Test in TTM4105 autumn 2006 – English – proposed answer

Note that the answers given below are more comprehensive than is actually required in the text. They are generally taken from the textbook.

The test consists of four problems with equal weight.

Problem 1 Synchronisation of satellite system

In the satellite system in the figure, 6 earth stations accesses the satellite by using TDMA (time division multiple access). The frame is constructed as shown in B) and consists of 6 bursts, that is, n = 6. The format of each bust is shown in C). The overall bit rate is 100 Mbps (mega bits per second) and the distance between adjacent bursts is 10 bits.

The satellite receives bursts from each earth station and retransmits the bursts back to the earth stations in the ordered form shown in figure B). All earth station can then communicate with one another.

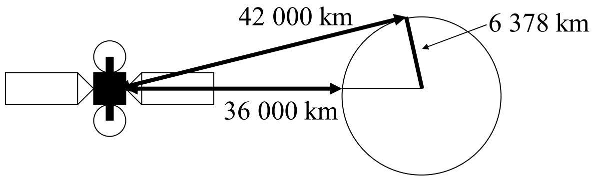

A) Geometry of a geostationary satellite

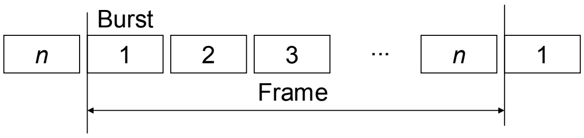

B) TDMA format

Carrier Clock	UW	Payload

C) Burst format

a) *If the speed of light is 300 000 km/s, how long time does it take for a signal to reach the satellite when it is sent from an earth station*
   - *located on the Equator just below the satellite;*
   - *located at a point at the edge of the coverage area?*

   All six earth stations calculate the time when they must send the burst to the satellite from the instant they receive the signal from the satellite.

   If the earth station on the Equator sends its burst at time t_0, when must then the earth station at the edge of coverage send its burst in order to avoid that the two bursts overlap when they arrive at the satellite but are received in adjacent timeslots?

   It is simplest to explain this by use of a sketch as in the lecture notes.

   Earth station at Equator just below the satellite: 36 000 km/300 000 km/s = 120 ms

   Earth station at edge of coverage: 42 000 km/300 000 km/s = 140 ms

An earth station at the edge of coverage must send the burst at the instant $t_0 - 20$ ms (the differential delay) earlier than the station at the Equator just below the satellite in order that the two bursts shall arrive at the same time in the satellite. In order to arrive in adjacent bursts, the burst must be shifted by one burst length (including the guard space); that is, at time $t_1 = t_0 - 20$ ms + one burst length. See the figure.

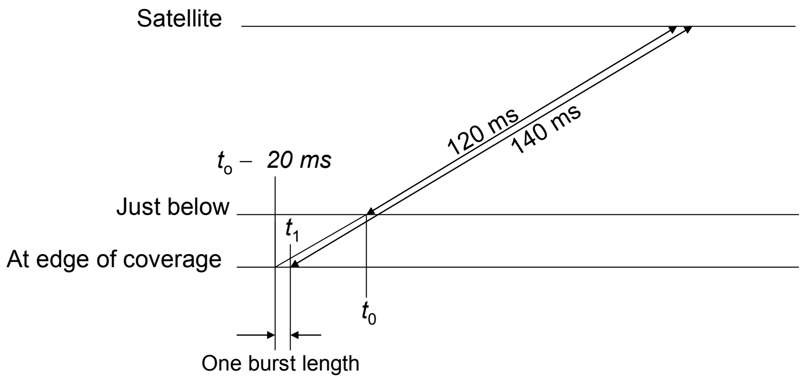

b) *Explain how the earth station synchronises the sending timer after synchronisation has taken place. (Hint: which properties must the unique word (UW) have in order to synchronise the frame format?)*

If the earth station has found the instant when the burst shall be sent, the position of the burst relative to the first burst in the frame (reference burst) can be computed accurately. The earth station keeps the burst in this position by using a slow phase-locked loop (bandwidth corresponding to about twice the roundtrip delay for reasons of stability) where the reference signal input to the loop is the position of the reference burst. This burst position is then compared with the actual position of the received burst of the station and the phase-locked loop will compensate from any deviations.

The unique words of the reference burst and the burst being corrected are used to derive the positions in time for the two bursts. This is done with an accuracy of one bit. The ten guard bits are required in order to allow fluctuations in the oscillators over a time equal to twice the roundtrip delay.

c) *Explain how the earth station is synchronised the first time it shall commence sending in a given timeslot.*

The method used is based on spread spectrum with a large spreading. We may, for instance, use a signal at rate 1 Mbps containing a unique word and a spreading code consisting of 100 chips per information bit as initial synchronisation signal. This gives a chip rate of 100 mega-chips per second (Mcps). The spread spectrum signal then fills the whole bandwidth of the satellite. This gives a coding gain of 20 dB; that is, the spread spectrum signal can be sent at a level 20 dB below the level of ordinary bursts and still be detected with the same quality as an ordinary burst.

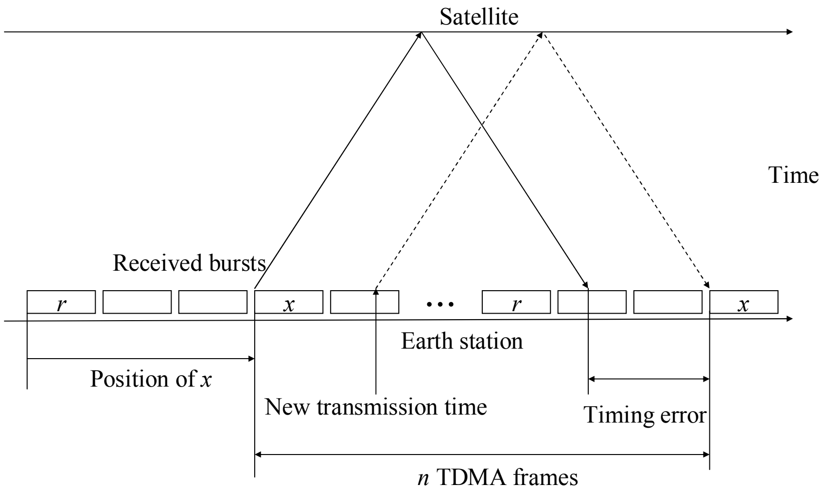

Initial synchronisation of burst transmission

The procedure is shown in the figure. Initially, the earth station monitors the received burst pattern from the satellite and determines the start of the TDMA frame (the reference burst *r*). From this information, the position of the timeslot that is assigned to the earth station (*x* in the figure) is calculated. The earth station sends the 100 Mbps spread spectrum signal in slot *x* and measures the instant when the signal is received back again from the satellite as shown in the figure. This allows the earth station to determine the timing offset from the correct transmission instance of burst *x* as shown.

The procedure may be repeated sending the spread spectrum burst at the derived instant in order to reconfirm the synchronisation. This procedure may be repeated several times.

However, this measurement can only be done with an accuracy of 100 bits of the 100 Mbps signal since the length if 1 bit of the 1 Mbps signal corresponds to 100 bits at 100 Mbps. To resolve this inaccuracy, the earth station may, after initial synchronisation is acquired, send a short burst at nominal power level and bit rate in the middle of the assigned burst and then do an accurate determination of the start of the burst. The station can then commence sending in the allocated burst slot and resume normal synchronisation procedures as described above.

Problem 2 Mobile communication

a) *What is handover and when is this mechanism used*?

Handover is the process of transferring a call in progress from one cell to another without interrupting the call. Handover may take place for two reasons: the communication channel is taken out of service because of equipment failure or maintenance; or the mobile terminal moves out of one cell and into another. In the first case the call is handed over to another resource in the same base station (intracell handover); in the second case the call is handed over to another base station (intercell handover).

b) *How is handover performed in GSM*?

In GSM, the mobile stations measures the quality (field strength and bit error rate) of the signal received in its own cell and in adjacent cells. The identity of the adjacent cells and the carrier frequency on which measurement may take place is signalled to the mobile station in the broadcast control channel. We saw in Section 8.5 that the mobile terminal when locked to a traffic channel neither sends nor receives information in the last timeslot of the 26-multiframe. During this interval, the mobile station tunes to the appropriate frequency of one of the neighbouring cells and measures the field strength of the received signal. The mobile station scans all the adjacent cells in some cyclic order. The measurement results are sent to the BSC together with the error rate and field strength of the traffic channel to which the terminal is locked.

The BSC compares the measurements and decides if a handover shall take place. The procedure is shown in Figure 8.25 for handover between base stations on the same MSC. The procedure is as follows.

When receiving the field strength and bit error rate measurements from the mobile station, the BSC will determine if handover is required and, if so, can possibly be performed. The message also contains the identity of the cell in which the measurement is taken. The BSC also measures the quality on the uplink. Both the measurements of the uplink and the downlink are used in order to establish the handover criterion. From the information received from the mobile station, the BSC can also determine to which BSC (the target BSC or BSC2 in the figure) the call should be handed over.

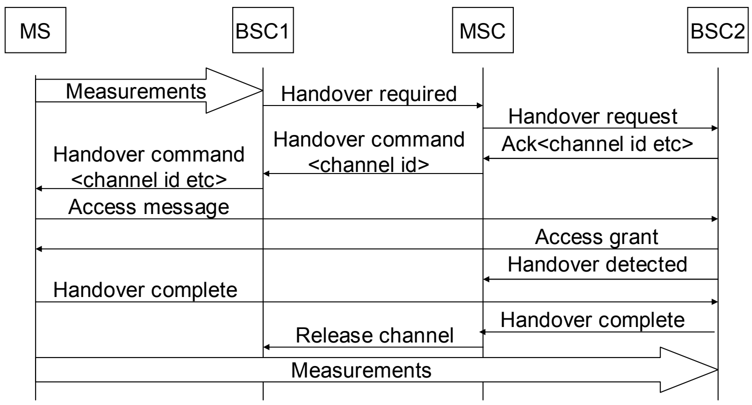

Figure 8.25 Handover in GSM

The BSC1 then requests the MSC to initiate the allocation of a radio channel at the target BSC2. If a channel is available, the MSC will send the handover command to the mobile station. This message contains the channel number (frequency and timeslot) of the new channel and other parameters such as the power level to be used in the new cell, location of the synchronisation channel, and a handover reference to be used in the response message. The mobile terminal may then initiates the timing advance procedure by sending a random access message. The access grant message contains the timing advance information the mobile terminal needs in order to continue communication on the new channel. The event is reported to the MSC such that through-connection of the new path can commence. Finally, the mobile terminal sends the handover complete message to the target BSC2 and commences normal operation, including measuring the quality of the new channel. The handover complete event is signalled to the MSC that orders the old BSC1 to release the radio channel.

c) *Explain how soft handover is done in UMTS (3G).*

The radio network controllers (RNC) under the same SGSN are interconnected. In the upper picture, the mobile station is served by RNC1. In the border zone between the two cells the signals from the mobile terminal is demodulate by both RNCs. RNC2 sets up a connection to RNC1 and transfers the message received from the mobile terminal together with field strength information. RNC1 then decides if and when a handover should take place.

When the call is handed over to RNC2, the call is still routed via RNC1 so that the management of the call is still with RNC1.

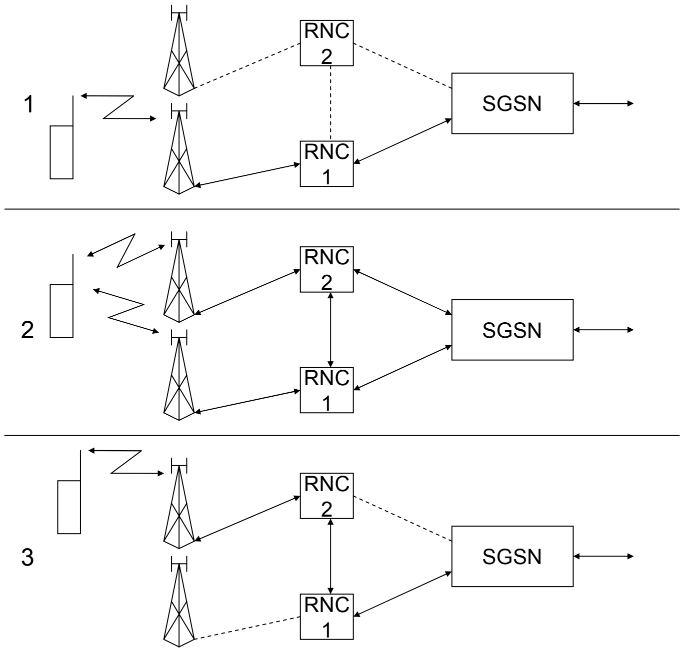

Figure 8.24 Soft handover in UMTS

The control of the call may be handed over to RNC2 if RNC1 finds this more appropriate, particularly if several handovers have taken place in succession and the distance between the RNCs is long. Note that the RNCs of UMTS have a more central role in controlling the network than the base station controllers of GSM.

d) *Why is soft handover possible in UMTS but not in GSM*?

UMTS uses DS-CDMA in a common frequency band so that the signal from a mobile terminal may be picked up by both Node Bs. This does not cause significant increase in interference.

The base stations in GSM operate in different frequency bands using TDMA in each band and can therefore not tune to the signals form mobile terminals in adjacent areas.

Problem 3 Multiplexing in SDH

a) *Show how signals with different data rates are multiplexed in SDH.*

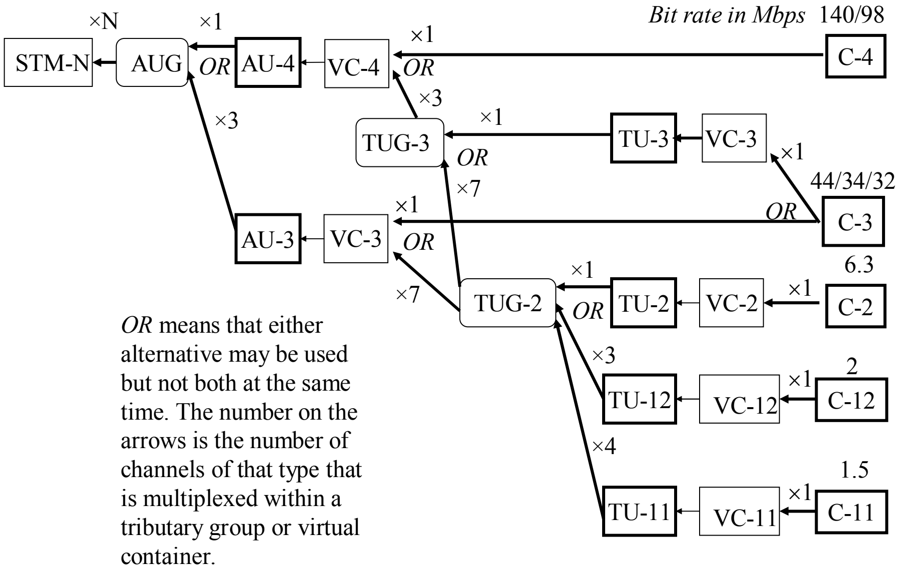

Figure 4.14 Multiplexing structure

The functions supported by the different units are described in Table 4.4. These are:

- The *container* (C) is the payload itself.
- The *virtual container* (VC) encapsulates the payload. The encapsulation consists of a header (the path overhead).
- The *tributary unit* (TU) contains pointers by which the position of the payload in the multiplex structure is uniquely identified. This entails insertion of stuffing bits in order to bring the input signal to a specified bit rate. The tributary unit also allows the tributary unit group to multiplex plesiochronous channels.
- The *administrative unit* (AU) is in charge of similar functions as the tributary unit (pointers and bit stuffing). The administrative unit takes care of the uppermost level of the multiplex hierarchy.
- The administrative unit group (AUG) is a multiplexer.
- The *synchronous transport module* (STM-*n*) multiplexes the signal further into the optical system.

The *tributary unit group* (TUG) multiplexes tributary units before they are presented to the next level of the hierarchy. This includes the addition of headers.

b) *Sketch the structure of Synchronous Transport Mosule (STM-1) in SDH. (Hint: remember the headers "section overhead", "pointer" and "line overhead.")*

The AU-4 of STM-1 contains the virtual container VC-4.

All virtual containers consist of two parts:

- Operation, administration and maintenance (OA&M) information is contained in a field called *path overhead*. This information follows the container until it is removed from the SDH system.
- The payload can either be a single PDH multiplex signal (container C) or a number of tributary unit groups (TUG) multiplexed in accordance with

principles defined for the VC. The TUGs consist of one or more tributary units (TUs) multiplexed in accordance with rules that apply for that TUG. The TUs contain pointers and other overhead information plus a virtual container as payload. The TU pointer identifies the beginning and the end of the VC frames. This together with Figure 4.14 illustrates the iterative structure of SDH. This structure then contains a set of linked pointers that finally identifies a single container uniquely.

The STM-1 envelope is shown in Figure 4.15. The structure of the envelope is drawn in the form of a matrix consisting of 270 columns of octets and nine rows. The envelope contains $9 \times 270 = 2430$ octets.

The transmission of bits is as follows. The first bit to be sent is the leftmost bit in row number 1. Then follows bit number two in this row and so on until the last bit of row number 1 has been sent. The next bit is then the leftmost bit in row number 2 and when the final bit of row number 2 has been sent, the next bit is the leftmost bit of row number 3 and so on. The final bit is the rightmost bit in row number 9.

The only magic about this arrangement is that it is easy to read, easy to understand and easy to specify.

The STM-1 envelope contains the section overhead and the link overhead and the AU-4 pointer field. The pointer field is in row number 4. If necessary, the AU adds stuffing bits in order to make the payload plus the path overhead consist of $9 \times 261 = 2349$ octets.

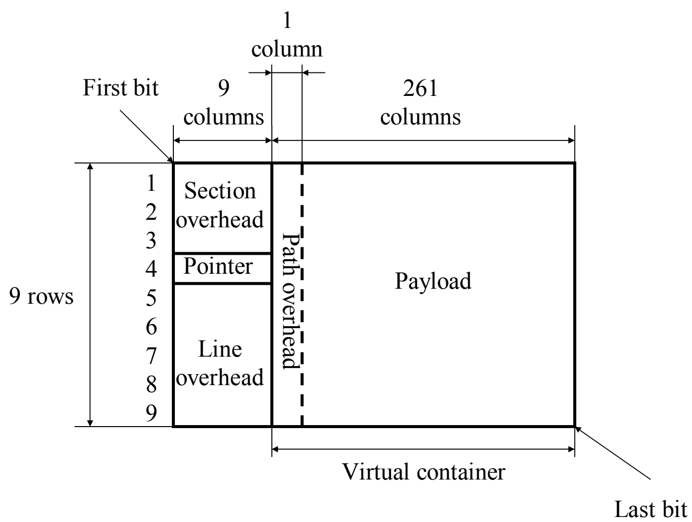

Figure 4.15 Structure of STM-1

c) *Explain how a virtual container (VC-4) is placed inside the STM-1 frames. What is floating payload? How can we accommodate a plesiochronous signal in such a payload without loosing information? (Hint: Start with the pointer mechanism and explain how it is used.)*

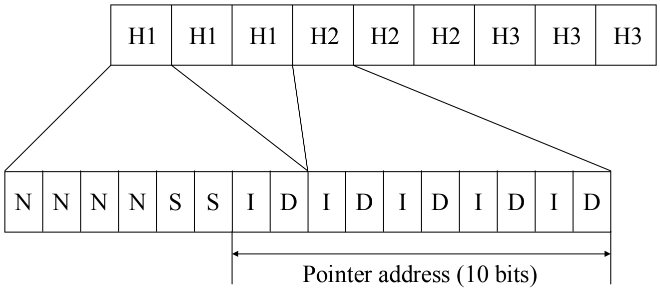

N = new data flag (0110 same address, 1001 new address)
S = type of AU (AU-4 or AU-3)
I = increment bit (normal: no adjustment; inverted: increment address by one)
D = decrement bit (normal: no adjustment; inverted: decrement address by one)

Figure 4.16 Pointer field

The STM-1 envelope contains the section overhead (3 × 9 octets = 27 octets), the link overhead (5 × 9 octets = 45 octets), the pointer field and the actual payload (C-4 or three TUG-3 plus the path overhead). The AU-4 pointer consists of ten bits in row 4. The value of the pointer multiplied by 3 gives the start position of the VC within the envelope. This gives a maximum value of the number of pointer addresses in the VC: $3 \times 2^{10} - 1 = 3071$. This is larger than the actual maximum length of the VC frame, which is 9 × 261 = 2349 octets.

The pointer field is shown in Figure 4.16. The field consists of nine octets divided into three groups H1, H2 and H3 where each group consists of identical formats. One H1 plus one H2 contain together the pointer address as shown in the figure, while H3 is used for rate adjustment. If the envelope contains one AU-4, then the first H1 and H2 octets are used as pointer while all three H3 octets are used for rate adjustment. If the envelope contains tree AU-3, all three sets of H1, H2 and H3 are used in order to assign individual pointers and adjustment fields to each AU-3.

The combination of the H1 and H2 octets consists of three parts as shown:

- The new data flag consisting of four bits that are normally encoded as 1001 indicating that the address in this envelope is the same as in the previous one. The value of the new data flag is the same even if rate adjustment takes place as described below. If a new pointer address is set, the new data flag is encoded as 0110 in the envelope in which the pointer address is changed. Such operations take place when instantiating the system, when the multiplex configuration is changed and after a disruption of communications.
- The *SS* bits indicate whether the envelope contains one AU-4 or one, two or three AU-3. This thus represents the multiplexing function of STM-1.
- The actual address of the pointer position is contained in the *I*-bits and *D*-bits. If the pointer value has to be increased because of rate adjustment, the *I*-bits are inverted in the envelope where the adjustment takes place. The pointer value is then increased by one in the next envelope but with the *I*-bits in the non-inverted form. The operation of the *D*-bit is similar and is used in order to decrease the pointer value by 1. How rate adjustment takes place is explained below.

Observe that a single bit error anywhere in the pointer field can be corrected since majority decision is used on the new data flag and that the value of the *SS* bits and the address bits are changing slowly in normal operation.

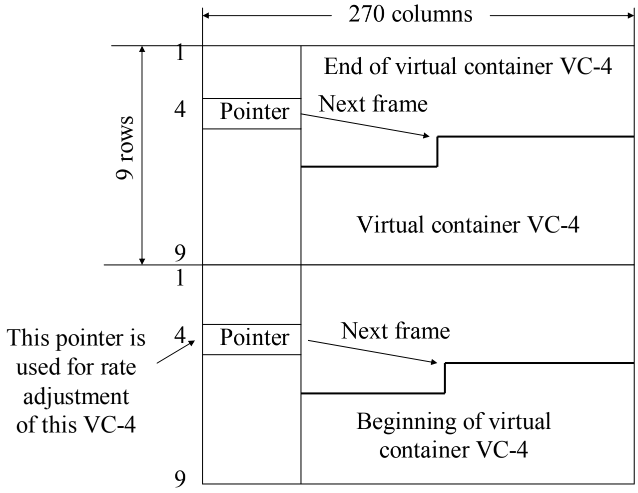

Figure 4.17 Floating VC-4

The start of the virtual container cannot be earlier than the first bit following the pointer. This implies that one VC-4 will be contained in two envelopes as shown in Figure 4.17. This allows the VC-4 to "float" within the envelope structure. The pointer that is within the VC-4 is used for rate adaptation of that VC-4 (see below).

The floating mechanism also allows bit rate adjustment. The adjustment must take place in groups of three octets because of the way in which the pointers are constructed as explained above.

The principle for bit adjustment is as follows. The input signal is buffered in an elastic store. Since an adjustment takes place in steps of three octets (24 bits) the size of the elastic store must be at least 2×24 = 48 bits in order to support adjustment in both directions.

If the input rate of the VC-4 signal is slower than the SDH rate, the number of bits in the buffer will be slowly reduced since the buffer is not filled up as fast as it is emptied. If the buffer has become three octets (24 bits) shorter, the envelope that is sent by the STM-1 envelop is made 24 bits *longer* than normal by letting the first three octets after the header be empty. The envelope contains a normal VC-4 frame of 2349 octets in addition to the 24 dummy bits. This operation is indicated in the pointer field by inverting the *I* bits as explained above. Now 24 bits more than normal are sent during the interval between two pointer values. In other words, the time interval between the pointers is made 24 bit durations longer. During this time the buffer will be filled up again to its nominal value of 24 bits. In the next envelope, the *I* bits are reverted to normal and the pointer value is incremented by one.

If the input rate is larger than the SDH rate, the buffer is filling up faster than it is emptied. When the buffer has become three octets longer than nominal, a VC-4 frame is inserted in an envelope that is three octets *shorter* than nominal but where three octets of the frame are inserted in the three H3 octets of the header. The number of information octets in the frame is still 2349 octets but they are sent in a time interval that is 24 bit durations shorter than normal. The *D* bits are inverted in the pointer of the same header that contains the adjustment bits indicating that rate adjustment is taking place. This operation reduces the buffer length by 24 bits. The pointer value is now decremented by one.

Problem 4 Miscellaneous

a) *Explain the three basic principles for statistic multiplexing: fixed frame structure, use of length indicators, and use of flags.*

Invariant frame structure

The frame structure may be fixed consisting of a synchronisation header and an information field of fixed length. If no information is sent, empty frames are sent. These frames are identical to information frames except that a parameter in the header is used to indicate that the frame does not contain information.

This principle is used in ATM where each frame (or cell) consists of a header containing 5 octets and a payload field of 48 octets as shown in Figure 4.19. The payload of individual cells may contain different information channels. The layer above ATM keeps track of which cells contain information belonging to the same and different communication sessions. In some applications, forward signalling is used to set up a connection between terminals in the same way as in the telephone network or ISDN. ATM provides a unique numbering of cells in order to support this routing function.

ATM supports information streams as shown in the upper part of the figure where cells (or frames) originating from different sources can be statistically multiplexed. The space between cells containing information is, if necessary, filled with empty cells (E).

The cell header contains the following parameters. The identifier field contains information that is used by the switching equipment to route the call through the network. The payload type is used to indicate whether the cell contains user-to-user information or network management information. CLP is the cell loss priority indicator. If CLP = 0, the cell must, if at all possible, not be discarded if the network is congested. If CLP = 1, the cell may be discarded in case of congestion.

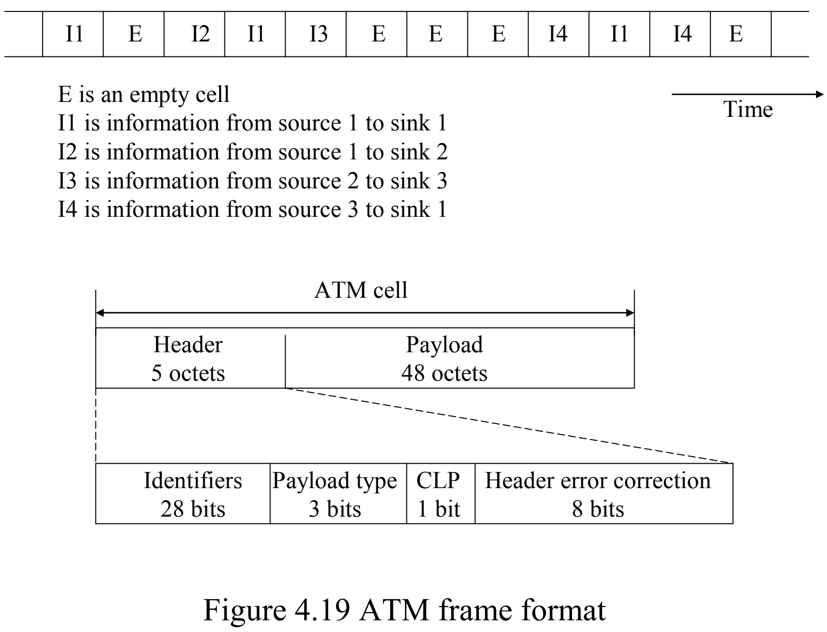

Figure 4.19 ATM frame format

The header error correction field serves two purposes: correcting bit errors appearing in the four first octets of the header in order to avoid misrouting of cells (this operation is called the error correction mode), and supporting synchronisation of cells (this operation is called the synchronisation mode) as explained in Section 3.6.

All bits of the first four octets of the header, except the CLP bit, are set to zero in empty frames. The CLP bit is set to 1.

Delimitation by frame-length indicators

The header is the same for all frames but the length of the frame varies. For this scheme to work, it is necessary to keep track of the header of each multiplexed signal by use of length indicators, more data indicators, end of message indicators, and alignment with an underlying envelope.

This is the multiplexing scheme of the ATM Adaptation Layers (AAL) named AAL2. We shall use AAL2 as a simple example in order to explain the principle (even though AAL2 is no longer used in ATM). The principle is also used in GSM but the procedures are more complex and intertwined with other functions.

In the case of AAL2, the underlying envelope is the ATM cell structure.

AAL2 supports low speed data channels with variable bit rates. More than one AAL2 channel may be multiplexed within one cell. This is done as shown in Figure 4.20.

The individual channels are segments of low-rate information streams (for example, several 28 kbps data channels and 64 kbps voice channels multiplexed in the 155 Mbps ATM stream). The streams are multiplexed within cells as shown using the cell header (H3) as fundamental timing reference. The header of the CPS-PDU (H2) consists of eight bits and contains a 6 bit pointer indicating where the H1 header of the first channel starts in the CPS-PDU. This is the major purpose of the CPS-PDU header. In the example, the pointer value in the first PDU indicates the start of channel 1 and the pointer value in the first PDU indicates the start of channel 1, and pointer value in the second PDU indicates the start of channel 4.

The bits following channel 5 in the example are padding bits containing no information and are used to encode empty bit positions at the end of the frame. These bits are set to binary zero.

The header H1 consists of three octets and contains several fields the two most important of which are the channel identifier field (the first 8 bits of the header) and the length indicator field (6 bits). The length indicator specifies the total number of octets in the payload field of the channel (maximum 63 octets). The channel identifier enumerates the individual channels in the multiplexing structure. The first bit of the identifier field cannot be zero in order to distinguish it from a padding field. In other words, if the first bit following a channel is zero (channel 5 in the figure), the remaining bits of the CPS-PDU are padding bits.

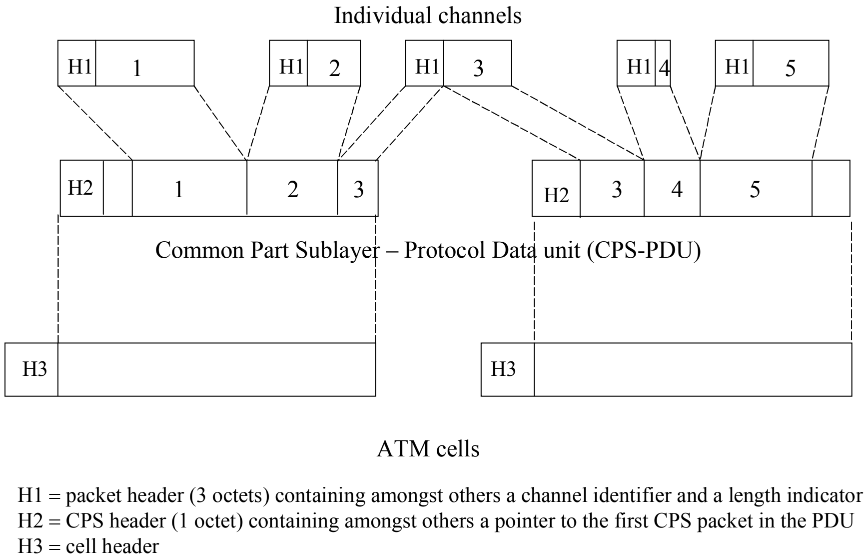

Figure 4.20 AAL2 multiplexing

The multiplexing format of AAL2 is very flexible, allowing the size of individual channels, the order in which the channels appear and the time between fragments of the same channels to vary during a communication session.

Note that the arrangement allows channels to be distributed over two ATM cells (channel 3).

Delimitation by flags

In these systems, a delimiter that cannot be simulated by the information signals is inserted between frames. This is the method used in HDLC-based networks (at the physical layer of some data networks and in the physical layer of Signalling System No 7).

The algorithm is as follows. Frames can be of arbitrary length and contain arbitrary strings of binary information. The delimiter between frames consists of the eight-bit pattern 01111110. This pattern is called a flag. Whenever the receiver detects a 0 followed by six 1s and then a 0, the receiver will interpret this as the end of one frame and the start of the next one. If there is no information to be sent, only flags are sent (corresponding to frames of zero length). In order to avoid that a flag appears in the information stream and is thus misinterpreted as the end of the frame, the sender insert a dummy 0 after every sequence of five contiguous 1s, that is the sequence 011111 becomes 0111110. The algorithm in the receiver is then such that whenever the sequence 0111110 is detected, the final 0 in the sequence is deleted and the remaining bits in the sequence are kept as part of the information stream.

Isochronism relative to octets cannot be maintained in systems employing this technology as explained in Section 3.2.

b) *Explain the basic principles of CSMA-CA (carrier sense multiple access with collision avoidance).*

Collision avoidance CSMA (CSMA/CA) is a control procedure that can be used in radio systems. The algorithm is as follows. If a source is ready to send and the medium is sensed idle during a certain short time interval, the source starts the transmission at the end of this short interval. If the channel is busy, the source waits until the medium becomes idle, draws a random back-off time (also known as *contention window*) and continues to monitor the medium while counting down the back-off time. If the medium becomes busy again, the source repeats this procedure until the message is sent. This may take several cycles. The method used is thus non-persistent CSMA with limited retransmission interval.

A more fair method is to give sources that have already been waiting a smaller back-off window the next time the medium becomes idle. One way to do this is to stop the back-off counter when the medium becomes busy again and restart it when the medium becomes idle. Then the source continues to count down the back-off counter. When the back-off counter expires, the packet is sent immediately (by definition, the medium is always idle when the back-off counter expires).

The load is further reduced by applying exponential back-off time where the collision window is doubled each time a collision is detected. This control mechanism is the same as in CSMA/CD described above.

Since the sources cannot detect collisions directly, collisions must be detected using other means. This may be handled by higher layers of the protocol (e.g., TCP) but then the source will not be aware of the collision since a retransmission by TCP will be regarded as a fresh packet. However, whenever TCP generates a retransmission, the control procedure of TCP will reduce the load on the radio channel.

Alternatively, acknowledgment packets may be sent directly on the medium. This procedure is possible by the WLANs of the IEEE 802.11 specification and is based on the particular inter-frame spacing shown in Figure 5.15.

The time intervals in the IEEE 802.11 WLANs using direct sequence CDMA are: SIFS = 10 μs, slot time = 20 μs, PIFS = 30 μs and DIFS = 50 μs (PIFS plus a slot time). Normal messages are sent after a waiting period equal to DIFS after the medium was sensed idle. Acknowledgement message is sent after a waiting time of SIFS (i.e., during the first *slot time*) and will therefore not collide with any other packets. The acknowledgement packet must reach all other sources in the area before the end of the normal waiting period DIFS in order to avoid that the acknowledgement packet collides with regular packets. The maximum roundtrip delay (the maximum time from sending the last bit of the information packet and until detecting the first bit of the acknowledgement packet) must thus be shorter than two slot times (40 μs). This corresponds to maximum distance between sources of 3000 m which is far more than the diameter of the coverage areas of practical WLANs.

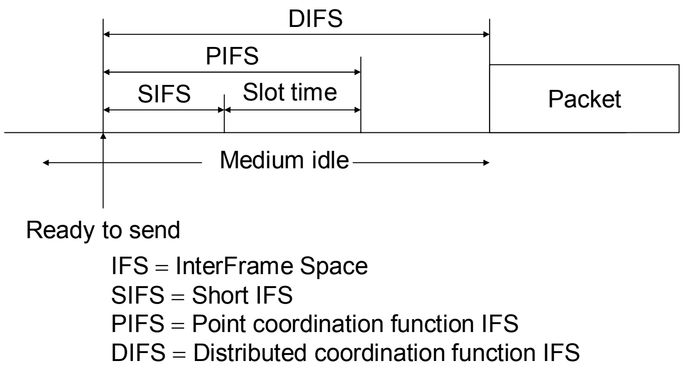

Figure 5.15 Inter-frame space

The collision probability can be further reduced if a reservation technique is employed where a source is using normal collision avoidance procedures to send a short request-to-send (RTS) packet. The receiver of the RTS packet returns a clear-to-send (CTS) packet during the collision-less interval, that is, just after the SIF in the same way as for acknowledgement packets. Since all sources will be aware of the reservation, the information packets for which reservation is made will not meet collisions. This method is specified as an option in IEEE 802.11 WLANs.

c) *What is a VSAT system and where (and why) is such a system used*?

Local area networks can be implemented using VSAT technology. VSAT networks have a number of applications where other solutions are impractical or impossible. Examples are:

- Interconnecting remote sites such as hydropower plants and dams with central operation, management and control systems.
- Providing temporary telecommunications capabilities for construction sites.
- Setting up local area networks for companies with operations widely distributed over large distances such as shopping chains.

The most important advantage of VSAT systems is that they are easy to install, expand and reconfigure.

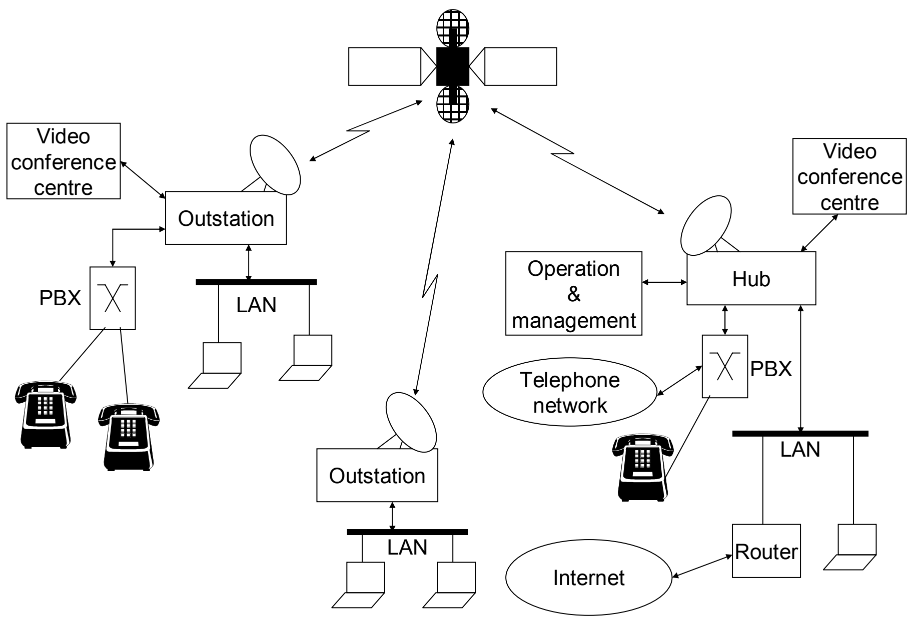

Figure 9.10 VSAT system

The most common configuration of a VSAT system is a star network consisting of a central station or hub and a number of outstations. Communication between outstations must then be relayed via the hub. The VSAT system can also be configured as a mesh network (or as a combination of star network and mesh network) where one of the stations is designated as the network controlling station. In this configuration, the outstations may communicate directly with one another. The VSAT system may offer transmission of data, voice and video signals. The typical configuration of a VSAT network is shown in Figure 9.10.

PBX designates a private branch exchange interconnecting the local telephone system of outstations and hubs.

TDMA is employed on the direction from outstation to hub where one or more timeslots are assigned to each outstation. TDM on a single carrier is used in the direction from the hub.

The most common frequency bands are in the 14/11 GHz range; that is, in a band around 11 GHz for the downlink and around 14 GHz for the uplink. The antenna size is in the range of 1.2 to 1.8 m.

d) *What do we mean by an isochronous signal*?

This term applies to a single signal. ITU defines the term as follows: The essential characteristic of a time-scale or signal such that the time intervals between significant instants either have the same duration or durations that are integral multiples of the shortest duration.

All signals consisting of bits of the same duration are isochronous. This means that all digital signals we are considering in this course are isochronous at this level of organisation. However, as we shall see below, there are other organising elements of the signal than the bits. The signal may not be isochronous with regard to these elements. Therefore, we must be careful when defining what is meant by the term in different contexts.

It is important to note that the bit rate will fluctuate along the isochronous signal.