Session 2001-2002 Exam 1

EG/ES 3567 Worked Solutions.

Please note that both exams have identical solutions, however the level of detail expected in ES is less, and the questions are phrased to provide more guidance on how to provide the solution.

Dr Gorry Fairhurst

G.Fairhurst@eng.abdn.ac.uk

Mark

1. (a) Sketch a diagram showing each of the layers in the Open Systems Interconnection (OSI) Reference Model. Include the position of each protocol layer in the diagram. [6 marks]

4 Marks for correct layering; 2 for detail in the diagram

The two lowest layers operate between adjacent systems connected via the physical link and are said to work "hop by hop". The protocol control information is removed after each "hop" across a link (i.e. by each System) and a suitable new header added each time the information is sent on a subsequent hop. The network layer (layer 3) operates network-wide and is present in all systems and responsible for overall co-ordination of all systems along the communications path.

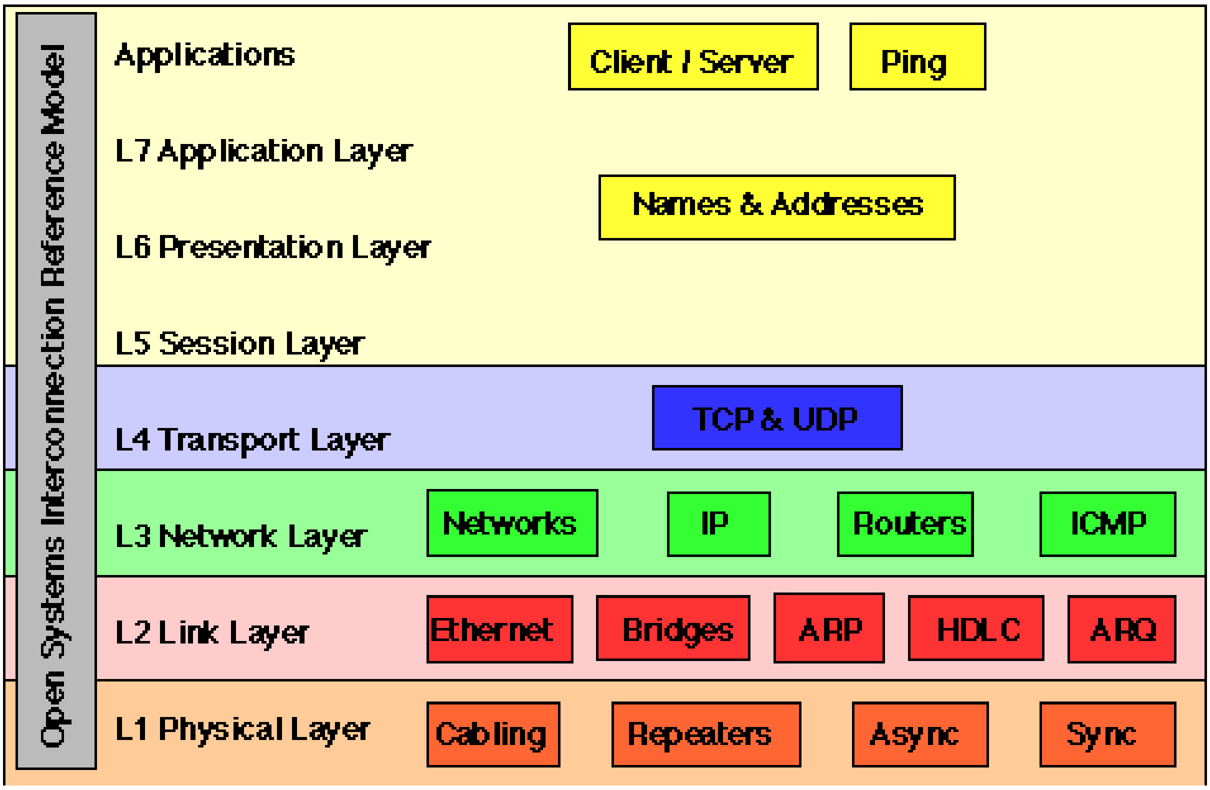

Example detail (text not required - but diagram must show key features) :

Physical layer: Provides electrical, functional, and procedural characteristics of thephysical links that transparently send the bit stream; only recognises individual bits, not characters or multicharacter frames.
Data link layer: Provides functional and procedural means to transfer data between network entities and correct/detect transmission errors; grouping of bits into bytes, synchronisation, media access control.
Network layer: Provides independence from link technology. Includes routing ; transfer data between end users.
The layers above layer 3 operate end-to-end and are only used in the End Systems (ES) which are communicating. The Layer 4 - 7 protocol control information is therefore unchanged by the IS in the network and is delivered to the corresponding ES in its original form. Layers 4-7 (if present) in Intermediate Systems (IS) play no part in the end-to-end communication.

6

Question Number	1	Solution	Page of 12

Mark

(b) The traceroute program may be used to determine an end-to-end Internet Path through a network. Explain (using appropriate diagrams) the packets that are exchanged when it used over the above path (figure 1). [8 marks]

<u>Marks summary: ICMP Echo request message; TTL set to 1; Router decrements TTL and discards; ICMP error message return (TTL exceeded); sender repeats this uses larger TTL; Next hop returns ICMP message; repeats; until final host accepts messgae and returns echo-response.</u>

Traceroute uses ICMP echo messages. These are addressed to trhe target IP address. The sender manipulates the TTL (hop count) value at the IP layer to force each hop in turn to return an error message.

It starts with a TTL of 1
The switches passes the packe the un-modified.
The router Y receives the packet (it is on the path to Z)
It decrements the packet TTL
This reduces to zero, the router generates an error message and returns this to the sender (W)

```
Client W ----ICMP echo src= W, dst=Z, TTL=1------> Router Y
                                                      |
 Client W <---ICMP error src= Y, dst=W, TTL=64---- Router Y
```

The client receives the ICMP error message and notes that Y is one hop away on the path to Z.

It then probes by sending the same packet with a TTL of 2.
The router Y receives the packet (it is on the path to Z)
It decrements the packet TTL
This is greater than zero, the router forwards this along the path to the destination.

```
Client W ----ICMP echo src= W, dst=Z, TTL=2------> Router Y
                                                      +--ICMP echo src= W, dst=Z, TTL=1-> Z

                                                      <-ICMP error src=  Z, dst=W, TTL=64-- Z
                                                      |
Client W <---ICMP error src= Y, dst=W, TTL=63---- Router Y
```

8 **(c) What is the difference between a Flat and a Hierarchical address organisation?Which organisation is used for Ethernet Medium Access Control addresses? [6 marks]**

Marks Summary: 2marks for clear description of: Hierarchy; Flat; and 2 for MAC=Flat

Flat = unstructurured - e.g. an assigned ID and a serial number. - The number itself only indicates the assigned address block, and odes not indicate anything about thelocation or operation of the equipment to which itis assigned.

Hierarchical = organised in a tree extending down wards from a root. (Suitable diagram may be appropriate)

6 Ethernet MAC uses a flat address sceheme with OUIs assigned by the IEEE on payment of a fee.

Mark

2 (a) The Ethernet Local Area Network (LAN) uses Carrier Sense Multiple Access with Collision Detection (CSMA/CD) to share the transmission medium. Define the following terms:

(i) Carrier Sense [4 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 (labelled CS in the figure).

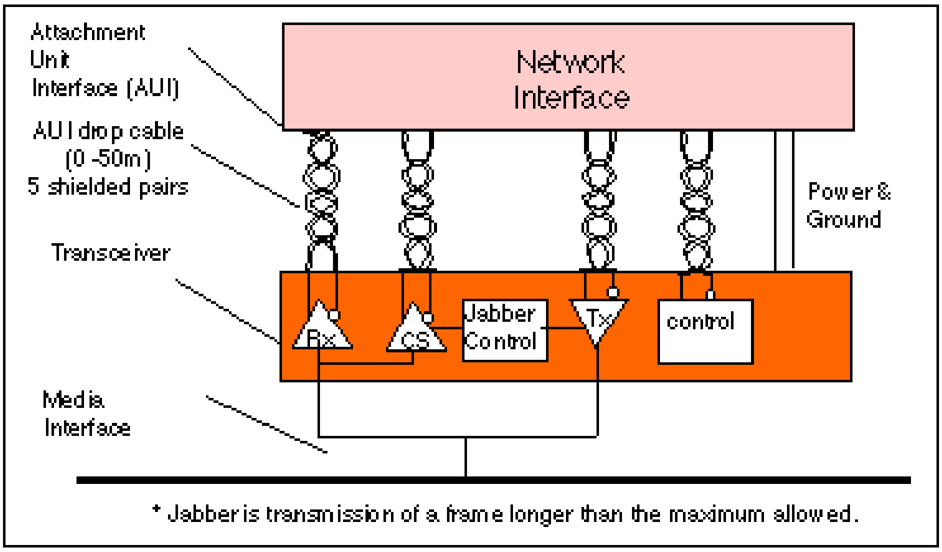

The individual bits are sent by encoding them with a 10 (or 100 MHz for fast Ethernet) clock using Manchester encoding. Data is only sent when no carrier is observed (i.e. no current present) and the physical medium is therefore idle.

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.

4

(ii) Collision Detection [4 marks]

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. 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.

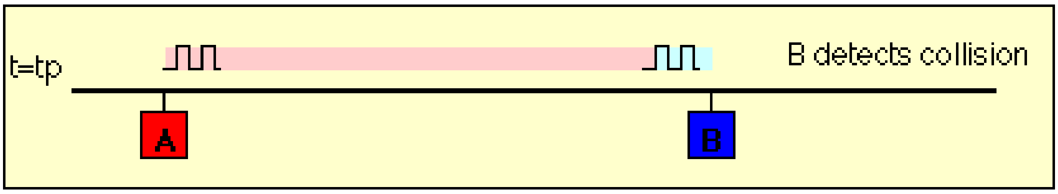

To ensure that no node may completely receive a frame before the transmitting node has finished sending it, Ethernet defines a minimum frame size (i.e. no frame may have less than 46 bytes of payload). The minimum frame size is related to the distance which the network spans, the type of media being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

Mark

fine 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.

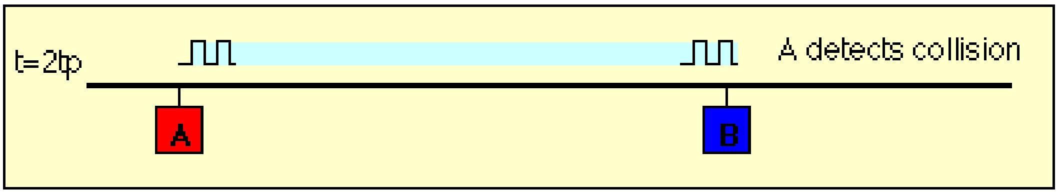

A short time later, computer B also transmits. (In this case, the medium, as observed by the computer at B happens to be idle too). After a period, equal to the propagation delay of the network, the computer B detects the other transmission from A, and is aware of a collision, but computer A has not 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.

4

(iii) Collision Domain [4 marks]

Traditional Ethernet uses a bus architecture in which all the computers connected to the cable share the capacity of the medium using CSMA/CD. In practice, most Ethernet networks employ hubs and repeaters, but these do not change thebasic rules of sharing. A network of repeaters and hubs is therefore called a "Shared Ethernet" or a "Collision Domain". The various systems sharing the Ethernet all compete for access using the CSMA/CD access protocol. This means that only one system is allowed to transmit within the Collision Domain at any one time. Each system has to share a proportion of the available network bandwidth.

Students should mention bridges/switches/routers separate collision domains. In contrast, the use of bridges, switches and routers separates each cable segment into an independent colission domain.

4

(b) Describe the phenomenon of Ethernet Capture [5 marks]

A drawback of sharing a medium using CSMA/CD, is that the sharing is not necessarily fair. When each node connected to the LAN has little data to send, the network exhibits almost equal access time for each node. However, if one node starts sending an excessive number of packets, it may dominate the network. Such conditions may occur, for instance, when one node in a LAN acts as a source of high quality packetised video. The effect is known as "Ethernet Capture".

Computer A dominates computer B. Originally both computers have data to transmit. A transmits first. A and B then both simultaneously try to transmit. B picks a larger retransmission interval than A (shown in red) and defers. A sends, then sends again. There is a short pause, and then both A and B attempt to resume transmission. A and B both back-off, however, since B was already in back-off (it failed to retransmit), it chooses from a larger range of back-off times (using the exponential back-off algorithm). A is teherfore more likely to succeed, which it does in the example. The next pause in transmission, A and B both attempt to send, however, since this fails in this case, B further increases its back-off and is now unable to fairly compete with A.

A similar situation may arise when many sources compete with one source which has much more data to send. Under these situations some nodes may be "locked out" of using the medium for a period of time.
The use of full duplex cabling or higher speed transmission (e.g. 100 Mbps Ethernet) eliminates this problem.

Question Number	1	Solution	Page of 12

Mark

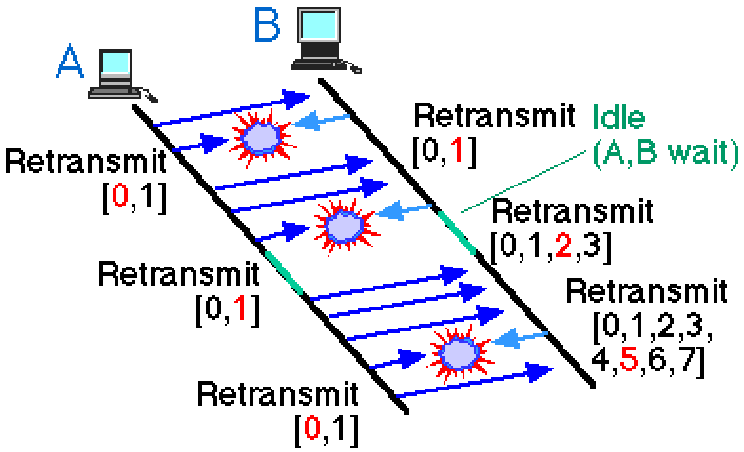

5

(c) A session using the User Datagram Protocol (UDP). It sends a series of packets over an Ethernet LAN. The payload of each UDP packet has a size of 530 B. Determine the size of the Ethernet frame using the information provided in the PDU header chart.

First determine theprotocol headers which contribute to the PDU size:

Preamble (8B) + MAC Header (14 B) + IP Header (20 B) + UDP(8 B) + UDP Payload (530 B) + CRC-32 (4 B)

8+14+20+8+530+4= 584 B (excluding the IFG)

3

Question Number	3	Solution	Page of 12

Mark

(a) Some protocols are said to provide a "reliable" service. What guarantees must a reliable protocol offer? [4 marks]

Reliable delivery has been succinctly defined as "Data is accepted at one end of a link in the same order as was 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) No corruption of the content

A frame must also be delivered within a reasonable period.

4

(b) The Trivial File Transfer Protocol (TFTP) may be used to provide a reliable service over an IP network. Explain in detail (using appropriate diagrams) how TFTP may receiver from missing IP packets. [8 marks]

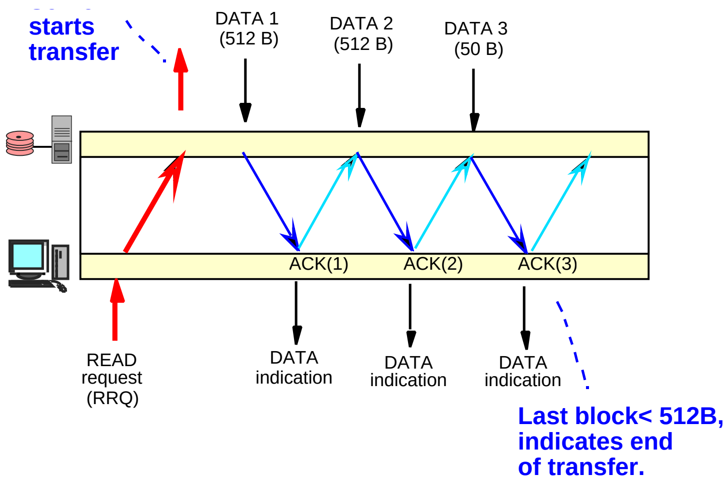

This is a very simple protocol to allow a client (often one being bootstrapped) to either get or put a file of data the protocol uses a stop and wait algorithm. The usual mode is to get a file by sending a read request (RRQ). The sender responds by sendingthe first data block. It numbers each block in turn, and when the block is received returns an acknowledgment. If a block is not received within a fixed period of time, a tiemr expires and the block is re-transmitted. The end of the file is signalled by reception of an incomplete (not full) block. The thrughput of tftp is limited becuase no window is used, and therefore over a long delay path, the protocol can work very very slowly.

The blue arrows show the sequence of data PDUs being sent across the link from the sender (top to the receiver (bottom). A Stop and Wait protocol relies on two way transmission to allow the receiver at the remote node to return PDUs acknowledging the successful transmission. The acknowledgements are shown in green in the diagram, and flow back to the original sender. A small processing delay may be introduced between reception of the last byte of a Data PDU and the corresponding ACK.

When PDUs are lost, the receiver will not be able to identify the loss. The transmitter must then rely upon a timer to detect the lack of a response. In the diagram, the second PDU of Data is corrupted during transmission. The sender is unaware of the packet loss, but starts a timer after sending each PDU. Normally an ACK PDU is received before this the timer expires. In this case no ACK is received, and the timer counts down to zero and triggers retransmission of the same PDU by the sender. The sender always starts a timer following transmission, but in the second transmission receives an ACK PDU before the timer expires, finally indicating that the data has now been received by the remote node.

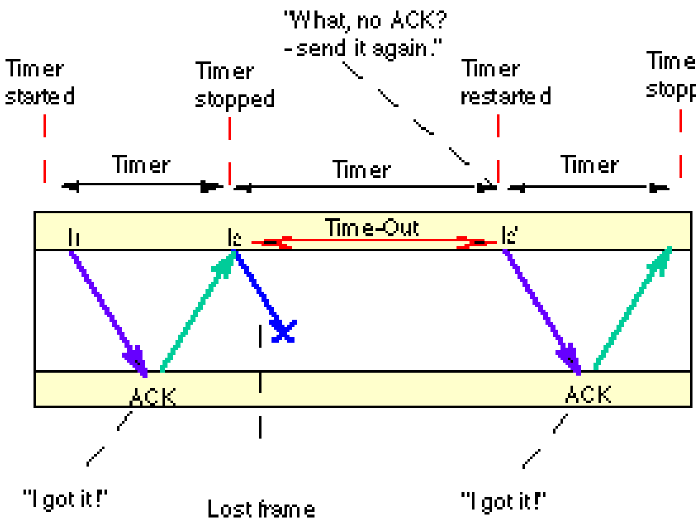

Mark

(c) Define the term Throughput. [2 marks]

Throughput is the number of bytes transferred per second by a protocol layert using the services of thelayer below.

It does not include protocol header information added by the layer itself of thelayers below. It is usally measured in bits per second.

2

(d) An end system sends 50 packets per second using the User Datagram Protocol (UDP) over a full duplex 100 Mbps Ethernet LAN connection. Each packet consists 1500B of Ethernet frame payload data. What is the throughput, when measured at the UDP layer? [6 marks]

Frame Size = 1500B

Packet has the following headers:

IP header (20B)
UDP header (8B)

Total header in each packet = 28B

Total UDP payload data is therefore 1500-28 = 1472B.

Total bits sent per second = 1472 x 8 x 50 = 588800 bps

6 or **588 kbps.**

Question Number	4	Solution	Page of 12

Mark

(a) The waveform above (voltage v. time) was observed on an oscilloscope when a byte was transmitted along an coaxial Ethernet cable. What is the value of the byte that was sent? [4 marks]

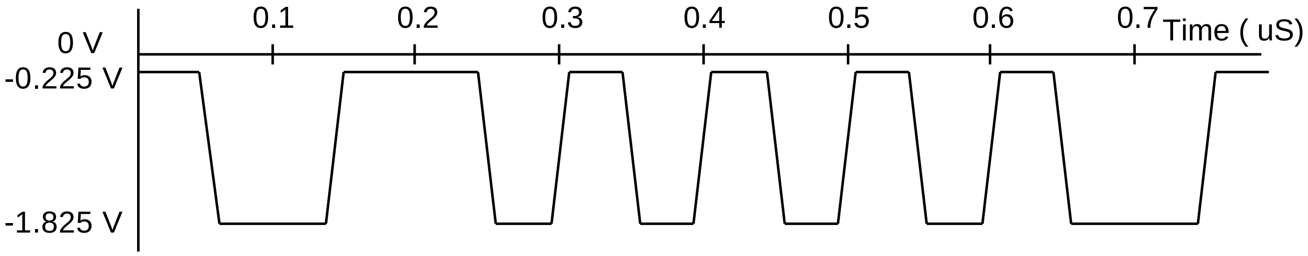

Manchester decoding:
Note transition at centre of each bit. (note also inversiuon of the signal).
1->0 == 0
0->1 == 1

Hence reading right toleft (Lsb first) <- 1 Mark for this.

4 1011 1110 = 0x7D

(b) What is the purpose of a Pre-amble and why is it sometimes needed for synchronous communications? [5 marks]

The purpose of the idle time before transmission starts is to allow a small time interval for the receiver electronics in each of the nodes to settle after completion of the previous frame. A node starts transmission by sending an 8 byte (64 bit) preamble sequence. This consists of 62 alternating 1's and 0's followed by the pattern 11. When encoded using Manchester encoding, the 62 alternating bits produce a 10 MHz square wave. The purpose of the preamble is to allow time for the receiver in each node to achieve lock of the receiver Digital Phase Lock Loop which is used to synchronise the receive data clock to the transmit data clock. At the point when the first bit of the preamble is received, each receiver may be in an arbitary state (i.e. have an arbitary phase for its local clock). During the course of the preamble it learns the correct phase, but in so doing it may miss (or gain) a number of bits. A special pattern (11), known as the start of frame delimiter, is therefore used to mark the last two bits of the preamble. When this is received, the Ethernet receive interface starts collecting the bits into bytes for processing by the MAC layer.

5

(c) Describe the role of the 2 byte Ethernet Frame Type field located in the Ethernet frame header, give two examples of protocols that use this field. [4 marks]

This is a service access point at thelink layer. (SAP)
This is used by the receiver to identify the type of payload being transported, and hence thenetwork layert interface to which the received packet is to be passed.

Two protocols which rely on this value are:

(i) IP - the network layer protocol of the Internet
(ii) arp - the address resolution protocol.
4 (IPv6 is another example, but not covered by the course)

Question Number		Solution	Page of 12

Mark

d) Explain the reason why there can be no more than 5 segments in series on the longest path between two systems using an Ethernet Local Area Network built using hubs. [5 marks]

The Ethernet Slot time is the maximum round-trip edge-to-edge delay that the CSMA/CD algorithm can accommodate.

Ethernet specifies this as 51.2 microseconds.

Two factors contribute to the slot time:
(i) The propagtion time of the signal in the physical medium
(ii) The processing delayof repeaters and hubs. (latency)

The latter is most significant, and standard components allow only 4 such devices along the longest path. Four devices - lead to the limit of five cable segments.

Note: the error conditions determine segment length, but isnce repeater is regenerative, this does not determine teh numebr of in-series segments. Marks may be given for understanding this.

4

e) Given that the Ethernet CRC-32 protects the integrity of frames sent across a Local Area Network, why does a transport protocol also include a checksum? [2 marks]

The link layer CRC protects the frame from corruption while being transmitted over the physical mediuym (cable). The CRC is removed by routers - as partr of the processing. A new CRC is added if the packet is forwarded by the router on another Ethernet link. While the packet is being processed by the router the packet data is protected by the CRC. Router processing errors may otherwise pass undetected.

The transport layer CRC therefore provides an end-to-end integrity check to ensure correctness of the data transferred.

<u>The main purpose is to detect problems that may arise in Inetermediate Systems (where there is no CRC on the data.</u>

2

Question Number	5	Solution	Page of 12

Mark

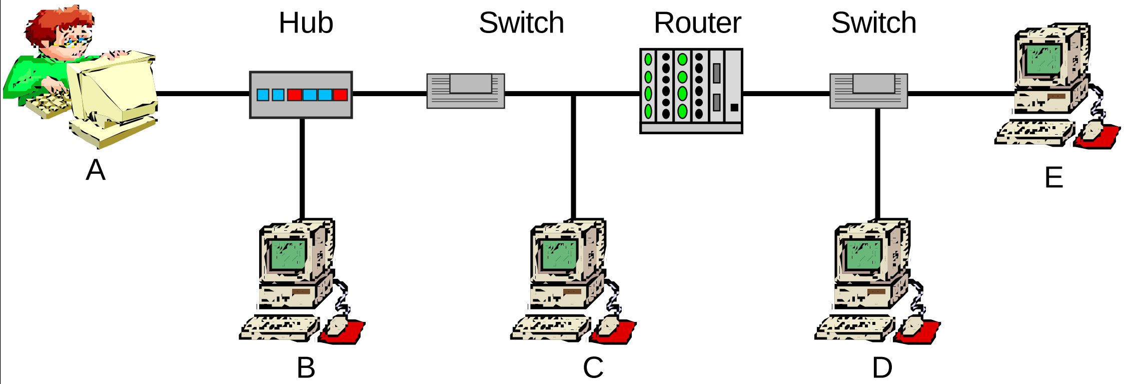

a) The end system A uses the Transmission Control Protocol (TCP) to send a packet to end system D with a payload of 100B, sketch the Ethernet frame that is sent. Ensure you show the addresses at both the MAC and IP layers. [6 marks]

Preamble (8B)
MAC Header (14 B)
  dst= D-mac address
  src = Router mac address <---<u>2 Marks</u> important detail
  type = 0x800 (IP)
IP Header (20 B)
  src = A IP address
  dst = D IP address
TCP Header (20 B)
DATA (100 B)
CRC-32 (4B)
6

b) An Internet Protocol packet is broadcast by B. Which End Systems receive this? [2 marks]

B is in the left broadcast domain
The frame is received by A,B,C

The right broadcast domain does not receive the frame because the router does not forward it.
The frame is NOT received by D,E
2

c) Outline the process used by A to determine the largest Maximum Transfer Unit MTU) supported by the router on the internet path to E. [4 marks]

<u>Look for: 1500B initial MTU; DF-bit set; ICMP message (with size); Refragment (still DF set); Cache result for a sepcific route; Try again in 10 (or so) minutes (in case route changes)</u>

A sends a maximum sized link frame (1500 B for Ethernet)
A sets theDOn'tFRagment boit in the IP header.

If Router has a lower MTU set, then router would be forced to fragment to the smaller MTU.
Instead router returns an ICMP error message indicating the smaller MTU value.

A receives the ICMP message, and caches this as the PATH MTU.
It starts a tiemr to reccheck-the PATH MTU of the route in the future.
While the timer is running it re-fragments this and all following packets to the new size.
4

Question Number	5	Solution	Page of 12

Mark

d) Explain the operation of a Switch, and describe how it performs Learning. [8 marks]

Summary:
MAC Sources address observed for learning
Associated with a port in the address table
MAC Destination address observed for forwarding
Learned addreses -> forward only to specified port
Discard frames to own address
Flood frames with unkonwn addresses to all ports
Aging required and re-learning when computers change the port they are connected to

A bridge or switch (as it is now more commonly called) is a LAN interconnection device which may be used to join two LAN segments, 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.

A bridge works within the data link layer (layer 2) of the OSI reference model. The format of PDUs at this layer in a LAN is defined by the Ethernet frame format (also known as MAC - Medium Access Control) consists of two 6 byte addresses and a one byte protocol ID / length field. The address field allows a frame to be sent to single and groups of stations. The MAC protocol is responsible for access to the medium and for the diagnosis of failure in either the hardware or the cabling.

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. This is called “learning”. 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 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).