(1a) Figure 1 shows three computers X, Y, and computer Z connected to two Local Area Networks (LANs). Outline the operation of the address tables within a bridge. Illustrate your answer by showing how the bridge in the figure recognises whether packets from computer X are to be forwarded from LAN A to LAN B. [6 marks]

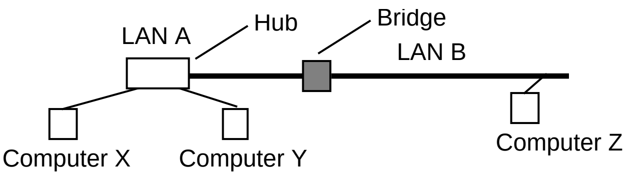

A bridge is a LAN interconnection device which may be used to join two LAN segments (A,B), constructing a larger LAN. A bridge is able to filter traffic passing between the two LANs and may enforce a security policy separating different work groups located on each of the LANs.

The bridge learns which MAC addresses belong to the computers on each conected subnetwork by observing the source address values which originate on each side of the bridge and storing them in an address table. This is called "learning". In the figure in the question, the source addresses X,Y are observed to be on network A, while the address of computer Z will be observed to be on network B.

The learned addresses are stored in the corresponding interface address table (there is a separate table for each active interface). Once this table has been setup, the bridge examines the destination address of all packet, and forwards them only if the address does not correspond to the source address of a computer on the local subnetwork. A system administrator may overide the normal forwarding by inserting entries in a filter table to inihibit forwarding between different workgroups (for example to provide security).

Packets with a source of X and destination of Y are received and discarded, since the computer Y is directly connected to the LAN A, whereas packets from X with a destiantion of Z are forwarded to network B by the bridge.

(1b) The IEEE 802.x LAN family supports many physical media, explain with the aid of diagrams the differences between 10BT (twisted pair) and 10B2 (coaxial cable) technologies [6 marks]

6

The basic difference between the two technologies are summarised below.

The 10BT cabling system uses a RJ-45 connector and 100 Ohm unshielded twisted pair cabling. This connects the computer directly (i.e. using a point to point link) to a wiring hub which acts as a media repeater. The maximum distance of a 10BT link is 100 m. It is normally used to connect work groups of users, sometimes by wiring an entire floor with outlets to each work area.

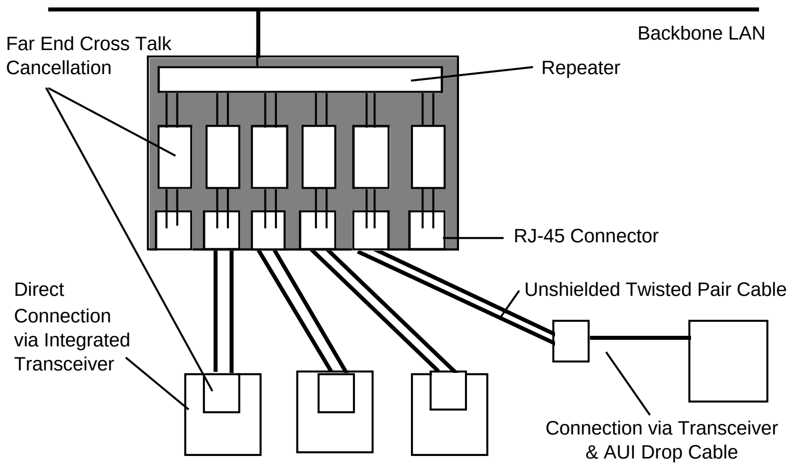
Far End Cross Talk Cancellation
Backbone LAN
Repeater
RJ-45 Connector
Direct Connection via Integrated Transceiver
Unshielded Twisted Pair Cable
Connection via Transceiver & AUI Drop Cable

Summary

Segment length 0.6m – 100m using cable which is flexible and very cheap

RJ-45 connector used which is often integrated into the computer or via external transceiver

Used mainly for workgroups, it is easy to manage

The 10B2 cabling system uses thin (RG-58U) co-axial cable which forms a shared bus. Upto 30 transceivers may be used to connect computers to form a bus. Each end of gthe bus must be terminated using a 50 Ohm termination resistor. this prevents refelection from the cable ends. Computers are connected via a "T" piece, which must be plugged directly into a NIC. 10B2 cabling may be used for backbone connections or to connect work groups. It is now fairly uncommon to find this type of cabling using to connect user's workstations, since 10BT has largely replaced this in corporate networks - since it is more flexible to use (supporting also telephone lines, video, 100BT).

Summary

10B2 uses 50 Ohm coaxial cable providing reasonable noise immunity
Segment length ≤ 185m, cable run needs careful installation
BNC-Type connector used with built-in or external transceiver

(1c) With the help of a frame transition diagram describe in detail the operation of the IP Address Resolution Protocol (arp) [8 marks]

6 The term address resolution refers to the process of finding an address of a computer in a network. The address is “resolved” using a protocol in which a piece of information is sent by a client process executing on the local computer to a server process executing on a remote computer. The information received by the server allows the server to uniquely identify the network system for which the address was required and therefor to provide the required address. The address resolution procedure is completed when the client receives a response from the server containing the required address.

The Ethernet network uses two hardware addresses which identify the source and destination of each frame sent by the Ethernet. The destination address (all 1’s) may also identify a broadcast packet (to be sent to all connected computers) or a multicast packet (msb=1) (to be sent only to a selected group of computers). The hardware address is also known as the Medium Access Control (MAC) address, in reference to the IEEE 802.x series of standards which define Ethernet. Each computer network interface card is allocated a globall unique 6 byte address when the factory manufactures the card (stored in a PROM). This is the normal source address used by an interface. A computer sends all packets which it creates with its own hardware source address, and receives all packets which match its hardware address or the broadcast address. When configured to use multicast, a selection of multicast hardware addresses may also be received.

The Ethernet address is a link layer address and is dependent on the interface card which is used. IP operates at the network layer and is not concerned with the addresses of individual links which are to be used. A protocol known as address resolution protocol (arp) is therefore used to translate between the two types of address. The arp client and server processes operate on all computers using IP over Ethernet. The processes are normally implemented as part of the software driver which drives the network interface card.

To reduce the number of address resolution requests, a client normally caches resolved addresses for a (short period of time. The arp cache is of a finite size, and would become full of incomplete and obsolete entries for computers that are not in use if it was allowed to grow without check. The arp cache is therefore periodically flushed of all entries. This deletes unused entries and frees space in the cache. It also removes any unsuccessful attempts to contact computers which are not currently running.

8

Question Number	2	Solution	Page of 10

Mark

(2a) With the help of diagrams, explain the operation of the framing provided by the HDLC protocol. Pay particular attention to define the following terms:

2/8 Flag

HDLC is a data link protocol which uses a unique bit sequence to delimit the start and end of each PDU transported by the data link layer service. In HDLC, frames are delimited by a sequence of bits known as a "flag". The flag sequence is a unique 8-bit sequence of the form 0111 110. The way in which this is performed is described in the text and diagrams which follow.

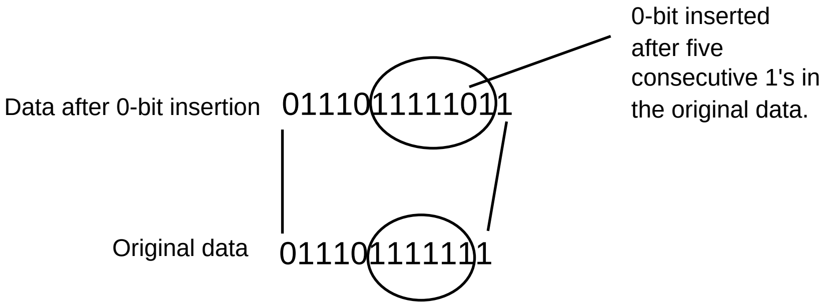
Data after 0-bit insertion
0111011111011
Original data
011101111111
0-bit inserted after five consecutive 1's in the original data.

2/8 Transparency

The flag sequence never occurs within the content of a frame becaus a technique known as 0-bit insertion is used to prevent random data synthesising a flag. The technique is said to make HDLC transparent, since any stream of bits may be present between the open and closin flag of a frame. The transparency is achieved by encoding the data by inserting a 0-bit after any sequence of 5 consecutive 1's within the payload as shown.

2/8 Hunt Mode

Normally a HDLC receiver starts in the idle state, waiting for the start of a frame. This is called "Hunt" mode, since the receiver is said to be hunting for a non-flag sequence. This may be achieved through a shift register and combinational logic as shown (a Finite State Machine (FSM) may also be used to implement this)

2/8 Abort & Idle

Valid frames are terminated by a closing flag. An error during transmission, or pre-emption by a higher priority frame, causes the frame to be terminated by an "abort" sequence: 0111 1111. A frame which is terminated by an abort is received by the receiver in the normal way, but marked as being "aborted". The frame is then discarded without further processing. This is shown below:

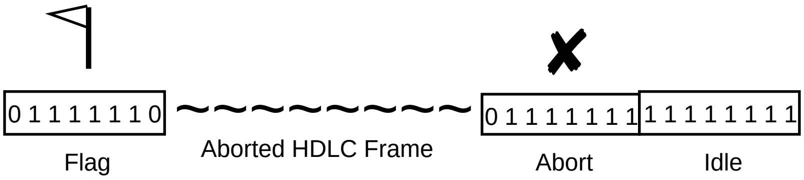

An abort sequence is often followed by a series of 1's. The sequence of all 1's may be used to fill the gaps between frames (or alternatively a continuous series of flags may be transmitted). The all 1's sequence is known as the "idle" sequence, since the line becomes idle (N.B. represented by 0 Volts).

8

Question Number	2	Solution	Page of 10

Mark

(2b) With reference to the IP network layer protocol describe the following terms:

(i) Network layer address [2 marks]

An address is a data structure understood by a network which uniquely identifies the recipient within the network. An IP address is a 32 bit value consisting of two parts, the network part (identifying the network to which the computer is attached) and the host part (which identifies the host within the local network). The IP network address is identified as the bit-wise logical AND of the netmask and the 32-bit IP address.

An address is a unique network identifier consisting of network part and host part. Each host has at least one address.

2

(ii) Fragmentation (also known as segmentation) [2 marks]

The layers which use the layers of a service below are not always aware of the maximum size of SDU which may be supported. In most cases packet networks limit the size of the maximum SDU at the network layer, but the actual maximum size will depend upon the network architecture which is being used. (LANs and MANs often allow comparatively large packets, whereas WANs often employ a much smaller maximum packet size). Many layers therefore support a segmentation (also known as fragmentation) service, which breaks large SDUs into a number of smaller SDUs. The corresponding peer protocol is responsible for reassembling the complete SDU before forwarding to the layer above.

2

(iii) The "more" flag in the IP header [2 marks]

Fragmentation / Segmentation allows large PDUs to be broken into smaller units. Each fragmented PDU carries a copy of the oiriginal PCI for the layer at which fragmentation occurs. The subsequent fragments usually modify the PCI to indicate that they continuation fragments. More is set in all but the last segment.

2

(iv) Maximum Transfer Unit (MTU) [2 marks]

The maximum transfer unit is the largest size of IP datagram which may be transferred using a specific data link connection The MTU value is a design parameter of a LAN and a mutually agreed value for most WAN links. The size of MTU may vary greatly between different links (from 128 B upto 10 kB) and is the reason why fragmentation/segmentation is used at intermediate systems.

2

(2c) Calculate the number of fragments which are sent when an IP datagram with payload of 3000 bytes is sent from a computer using a network connection with an MTU of 512 bytes. Ensure that your answer specifies the number and size of each of the IP datagrams sent [4 marks]

1 Total size of initial PDU = 3000 + 20 B (PCI) = 3020.
MTU 512 B < 3020 B - therefore fragmentation is required.

1 Fragment payload size = 512- 20 B = 492 B
1 Total number of packets sent via network B = round(3000/492) = 7 packets.
1 First six packets of size 512 B, Last of size 48+20 B =68 B.

4

3a) The OSI Reference Model

The OSI reference model specifies standards for describing "Open Systems Interconnection" with the ter 'open' chosen to emphasise the fact that by using these international standards, a system may be define which is open to all other systems obeying the same standards throughout the world. The definition of a common technical language has been a major catalyst to the standardisation of communications protocols and th functions of a protocol layer.

The seven layers of the OSI reference model showing a connection between two end systems communicatin using one intermediate system.

The structure of the OSI architecture is given in the figure above, which indicates the protocols used to exchange data between two users A and B. The figure shows bidirectional (duplex) information flow; information in either direction passes through all seven layers at the end points. When the communication is via a network of intermediate systems, only the lower three layers of the OSI protocols are used in the intermediat systems.

The OSI layers may be summarised by:

Physical layer: Provides electrical, functional, and procedural characteristics to activate, maintain, and deactivate physical links that transparently send the bit stream; only recognises individual bits, not characters c multicharacter frames.

Data link layer: Provides functional and procedural means to transfer data between network entities an (possibly) correct transmission errors; provides for activation, maintenance, and deactivation of data link connections, grouping of bits into characters and message frames, character and frame synchronisation, errc control, media access control, and flow control.

Network layer: Provides independence from data transfer technology and relaying and routing considerations; masks peculiarities of data transfer medium from higher layers and provides switching and routin functions to establish, maintain, and terminate network layer connections and transfer data between users.

Transport layer: Provides transparent transfer of data between systems, relieving upper layers from concern with providing reliable and cost effective data transfer; provides end-to-end control and information interchange with quality of service needed by the application program; first true end-to-end layer.

Session layer: Provides mechanisms for organising and structuring dialogues between application processes; mechanisms allow for two-way simultaneous or two-way alternate operation, establishment of major an minor synchronisation points, and techniques for structuring data exchanges.

Presentation layer: Provides independence to application processes from differences in data representation, that is, in syntax; syntax selection and conversion provided by allowing the user to select a "presentatio context" with conversion between alternative contexts.

Application layer: Concerned with the requirements of application. All application processes use the service elements provided by the application layer. The elements include library routines which perform interprocess communication, provide common procedures for constructing application protocols and for accessing th services provided by servers which reside on the network.

The communications engineer is concerned mainly with the protocols operating at the bottom four layer (physical, data link, network, and transport) in the OSI reference model. These layers provide the basic communications service. The layers above are primarily the concern of computer scientists who wish to build distributed applications programs using the services provided by the network.

10 [4 marks for OSI stack + 6 marks for detail = 10 marks]

Mark

(3b) The Universal Datagram Protocol (UDP) is a simple transport protocol supported b the Internet Protocol (IP) suite. The format of a UDP packet is shown above. Explai the function of each of the component fields of the UDP packet header. [6 Marks]

Ports are used as a service access point to assocaite the transport connection with the processes handling th middleware within the end systems.

The length field indicates the total length of the UDP PDU (including the length of the UDP hedaer itself).

The checksum verifies the integrity of the received PDU as receceived at the end system. It includes all th usre data, the UDP header and also the key parts of the IP header (e.g. the src and dst network addresses). A value of zero indicates that the sender has not calculated a checksum - and therefore that the receiver ma not perform this integrity check. Although this is allowed, it is strongly recommended that a checksum i used in every UDP PDU.

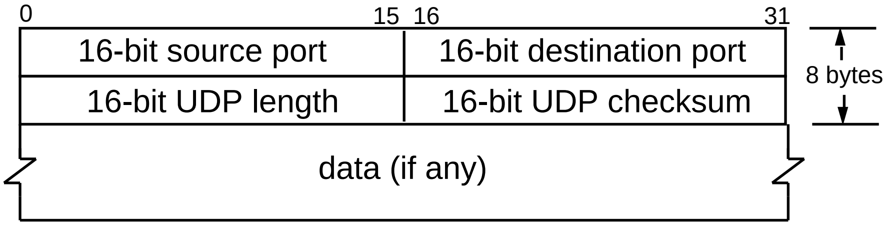

0 … 15	16 … 31
16-bit source port	16-bit destination port
16-bit UDP length	16-bit UDP checksum
data (if any)	

6

(3c) A UDP packet containing 150 B of payload data is transmitted using IP over an Ether net LAN. Draw a diagram showing the transmitted frame, including all protocol headers. What is the total size of the frame sent using the Ethernet LAN? [6 Marks]

Students should provide a protocol header diagram showing the following PCI:

IFG (9.5-10.6 microssec) + Preamble (8B) + MAC (src+dst+type = 14B) + IP Header (20B) +
+ UDP Header (8 B) + IP payload (150B) + CRC-32 (4B)

Students should calculate the total PDU size:

Total size =
8+14+20+ 8 +150+4 = 204 B (ignoring IFG).

6

Mark

(4a) Explain the terms Wide Area Network,[2 Marks] Metropolitan Area Network, [2 Marks] and Local Area Network,[2 Marks].

The Wide Area Network (WAN) usually refers to a network which covers a large geographical area, and use telecommunications circuits to connect the intermediate nodes. A major factor impacting WAN design and performance is a requirement that they lease communications circuits from telephone companies or other communications carriers. This restricts the communications facilities, and transmission speeds, to those normally provided by such companies. Transmission rates are typically 56 kbps, 64 kbps, 2 Mbps, 34 Mbps, 45 Mbps, or sometimes considerably less (e.g. 28.8 kbps, 9.6 kbps, or slower)

The characteristics of the transmission facilities lead to an emphasis on efficiency of communications techniques in the design of WANs. Flow control to limit traffic and avoid excessive delays is important, as is recovery from transmission errors. Since the topologies of WANs are likely to be more complex than those of LANs, routing algorithms also receive more emphasis. Many WANs also implement sophisticated monitoring procedures to account for which users consume the network resources. This is, in some cases, used to generate billing information to charge individual users.

The Local Area Network (LAN) is by far the most common type of data network. As the name suggests, a LAN serves a local area (typically the area of a floor of a building, but in some cases spanning a distance of several kilometres). Typical installations are in industrial plants, office buildings, college or university campuses, or similar locations. In these locations, it is feasible for the owning organisation to install high quality high-speed communication links interconnecting nodes. Typical data transmission speeds are one to 100 megabits per second. A wide variety of LANs have been built and installed, but a few types have more recently become dominant. The Ethernet network familly (also known as IEEE 802.3) which operates at 10 Mbps is the most common LAN technology.

The Metropolitan Area Network (MAN) is a relatively new class of network. There are three important features which discriminate MANs from LANs or WANs.

First, the network size falls intermediate between LANs and WANs. A MAN typically covers an area of between 5 and 50 km diameter. Many MANs cover an area the size of a city, although in some cases MANs may be as small as a group of buildings or as large as the North of Scotland.

Second, a MAN (like a WAN) is not generally owned by a single organisation. The MAN, its communications links and equipment are generally owned by either a consortium of users or by a single network provider who sells the service to the users. This level of service provided to each user must therefore be negotiated with the MAN operator, and some performance guarantees are normally specified.

6 The third difference between a LAN and a MAN is that the MAN often acts as a high speed network to allow sharing of local resources (as in a LAN), but also allows connection to other networks via a WAN.

(4b) Discuss the use of fibre optic cabling in Wide Area Networks, suggesting reasons for the trend to increasingly replace copper conductors with optical fibre. [4 Marks]

Fibre Optic cable propagate the signal as a pulse of light along a transparent medium. The refractive index of the medium varies as a function of the distance from the centre of the fibre, resulting in the light being guided along the fibre. The outside of the fibre is protected by cladding and may be further protected by additional layers of treated paper, PVC or metal. This required to protect the fibre from mechanical deformation and the ingress of water.

The principle reasons for the increasing use of fibre in WANs and MANs are:

Lower signal loss per unit distance (resulting in longer distances between repeaters)
Higher Capacity (allowing operation at higher data rates)
4 Smaller physical size (allowing more fibres in a duct or trunk)

(4c) Various types of equipment may be used to connect parts of a large network. Summarise the differences between a router, repeater and a bridge. [8 Marks]

The fundamental differences between the three types of equipment are outline below. Each operates at a separate layer of the OSI reference model (see also figure).

LAN Repeaters – Join LAN segments
Very cheap
Regenerate the signal and timing information
Allow multiple types of media to be connected
Work below the MAC Layer (Support all protocols)
Build one single LAN

Bridges – Separate work group traffic
Cheap
Allow multiple types of media to be connected (also known as a “hub” or “switch”)
Work at the MAC Layer (Support all protocols)
May provide filtering to implement simple security policies
Build one single IP network

Routers – Connect IP networks
More Expensive
Work at Network Layer (e.g. IP) and support one or more protocols
Connect separate networks into an internet
Connects different link protocols and media (e.g. HDLC to Ethernet)
Need more detailed configuration but may protect networks from unauthorised access 8

(4d) Every computer in a LAN requires a unique Medium Access and Control (MAC) address. Explain how these unique addresses are allocated.

Each NIC is assigned a unique address contained in a prom/eprom on the card. Addresses are 6-bytes long, with half the address space reserved for multicast/broadcast and the all-zero address reserved for "unkown" source addresses. Addresses are assigned in a groups using flat addressing. The manufacturers pay for for groups of addresses. Since each NIC is globally unique, this satisfies the need to have a locally unique address. 2

ark

5a) Describe the features required in a communications protocol to provide a Reliable Service.

Reliable delivery has been succinctly defined as "Data is accepted at one end of a link in the same order as wa transmitted at the other end, without loss and without duplicates." This implies four constraints:

(i) No loss (at least one copy of each frame is sent)
(ii) No duplication (no more than one copy is sent)
(iii) FIFO delivery (the frames are forwarded in the original order)
(iv) A frame must be delivered within a reasonable period

For a communications protocol to support reliability, requires that the protocol numbers the PDUs that ar transmitted, implements an error recovery procedure (e.g. checkpointing or go-back-N), and provides erro free procedures for link management.

There is very little data which is so important that it must be sent no matter how late. Layered protocols usual ly also employ timers at each level, governing this interval. The service provided by a protocol layer may b unreliable for various reasons including:

(i) Corruption of bits within the physical medium or the interface to the physical media.
(ii) Faulty bit-timing resulting in erroneous decoding of the value of a received bit.
4 (iii) A software error within the software used to implement the communications protocol.
(iv) Insufficient buffer space within the communications equipment.

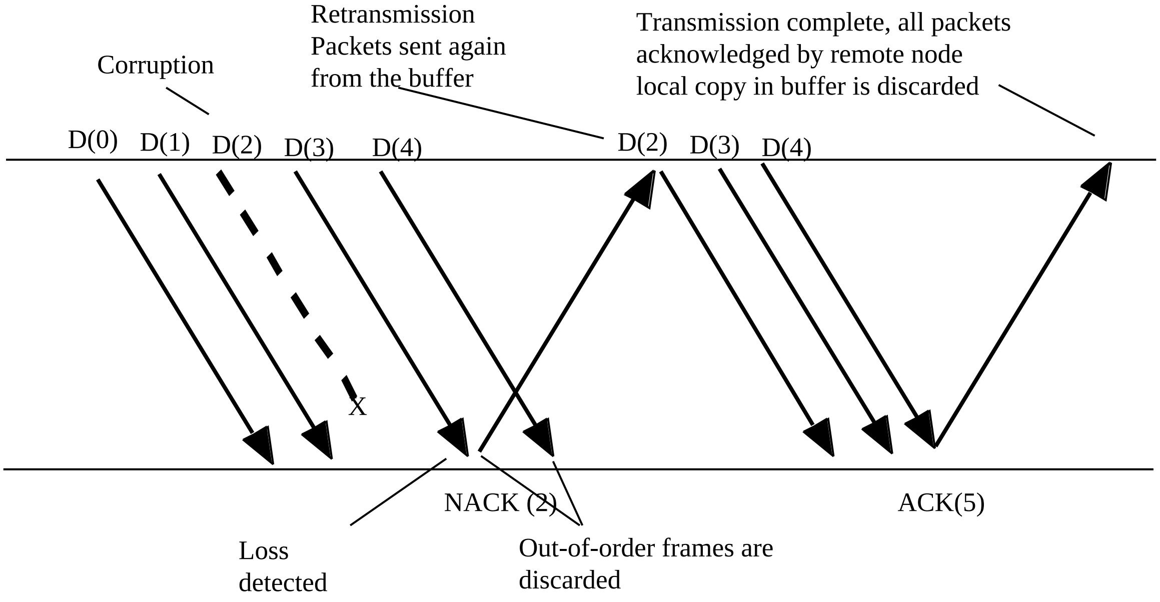

5b) Go_Back_N Recovery

Description to support diagram.

The recovery of a corrupted I-frame proceeds in three stages:

First, the corrupted frame is discarded at the remote node's receiver.
The loss of the I-frame is revealed when a correct (but out-of-sequence) I-frame is received.
Second, the remote node requests retransmission of the missing I-frame(s).
The final stage consists of retransmission of the lost I-frame(s).

Question Number	5	Solution	Page of 10

Mark

A remote node may request retransmission of corrupted I-frames by initiating Go-Back-N error recovery b
sending a REJ (reject) frame. The remote node sends a REJ frame to instruct the sending station to begin re
transmission of I-frames at the frame number indicated (in the ACK value of the NAK(REJ) packet). Sinc
the remote node does not store any out-of-sequence frames, the ACK value corresponds to the next expecte
in sequence I-frame (i.e. the receive state variable (V(R)). The receiver continues to discard received I-frame
until one is received with the expected sequence numnber (i.e. V(R)=N(S)).

Upon receipt of a REJ frame (by the local node), the transmitter winds-back the sequence of I-frames pendin
transmission to the indicated I-frame (i.e. the send state variable (V(S)) is assigned the value of the receive
sequence number (N(R))). The transmitter then retransmits the requested I-frame followed by all successive I
frames. This is sometimes known as "wind back" of the transmitter.

6 5c) **Explain why event timers form an important part of a reliable protocol.**

Unless it is guaranteed that everything arrives punctually (almost impossible in the real world, let alone th
communications world!), a mechanism must be provided for recovering from the loss of one or more frame
(e.g. to errors on a physical link). For this purpose, timers are integrated into a protocol.

Timers are pieces of hardware or software (usually provided as routines in an operating system) which cour
down from some preset timer value. Timers are started and stopped whenever particular types of events oc
cur. In most cases, timers are stopped before their time out period expires, and no action is taken. If the cour
should reach zero, then the timer is said to have "expired". This normally triggers a specific event within
5 protocol. HDLC identifies a number of timers that may be used by an HDLC link (ABM Mode) to provid
reliable data transmission.

5d) **Describe the Internet Control Message Prococol (ICMP) [5 marks]**

The ping ultility allows a client to generate ICMP echo request messages which are encapsulated in IP data
grams. Each system in the Internet operates an ICMP echo server which, which responds to ICMP echo re
quests by generating an ICMP echo reply back to the originator. The returned message has an Ip header wit
the source and destination addresses reversed and carries the same payload as originally sent. By including
sequnec number and timing the response, a client can determine whether the network path is operationald an
also measure the current round trip delay and packet loss rate.