Question Number	1	Solution	Page of 12

Mark

1. (a) Ethernet is now the most common networking technology for the construction of Local Area Networks (LANs) and supports a range of physical media. Explain the differences between 10BT and 10B2 cabling. [8 marks]

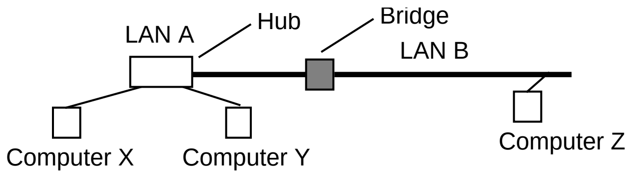

The basic difference between the two technologies are summarised below.

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.

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 the bus must be terminated using a 50 Ohm termination resistor. this prevents reflection 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

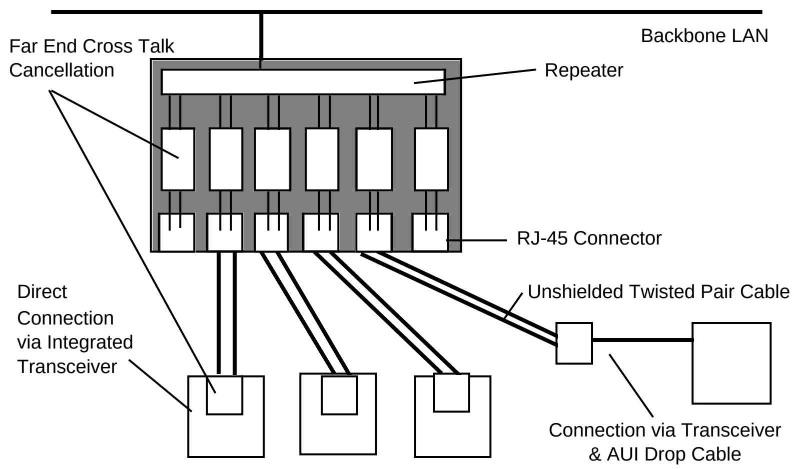

8

Question Number	1	Solution	Page of 12

Mark

(b) Ethernet supports a protocol known as Carrier Sense Multiple Access with Collision Detection (CSMA/CD). Explain how CSMA/CD works, giving an example of how it ensures a low probability of collision when two nodes attempt to transmit at the same time? [10 marks]

Ethernet uses a refinement of ALOHA, known as CSMA, which improves performance when there is a higher medium utilisation. When a node has data to transmit, the node first listens to the cable (using a transceiver) to see if a carrier (signal) is being transmitted by another node. This may be achieved by monitoring whether a current is flowing in the cable (each bit corresponds to 18-20 milliAmps (mA)). The Ethernet transceiver contains the electronics to perform this detection. Data is only sent when no carrier is observed (i.e. no current present) and the physical medium is therefore idle.

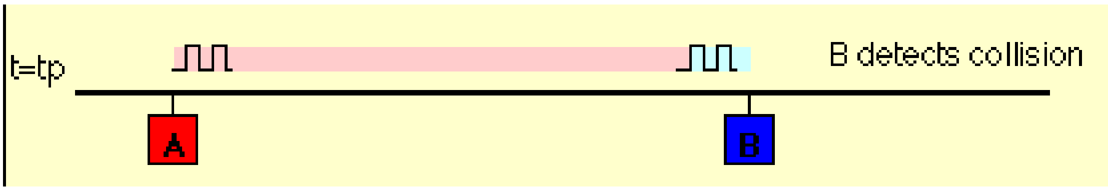

However, this alone is unable to prevent two nodes transmitting at the same time. If two noes simultaneously try transmit, then both could see an idle physical medium (i.e. neither will see the other's carrier signal), and both will conclude that no other node is currently using the network. In this case, both will then decide to transmit and a collision will occur. The collision will result in the corruption of the data being sent, which will subsequently be discarded by the receiver since a corrupted Ethernet frame will not have a valid 32-bit MAC CRC at the end.

A second element to the Ethernet access protocol is used to detect when a collision occurs. Each transmitting node monitors its own transmission, and if it observes a collision (i.e. excess current above what it is generating, i.e. > 24 mA) it stops transmission immediately and instead transmits a 32-bit jam sequence.

To ensure that no node may completely receive a frame before the transmitting node has finished sending it, Ethernet defines a minimum frame size (i.e. no frame may have less than 46 bytes of payload). The minimum frame size is related to the distance which the network spans, the type of media being used and the number of repeaters which the signal may have to pass through to reach the furthest part of the LAN. Together these define a value known as the Ethernet Slot Time.

When two or more transmitters each detect a corruption of their own data (i.e. a collision), each responds in the same way by transmitting the jam sequence. At time t=0, a frame is sent on the idle medium by computer A. A short time later, computer B also transmits. (In this case, the medium, as observed by the computer at B happens to be idle too). After a period, equal to the propagation delay of the network, the computer B detects the other transmission from A, and is aware of a collision, but computer A has not yet observed that computer B was also transmitting. B continues to transmit, sending the Ethernet Jam sequence (32 bits).

After one complete round trip propagation time (twice the one way propagation delay), both computers are aware of the collision. B will shortly cease transmission of the Jam Sequence, however A will continue to transmit a complete Jam Sequence. Finally the cable becomes idle.

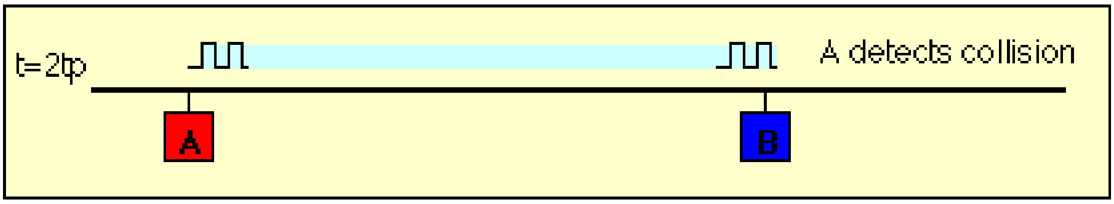

10

Question Number		Solution	Page of 12

Mark

(c) What is the purpose of the "Jam sequence"? [2 marks]

2

The Jam sequence is sent when a node detects that another node is using the shared medium. When this is detected, the sender stops transmission, and instead sends a sequence of 32 bits, known as the "Jam sequence". The purpose of this sequence is to ensure that any other node which may currently be receiving this frame will receive the jam signal in place of the correct 32-bit MAC CRC, this causes the other receivers to discard the frame due to a CRC error.

Mark

2. (a) The following frame transition diagram (figure 1) shows an exchange of Ethernet frames between two computers, A and B connected via a 10BT Hub. Each frame sent by Computer A contains 1500 B of Ethernet payload data, while each frame sent by Computer B contains 40 B of Ethernet payload data.
Calculate the average Utilisation of the media during this exchange. [6 marks]

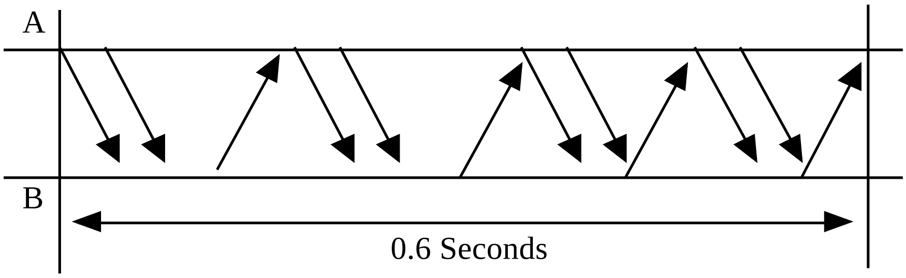

FIGURE 1: Frame Transition Diagram for Communication Between A and B

No Frames from A = 8
Ethernet MAC Frame Payload = 1500B (comprised of = 20 B (IP) + 20 B (TCP) +1460 B DATA)
Total A Frame Size = 8 B (Preamble) + 14 B (Mac) + 1500 B + 4 B (CRC-32)
= 8+14+1500+4 = 1526 B= 12208 b

No of Frames from B = 4
Total B Frame Size = 40B (20 B (IP) + 20 B (TCP) + NO DATA)
Total B Frame Size = 8 B (Preamble) + 14 B (Mac) + 40 + 6 B PAD + 4 B (CRC-32)
= 8+ 60 + 4 = 72 B = 576 b

(I have ignored Inter-Frame Gap, IFG, which could be included as overhead).
Total Utilised Bandwidth in this period = 1526 x 8 x 8 + 72 x 4 x 8 b = 97664+ 2304 b = 99968

Utilisation = (99968 x 100)/(.6x 10E7) = 1.7 %

6 **(b) Is the exchange in figure 1 best described as Full Duplex, Half Duplex, or Simplex? [2 marks]**

Half Duplex - the two ends alternately take the opportunity to send.

2 **(c) What is the throughput of the transfer from A to B measured at the TCP layer? [6 marks]**

The MAC payload is therefore 1500B which contains the following PDU:

(20 B (IP) + 20 B (TCP) +1460 B DATA)

Volume of data sent per second is = 1460* 8*8 /0.6 = 155.733 kbps
(averaging over the 0.6 second period).

(d) Two Ethernet LANs are connected by an Ethernet bridge. Explain how the bridge automatically recognises which packets are to be forwarded and which are to be discarded. [4 marks]

6 Abridge learns by observing the MAC source addresses belong to the computers on each connected subnetwork by observing the source address values which originate on each side of the bridge. This is called "learning". In the figure in the question, the source addresses X,Y are observed to be on network A, while the ad-

Mark

dress of computer Z will be observed to be on network B.

A bridge stores the hardware addresses observed from frames received by each interface and uses this infor-

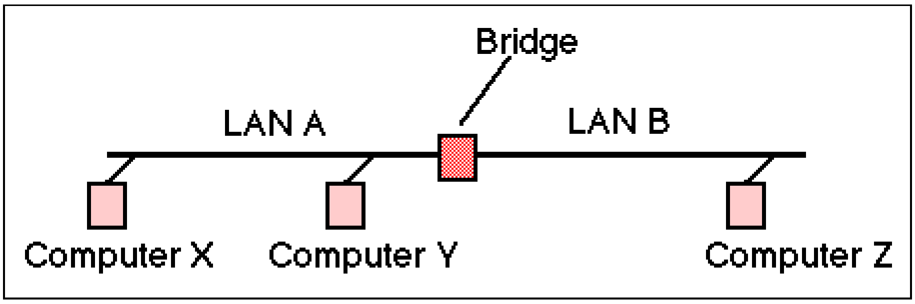

mation to learn which frames need to be forwarded by the bridge. Packets with a source of X and destination of Y are received and discarded, since the computer Y is directly connected to the LAN A, whereas packets from X with a destination of Z are forwarded to network B by the bridge

The learned addresses are stored in the corresponding interface address table. Once this table has been setup, the bridge examines the destination address of all frames, and forwards them only if the address does not correspond to the source address of a computer on the local subnetwork.

(e) What complication arises when a second Ethernet Bridge is connected in parallel with the first? [2 marks]

Packets are forwarded initially by both bridges - the destination will receive two copies.

4

Bridges see the source address on both interfaces (they may learn a false location of the node within the network. This could lead to them forwarding each others forwarded packets - resulting in a forwarding loop. The level of traffic on the LANs grows as a result of this error.)

(Spanning Tree is an algorithm which elects a bridge as a root of a forwarding tree, by sending and receiving spanning tree frames. The ST algorithm sets one bridge to the “blocked” mode. A tree is then built which assures that there are no loops (as above). If a loop is found, one of the two paths is disabled by setting the bridge to the “blocking” mode. In blocking mode, the bridge will not forward any Ethernet frames. The bridge continues to listen to ST frames, so that it can detect if another bridge has stopped forwarding, and therefore if there is a need to remove the blocked state. Spanning Tree provides a robust solution, protecting from individual bridge failure.)

2

Question Number	3	Solution	Page of 12

Mark

3. (a) Sketch the Open Systems Interconnection (OSI) Reference Model and describe the services provided by each layer. [8 Marks]

The OSI reference model specifies standards for describing "Open Systems Interconnection" with the term 'open' chosen to emphasise the fact that by using these international standards, a system may be defined 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 the functions of a protocol layer.

The seven layers of the OSI reference model showing a connection between two end systems communicating 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 intermediate systems. The OSI layers may be summarised by:

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.

Data link layer: Provides functional and procedural means to transfer data between network entities and (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, 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 routing 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 and 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 "presentation 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 the services provided by servers which reside on the network.

8

(b) What layer of the OSI reference model best describes the Internet Protocol (IP)? [1 mark]

Network Layer (layer 3 of the OSI RM).

1

(c) The Internet Protocol (IP) may be used over both HDLC and Ethernet links. For each

Mark

type of link explain how the link layer identifies the first and last byte of the frame data.

(i) First Byte of Frame Header in an Ethernet frame

Start of the frame is detected by a signal (current) in the medium. The receiver then attempts to recover the data, but first the receiver DPLL must acquire lock. After detecting bit timing, the receiver looks for the SFD (ending 11). The following byte is the first of the frame.

(ii) Last Byte of Frame Checksum in an Ethernet frame [5 marks]

The end of the frame is detected by an absence of signal (current) in the medium. When the receiver senses this, it treats the last 32 bits as the frame CRC value. It then gets the accumulated CRC value and accepts or discards the frame, based on the CRC value. The receiver hardware will also discard frames which are too short (RUNTS) or too long (JABBER). It also discards any frame which contains residual bits (i.e. where the total number of bits between the start and end of frame do not correspond to an integral number of bytes.

5

(d) Explain how the link layer identifies the first and last byte of a frame data when an IP packet is sent over a link using the High Level Data Link Control (HDLC) protocol.

(i) First Byte of Frame Header in HDLC

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 flag sequence never occurs within the content of a frame because 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 closing 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. In HDLC, the gaps between frames are filled by an idle sequence (usually continuous flag bytes). The start

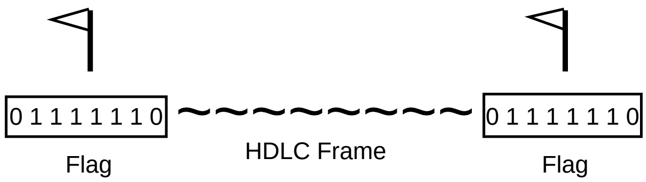

of a frame is detected by reception of a byte which does not have the flag value (0111 1110).

(ii) Last Byte of Frame Checksum in HDLC [6 marks]

In HDLC, the end of a frame is marked by the following byte being a flag byte with the value 0111 1110. The receiver therefore uses the terminating flag to deduce that the previous two bytes were the CRC value of the-frame. This requires a FIFO at the receiver to hold each byte until the next is received. The scheme also allows for an arbitrary number of bits per frame (i.e. a frame does not necessarily contain an integral number of bytes, although in practice this is normally the case.

6

Mark

4. (a) When an Internet Protocol (IP) packet is carried in a link layer frame, it contains addresses at both the link layer and network layer.

(i) How does the node identify its own hardware address? [4 marks]

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 globally unique 6 byte address when the factory manufactures the card (stored in a PROM). This is the normal source address used by an interface. The IEEE manages blocks of sequentially assigned addresses which the manufacturers of Ethernet equipment are required to purchase.

4

(ii) How does the node identify the hardware address of the intended recipient? [4 marks]

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 therefore 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 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.

4

(iii) How does the node identify it’s own IP 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. The address is configured by user/network manager.

2

(iv) What protocol is used to find out the IP address of the intended destination end system, when only the name of the system is known? [2 marks]

The sender uses the DNS protocol. In the DNS, there are a set of root domain servers (rather like the old Stanford computer), but they don't actually store much information. Instead they contain the IP addresses of other servers which have information about specific groups of addresses known as "domains". The root server is said to delegate responsibility for each domain to a lower domain server. In turn, each of these servers may delegate other domains to other servers. Before long, there were many many domain servers each responsible for the groups of users in a local area. Each server maintained pointers allowing them to find out information about other domains by sending query messages to the other domain servers. In this way, any DNS server can resolve the name of any computer to an IP address of any user irrespective of whether that user is in the same local domain or is registered with some remote domain.

2

Question Number	4	Solution	Page of 12

Mark

(b) A path between two End Systems consists of three Ethernet links with an Maximum Transmission Unit (MTU) of 1500 B, 1450 B, and 1500 B. By first calculating the Path-MTU, determine how many IP packets are received when a UDP datagram of 8 KB is sent over such a path using (i) MTU Discovery, and[4 marks] (ii) IP Fragmentation.
[4 marks each]

(i) With Path-MTU Discovery the sender originally sends a full-sized packet (i.e. with the maximum size dicated by its local MTU for the interface over which the packet is sent). The packet has the "Don't Fragment" (DF) bit set in the header. A router along the transmission path which has a smaller MTU than the frame size, discards the frame (since the DF-bit was set). It then returns an ICMP error message indicating the actual MTU of the link which caused the discard. The sender, on receiving an ICMP error message reduces the Path-MTU size for the specified destination. The packet is then fragmented again by the sender. All subsequent packets are fragmented according to the Path-MTU size.

4 PMTU = Min(1500, 1450, 1500) = 1450

IP Payload= PMTU-(IP Header)
= 1450 - 20 = 1430 B

No Fragments (after discovery of the PMTU) is:

4

(8000+8)/ (1430) = 5.6 = 6 packets (or 5 full fragments of 1450 B, plus one incomplete fragment)

(N.B.With IP router fragmentation:

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.

Original end system sends 6 packets.
Each of the first 5 packets exceed the MTU of the middle segment and are fragmented a second time.)

Question Number	5	Solution	Page of 12

Mark

5. (a) What is a router? [6 marks]

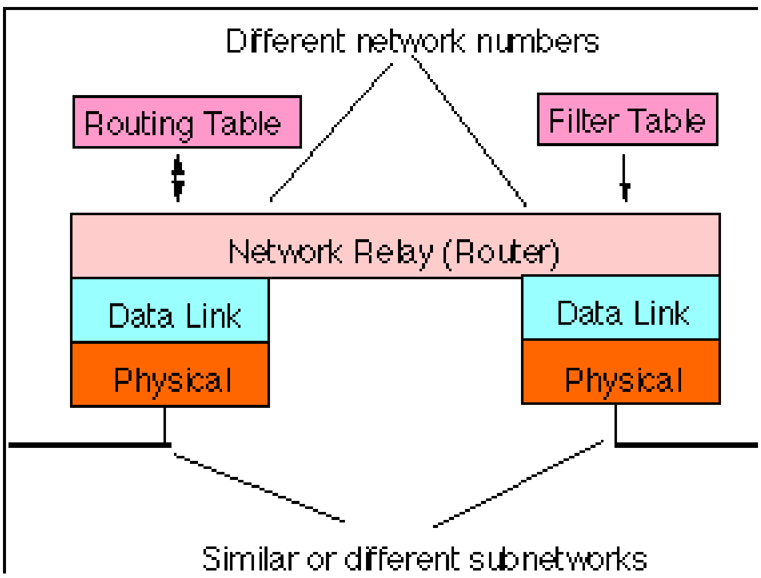

A router is an Intermediate System (IS) which operates at the network layer of the OSI reference model. Routers may be used to connect two or more IP networks, or an IP network to an internet.

A router is most suited for the connection of a LAN to a MAN. The router allows two separately administered networks to communicate without forming one homogenous network. The two networks may have different media, and belong to different IP networks (in the case of IP). The router also provides routing of packets to destinations reachable via the MAN and can control access to/from the MAN.

A router consists of a computer with at least two network interface cards supporting the IP protocol. The router receives packets from each interface and forwards the received packets to an appropriate output interface. The router uses the IP address, along with routing information held within the router and stored in a routing table, to determine the destination for each packet. A filter table may also be used to ensure that unwanted packets are discarded. The filter may be used to deny access to particular protocols or to prevent unauthorised access from remote computers.

Routers are often used to connect together networks which use different types of links (for instance an HDLC link connecting a WAN to a local Ethernet LAN). The optimum (and maximum) packet lengths (i.e. the Maximum Transfer Unit (MTU)) is different for different types of network. A router may therefore uses IP to provide segmentation of packets into a suitable size for transmission on a network.

Routers :
- Are more expensive than Bridges or Switches
- Work at Network Layer (e.g. IP) and support one or more protocols
- Connect separate networks into an internet
- May protect networks from unauthorised access

6

(b) The following packet was received by a router from an Ethernet interface.

```
  0: 0100 5e02 dc3e 00d0 bbf7 c6c0 0800 4500
 16: 00cc e206 0000 7111 a1a9 84b9 8476 e002
 32: dc3e 7982 7982 00b8 08a0 8005 dbc6 d721
 48: 69c0 0752 bb5f fe39 3600 8808 b120 8933
 64: 6219 9118 5128 ffc8 1321 bc10 933e aa23
 80: 3233 ba00 e892 a00c 1a3c 0a28 37ab 012d
 96: aca5 4819 9088 0b39 64ba 43a0 b9a8 04b3
112: 88b8 4bf8 3940 d024 0a98 8b0b 1703 0a3a
128: 8820 a381 a21f 3bc0 9298 e893 90bd 042a
144: 0a88 3287 59ab e980 1211 4002 2208 98b1
160: 7039 0b26 e898 99ab b118 a1aa a702 9ac4
176: 9128 ca21 7822 2971 090a 2194 98d0 27bb
192: 0958 8092 993f b3b0 2922 337a 0f88 8810
208: 8a29 0183 fb15 b888 0d4c
```

Question Number	5	Solution	Page of 12

Mark

Decodes to:

```
ETHER:  ----- Ether Header -----
ETHER:  Packet 33 arrived at 14:14:18.73
ETHER:  Packet size = 218 bytes
0100 5e02 dc3e
ETHER:  Destination = 1:0:5e:2:dc:3e, (multicast)
00d0 bbf7 c6c0
ETHER:  Source      = 0:d0:bb:f7:c6:c0,
0800
ETHER:  Ethertype = 0800 (IP)
        IP:   ----- IP Header -----
        45
        IP:   Version = 4
        IP:   Header length = 20 bytes
        00
        IP:   Type of service = 0x00
        IP:         xxx. .... = 0 (precedence)
        IP:         ...0 .... = normal delay
        IP:         .... 0... = normal throughput
        IP:         .... .0.. = normal reliability
        00cc
        IP:   Total length = 204 bytes
        e206
        IP:   Identification = 57862
        0000
        IP:   Flags = 0x0
        IP:         .0.. .... = may fragment
        IP:         ..0. .... = last fragment
        IP:   Fragment offset = 0 bytes
        71
        IP:   Time to live = 113 seconds/hops
        11
        IP:   Protocol = 17 (UDP)
        a1a9
        IP:   Header checksum = a1a9
        84b9 8476
        IP:   Source address = 132.185.132.118, simonl.kw.bbc.co.uk
        e002  dc3e
        IP:   Destination address = 224.2.220.62, 224.2.220.62
                IP:   No options
                UDP:  ----- UDP Header -----
                UDP:
                7982
                UDP:  Source port = 31106
                7982
                UDP:  Destination port = 31106
                00b8
                UDP:  Length = 184
                08a0
                UDP:  Checksum = 08a0
```

Question Number	5	Solution	Page of 12

Mark

What protocol is used at the transport layer? [2 marks]

The Internet Protocol (IP) provides two transport layer protocols:

The Universal Datagram Protocol (UDP) (Best Effort Service) (SAP=17)
The Transmission Control Protocol (TCP) (Reliable Service) (SAP=6)

The IP protocol type shows that this is UDP (i.e. 11 in hexadecimal). The UDP header consists of 8 bytes of data.

2

(b)Determine the hardware destination address, and IP network destination address. By examining the addresses used, determine whether the packet is multicast or unicast. [4 marks]

Unicast is the term used to describe communication where a piece of information is sent from one point to another point. In this case there is just one sender, and one receiver. Unicast transmission, in which a packet is sent from a single source to a specified destination, is still the predominant form of transmission on LANs.

Multicast is the term used to describe communication where a piece of information is sent from one or more points to a set of other points. 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 appearance of a multicast address on the cable is therefore as shown below (bits transmitted from left to right):

```
0                        23 IP Multicast Address Group  47
 |                        | <--------------------------->|
 1000 0000 0000 0000 0111 1010 xxxx xxx0 xxxx xxxx xxxx xxxx
 |                              |
 Multicast Bit                        0 = Internet Multicast
                                1 = Assigned for other uses
```

The address used is 0100 5e02 dc3e (N.B. Note that BYTE order rather than bit order means that the multicast mode is signalled in the least significant bit of the first byte of address.

4

(c) All transport layer Protocol Data Units (PDUs) include a checksum, what is the checksum value in the above frame (figure 2). [2 marks]

Student has to calculate size of MAC, IP, and UDP Headers. The UDP Checksum is 08a0 in hexadecimal (2208 in decimal).

2

(d) What is the purpose of this checksum? (In what ways does it differ from the Cyclic Redundancy Check (CRC) used at the link layer). [6 marks]

The UDP Checksum is a checksum to verify that the end to end data has not been corrupted by routers or bridges in the network or by the processing in an end system. If this check is not required, the value of 0x0000 is placed in this field, in which case the data is not checked by the receiver. At the final destination, the UDP protocol layer receives packets from the IP network layer. These are checked using the checksum (when enabled this checks correct end-to-end operation of the network service) and all invalid PDUs are discarded. The checksum includes all the user 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 may not perform this integrity check. Although this is allowed, it is strongly recommended that a checksum is used in every UDP PDU, since it provides an End-to-End reliability guarantee.

6