks. Hence the range of services available is only a subset of those listed in Section 11.3.3.

11.5 High-speed PSTN access technologies

As we saw in Figure 1.1(c), the access network of a PSTN, in addition to supporting the plain old telephone service (POTS) for which it was designed, now supports a number of additional services. For example, a range of low bit rate data applications such as fax are supported by means of low bit rate – less than 56 kbps – modems. Also, as we saw in Section 7.2.3, twisted-pair lines in the PSTN access network are used as the access lines for an ISDN. Bit rates of between 144 kbps (basic rate) and 1.544/2.048 Mbps (primary rate) over several miles/kilometers are obtained using baseband transmission. The access line is then known as a digital subscriber line (DSL). In the case of a basic rate line, this is called an **ISDN DSL (IDSL)** and, in the case of a primary rate line, a **high-speed DSL (HDSL)**. An IDSL uses a single pair and an HDSL two pairs. In addition there is a simpler version of HDSL which operates over a single pair. This is known as **single-pair DSL (SDSL)** and bit rates of up to 1.544/2.048 Mbps are supported depending on line length.

Both the basic rate and the primary rate lines of an ISDN are symmetric; that is, they operate with an equal bit rate in both directions. However, as we saw earlier in Section 11.2.1, with most interactive applications the information flow is asymmetric and involves a low bit rate channel from the subscriber for interaction purposes and a high bit rate channel in the downstream direction for the return of the requested information. Asymmetric ratios of from 10:1 to in excess of 100:1 are common. In order to exploit this, many cable operating companies have introduced cable modems that support various applications of this type.

In addition, many **telecommunication operating companies (telcos)** have introduced additional types of DSL technologies to meet these same requirements over the twisted-pair lines used in most PSTN access networks. Unlike the various DSL technologies associated with an ISDN, these have been designed to enable the signals associated with the existing telephony services to coexist with those associated with the newer high-speed interactive services on the same twisted-pair line. Two types are used: the first known as **asymmetric DSL** (**ADSL**) and the second **very-high-speed DSL** (**VDSL**). In the case of ADSL, the high-speed asymmetric channel is designed to coexist with the existing analog telephony service. In the case of VDSL, the high-speed channel, in addition to operating at a higher bit rate than that of an ADSL, can operate in either an asymmetric or symmetric mode and is designed to coexist with either analog telephony or basic-rate ISDN services. In this section we discuss both these technologies.

11.5.1 ADSL

The standard relating to ADSL was produced by the ANSI in 1995 and is defined in **T1.413**. It was defined originally to meet the requirements of broadcast-quality video-on-demand (VOD). Hence the standard allows for a bit rate of up to 8 Mbps in the downstream direction – that is, from the local exchange/end office (LE/EO) to the customer premises – and up to 1 Mbps in the upstream direction. In practice, however, high-speed access to the Internet proved to be more popular than VOD. As a result, since the bit rate and QOS requirements associated with Internet applications are less than those of VOD, a variant of the original ADSL standard known as **ADSL-Lite** (or sometimes **G-Lite**) has been defined. This has been developed within the ITU and is defined in standard **G.992.2**. It provides a downstream bit rate of up to 1.5 Mbps and an upstream bit rate of up to 384 kbps. As with ADSL, the actual bit rates achievable depend on the length and quality of the line. Nevertheless, the lower bit rates associated with ADSL-Lite means that it can be used over longer distances and with poorer-quality lines than ADSL. In addition, as we shall see, ADSL-Lite can be used with a passive network termination at the customer premises rather than the active termination that is required with ADSL.

Access network architectures

As we saw in Section 7.2.1, in the earliest PSTN access networks the transmission lines used were made up entirely of interconnected sections of unshielded cable containing multiple twisted-pair wires. Over a period of time, however, sections of these cables have been replaced with optical fiber cable. The architecture of a typical PSTN access network is as shown in Figure 11.22.

As we can see, the amount of twisted-pair cable used varies from zero with **fiber-to-the-home** (**FTTH**) or **fiber-to-the-building** (**FTTB**), short (less than 300 yards/meters) drop cables with **fiber-to-the-kerb/curb** (**FTTK/C**), one or two cable sections with **fiber-to-the-cabinet** (**FTTCab**), and all twisted-pair for direct-to-building cable runs of up to 2.5 miles (4km). Hence, since the objective of ADSL is to provide high-speed interactive services over the twisted-pair portion of the access line, ADSL is designed on the assumption that the maximum length of twisted-pair cable is less than 2.5 miles (4 km).

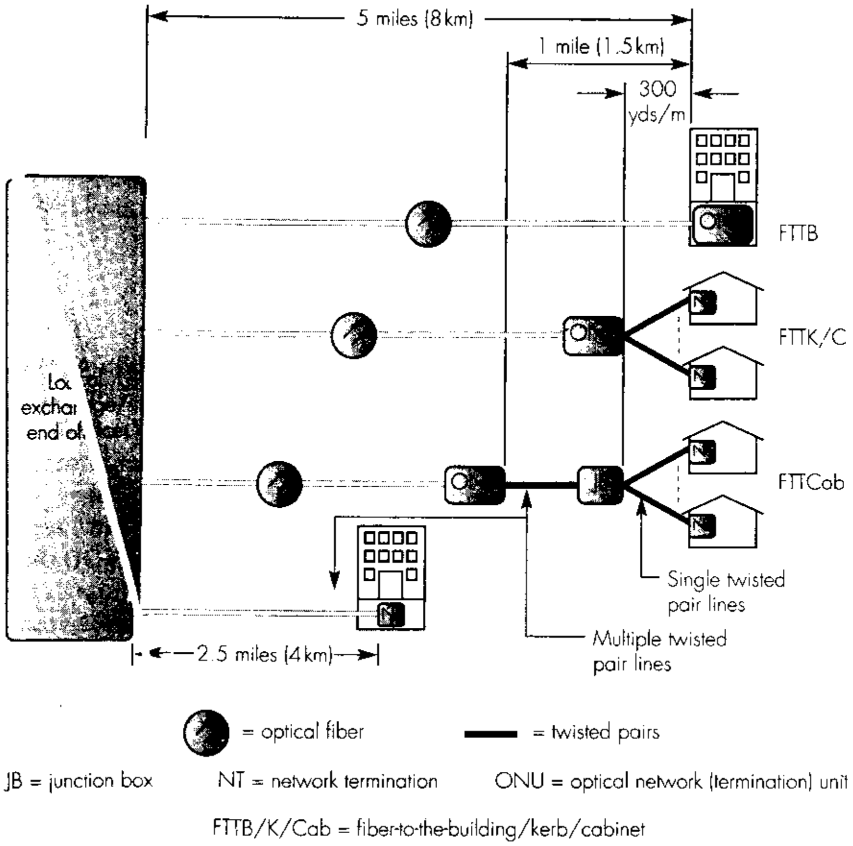

Figure 11.22 Typical modern access network architecture.

Connection alternatives

As we saw in Section 2.5.1, the lower frequency band up to 4kHz of the bandwidth available with a (single) twisted-pair line is used for analog telephony (POTS). So in order for the two ADSL/ADSL-Lite signals to coexist with the POTS signal on the same line, modulated transmission must be used to take the signals away from the lower frequency band. In the case of ADSL, the two signals are transmitted in the frequency band from 25 kHz through to 1.1 MHz and, for ADSL-Lite, the upper frequency is limited to 500 kHz. The components that are used to deliver both services over the same access line are shown in Figure 11.23. The arrangement shown in part (a) relates to a typical ADSL installation and that in part (b) an ADSL-Lite installation. In both cases it is assumed that the access line is all twisted-pair and this terminates in the LE/EO. With the various fiber access alternatives, the same line termination arrangement is located in the **optical network (termination) unit (ONU)**.