

Notes:

- (i) Candidates are permitted to use approved calculators
- (ii) Candidates are not permitted to use the Engineering Mathematics Handbook
- (iii) An information sheet of protocol headers is provided

Candidates should attempt THREE questions. All questions carry 20 marks.

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. [4 marks]

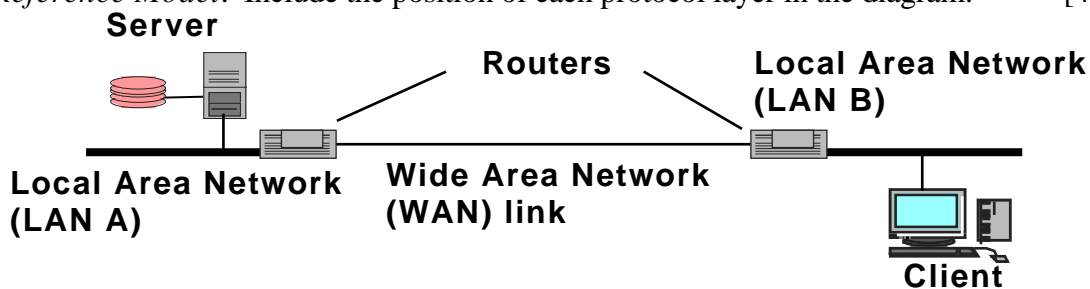


Figure 1: Two computers (server and client) connected via a network.

- (b) The *ping* program may be used to validate an end-to-end *Internet Path* through the above network (figure 1). Explain (using appropriate diagrams) the packets which are exchanged during a test. [8 marks]
- (c) In figure 1, the *Maximum Transfer Unit (MTU)* on the WAN link is 576 B. Explain how *Path MTU (PMTU)* discovery may be used by end systems on LAN A to discover a maximum packet size to send to end systems connected to LAN B. [8 marks]
2. (a) Describe the Ethernet transmit process, and explain the algorithm used to ensure retransmission following a collision in the transmission medium. [10 marks]

```

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

Figure 2: Transmitted Ethernet Frame

- (b) Figure 2 shows the hexadecimal dump of a packet sent using an Ethernet interface by a computer. What is the computer's own hardware address? [2 marks]
- (c) What is the value of the IP header checksum shown in figure 2? [2 marks]
- (d) Explain why checksums and *Cyclic Redundancy Checks (CRCs)* are applied at a number of protocol levels in a typical packet? [4 marks]
- (e) What can cause the loss of a packet by a network router? [2 marks]

continued over

3. (a) Some protocols are said to provide a “reliable” service. What guarantees must a reliable protocol offer? [4 marks]
- (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 recover from missing IP packets. [8 marks]
- (c) Define the term “Throughput” [2 marks]
- (d) An end system sends 50 packets per second using the *User Datagram Protocol (UDP)* over a 10B2 Ethernet LAN. Each packet consists 1500B of MAC payload data. What is the throughput, when measured at the UDP layer? [6 marks]
4. (a) Before an end system may communicate over a Local Area Network (LAN) it must first perform name resolution, and hardware address resolution. Explain the frames / packets exchanged when performing:
 - (i) Name resolution [6 marks]
 - (ii) Hardware address resolution [6 marks]
- (b) An end system uses the *Transmission Control Protocol (TCP)* transport protocol to communicate with another system a 10 Mbps LAN. The sender transmits 10 packets per second with 120 B of TCP data, and receives 5 packets per second of TCP protocol control information with no transport layer data. Calculate the utilisation of the network. [8 marks]
5. (a) Which of the following cable technologies may be used to support a 100 Mbps Ethernet Local Area Network?
 - (i) Coaxial cable
 - (ii) Fibre optic cable
 - (iii) Unshielded twisted pair cable [3 marks]

(b) Consider the following network shown in figure 3:


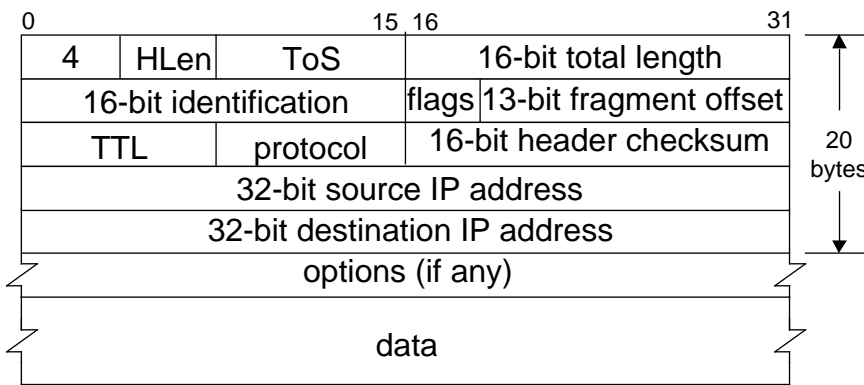


Figure 3: A network connecting 5 End Systems (V,W,X,Y,Z).

- (i) If W sends an IP packet to X which systems receive this packet at their physical interface? [3 marks]
- (ii) If X sends an IP packet to Y, which system’s *Medium Access Control (MAC)* address will be inserted in the source address field of the packet which is received by Y. [2 marks]
- (iii) If V sends an IP broadcast packet, which *End Systems* will receive it? [2 marks]
- (iv) Explain how the switch determines whether to forward a packet received from Z and destined for V. [5 marks]
- (c) Explain how the router determines whether to forward a packet received from Z and destined for V. [5 marks]

PDU Header Chart



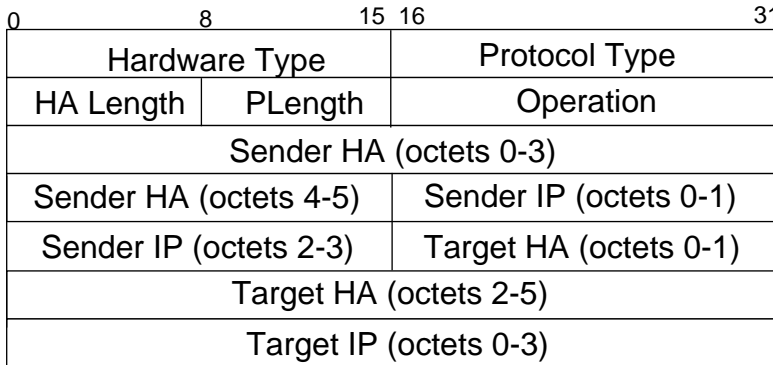
IP Protocol Types

0	IP
1	ICMP
2	IGMP
6	TCP
17	UDP

Flags

--X	More
-X-	Don't Fragment
X--	Unused


Internet Protocol Datagram (Ethernet Type = 0x800)



Operation ARP Message

1	ARP request
2	ARP reply
3	RARP request
4	RARP reply


ARP / RARP Packet (Ethernet Type = 0x806)



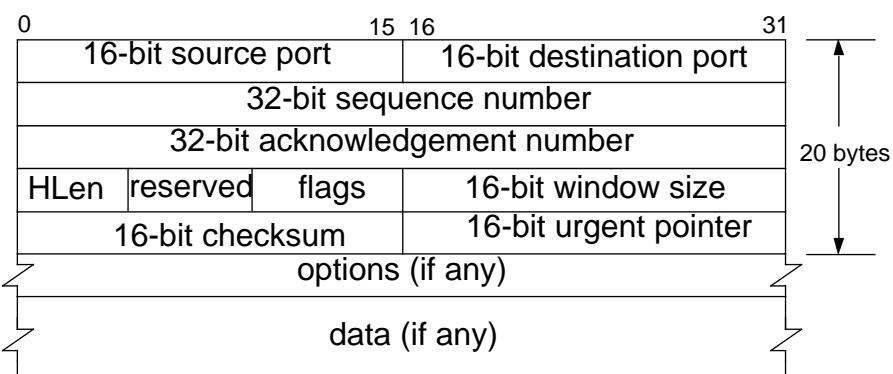
ICMP Message Type

0	Echo reply
3	Destination unreachable
4	Source quench
5	Redirect
8	Echo request

ICMP Message



UDP Packet



Well-Known TCP Server Ports

Port (decimal)	Service
23	Telnet
25	Mail
69	TFTP
8	WWW (http)

TCP Packet

Question Number	Solution	Page of 12
Mark	<h1>Worked solutions for EG/ES 3567</h1> <h2>2000 / 2001</h2> <p>Please note:</p> <p>ES paper questions simplified (some hints are given for the answers), and marking will be less strict according to normal marking of ES papers.</p> <p>A common worked solutions is provided for both EG & ES papers.</p>	

Question Number	1	Solution	Page of 12
-----------------	---	----------	------------

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. [4 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 (Students should draw a diagram):

- Physical layer (bottom layer)
- Lnk layer
- Network layer
- Transport layer
- Session layer
- Presentation layer
- Application layer (top layer)

4



The ping program may be used to validate an end-to-end Internet Path through the above network (in figure 1). Explain (using appropriate diagrams) the packets which are exchanged during a test.

The ping program allows a client to generate ICMP echo request messages which are encapsulated in IP datagrams. Each message contains an 8-bit type code which identifies the types of message. In this case two types of message are involved the ECHO request (sent by the client) and the ECHO reply (the response by the server). Each message may contain some optional data. When data are sent by a server, the server returns the data in the reply which is generated. ICMP packets are encapsulated in IP for transmission across an internet.

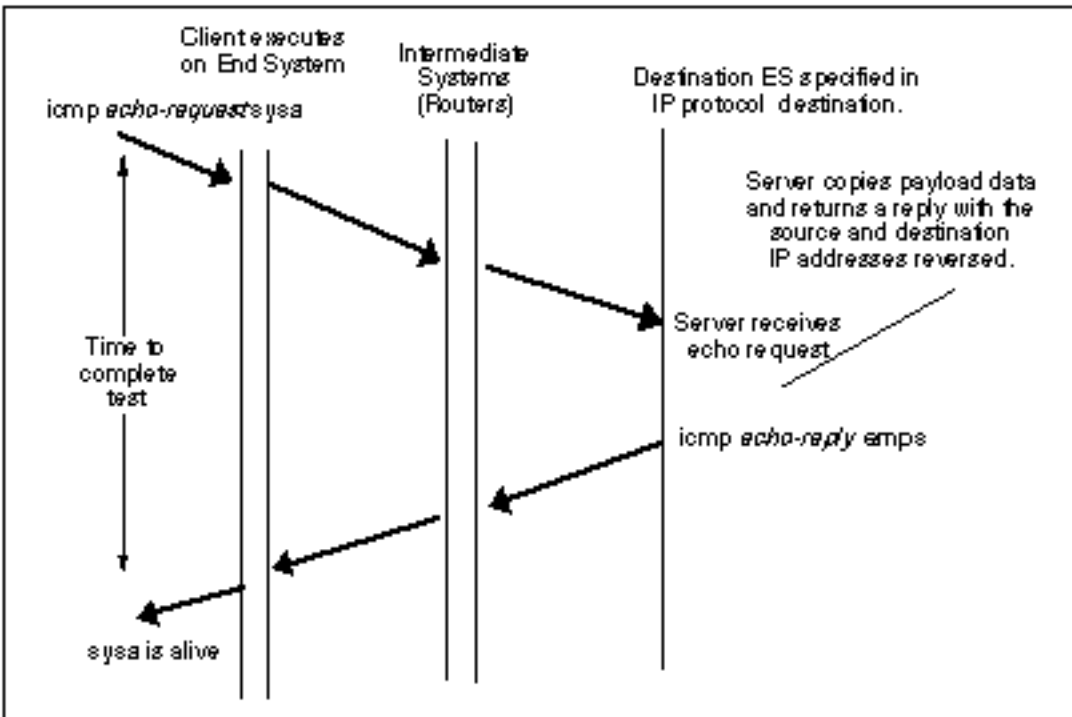
Each system in the Internet operates an ICMP echo server which, which responds to ICMP echo requests by generating an ICMP echo reply back to the originator. The returned message has an IP header with the source and destination addresses reversed and carries the same payload as originally sent. By including a

Question Number	1	Solution	Page of 12
-----------------	---	----------	------------

Mark

sequence number and timing the response, a client can determine whether the network path is working and also measure the current round trip delay and packet loss rate.

The "ping" program contains a client interface to ICMP. This may be used by a user to verify an end to end connection is working using the ICMP ECHO REQUEST/REPLY messages. The -s option of "ping" also collects some performance statistics (i.e. the measured round trip time and the number of times the remote server fails to reply. Each time an echo reply packet is received a single line of text is displayed. Each echo request packet contains a sequence number (starting at 0) which is incremented after each transmission, and a timestamp value indicating the transmission time. The text printed by ping shows the received sequence number, and the measured round trip time (in milliseconds).



Example transition diagram: Use of the ping program to test whether the computer "sysa" is operational.

The operation of ICMP is illustrated in the frame transition diagram shown above. In this case there is only one Intermediate System (IS) (router).

8

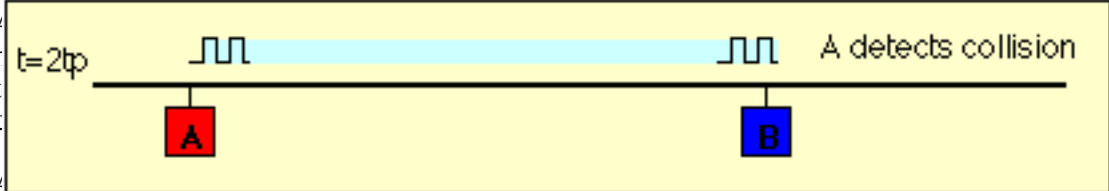
(c) In figure 1, the Maximum Transfer Unit (MTU) on the WAN link is 576 B. Explain how Path MTU (PMTU) discovery may be used by end systems on LAN A to discover a maximum packet size to send to end systems connected to LAN B.

Both A and B will have an MTU of 1500B, if they were Ethernet LANs. This is the largest size of IP packet which may be carried over an Ethernet LAN.

With Path-MTU Discovery the sender originally sends a full-sized packet (i.e. with the maximum size dictated 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.

In this case, the sender is directly connected to an Ethernet LAN and must therefore have a MTU of 1500B, the largest frame payload size allowed in an Ethernet LAN.

Question Number	1	Solution	Page of 12
Mark		<p>A router along the transmission path (in this case the first router) 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 (576 B in this case).</p> <p>The second router along the path now receives packets with a maximum size of 576B, and has an MTU of 1500B. Since $576 < 1500$ this router does not need to fragment.</p> <p>Note: reassembly ONLY occurs by the end system which receives the fragments - i.e., the end system that has the address specified in the IP destination address field of the packet's network layer header.</p> <p>When sender, receives an ICMP error message. It checks the type and the packet header which caused the error (this is returned in the payload field of the ICMP message). If the type field of the ICMP message indicates this packet was discarded because the packet would otherwise have required fragmentation, the sender now knows that packets sent to this destination IP address require fragmentation. The sender therefore reduces the Path-MTU size for the specified destination. The packet is then fragmented again by the sender. The fragments and all subsequent packets are fragmented according to the Path-MTU size specified in the ICMP message received. These still carry the Don't Fragment (DF) flag.</p> <p>If there are further routers along the path, with a smaller MTU, then they now receive the fragments with the discovered path MTU. If this now happens to be too large, then they also will discard the fragments and generate an ICMP message indicating the smaller MTU. In practice, this seldom happens more than once, but in theory it could happen many times, as each down-stream MTU size is discovered.</p> <p>A timer is started, so that some time later, the sender can generate an unfragmented full sized (1500B) packet to probe to see if the path MTU has increased. This would occur, if the path (i.e. routers being used) changed, and the packets no longer travel along the path with a small MTU.</p> <p>In this case, $PMTU = \text{Min}(1500, 576, 1500) = 576$</p>	
8		Appropriate diagrams are desirable, and will receive marks.	

Question Number	2	Solution	Page of 12
Mark		<p>(a) Describe the Ethernet transmit process, and explain the algorithm used to ensure re-transmission following a collision in the transmission medium. [10 marks]</p> <p>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 (loften abelled CS).</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>A diagram of the above is very desirable and will receive 2 marks, if provided.</p> <p>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.</p> <p>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.</p>  <p>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.</p> <p>Other appropriate other diagrams will also receive marks.</p>	
10			

Question Number	2	Solution	Page of 12
Mark		<p>Figure 2 for use by Parts b & c: 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</p> <p>Full decode (not required in student answer): ETHER: ---- Ether Header ---- ETHER: 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) ETHER: IP: ---- IP Header ---- IP: 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</p>	

Question Number	2	Solution	Page of 12
Mark		<p>(b) Figure 2 shows the hexadecimal dump of a packet sent using an Ethernet interface by a computer. What is the computer's own hardware address? [2 marks]</p> <p>0x00d0 bbf7 c6c0 <- This is the unique MAC address, copied from the PROM of the Sender's Ethernet NIC Student must show how this value was arrived at.</p> <p>(c) What is the value of the IP header checksum shown in figure 2? [2 marks]</p> <p>a1a9 (2 bytes) <- This is the integrity check applied (signature) applied to the network layer protocol header. Student must show how this value was arrived at.</p> <p>(d) Explain why checksums and Cyclic Redundancy Checks (CRCs) are applied at a number of protocol levels in a typical packet? [4 marks]</p> <p>A powerful method for detecting errors in the received data is by grouping the bytes of data into a block and calculating a Cyclic Redundancy Check (CRC). This is usually done by the data link protocol and calculated CRC is appended to the end of the data link layer frame.</p> <p>A CRC is calculated by performing a modulo 2 division of the data by a generator polynomial and recording the remainder after division. Ethernet uses a 4B or 32-bit CRC. The 32-bit CRC added at the end of the frame provides error detection in the case where line errors (or transmission collisions in Ethernet) result in corruption of the MAC frame. Any frame with an invalid CRC is discarded by the MAC receiver without further processing. The MAC protocol does not provide any indication that a frame has been discarded due to an invalid CRC.</p> $\text{CRC-32} = x^{32} + x^{26} + x^{16} + x^{12} + x^{11} + x^{10} + x^5 + x^4 + x^2 + x + 1$ <p>Note - it is useful for students to say that this division may be performed in software, it usually performed using a shift register and X-OR gates. The hardware solution for implementing a CRC is much simpler than a software approach.</p> <p>(e) What can cause the loss of a packet by a network router? [2 marks]</p> <p>A router provides no delivery guarantees and therefore provides only a <u>best effort</u> service.</p> <p>The router may fail to send (forward) a packet for various reasons. These include:</p> <ul style="list-style-type: none">(i) Hardware failure (e.g. a processor failure, RAM error, internal hardware, power failure)(ii) Software failure (e.g. a failure to interpret IP network packet options, or a fault in the software)(iii) Transmission error (e.g. a bit error on a WAN link, or excessive congestion in a LAN)(iv) A routing error (e.g. when there is no known route to a destination, N.B. this may be transient) <p>In addition, the following normal actions also lead to the packet not being forwarded:</p> <ul style="list-style-type: none">(i) When the Time To Live (TTL) value becomes zero (the packet is discarded)(ii) When a packet arrives which is too large to be forwarded over the next link., and the packet has the Don't Fragment (DF) bit set in its header.(iii) When the packet's destination address is the same as the router's own IP address (the packet was sent to the router itself)(iv) When the packet is an IP broadcast packet.(v) When the packet matches an entry in the router's filter list (i.e. is illegal) <p>Full list above is not required in this answer, a selection of 4 or more of the above should suffice to receive 2 marks, 1 mark for a list of 2.</p>	

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]

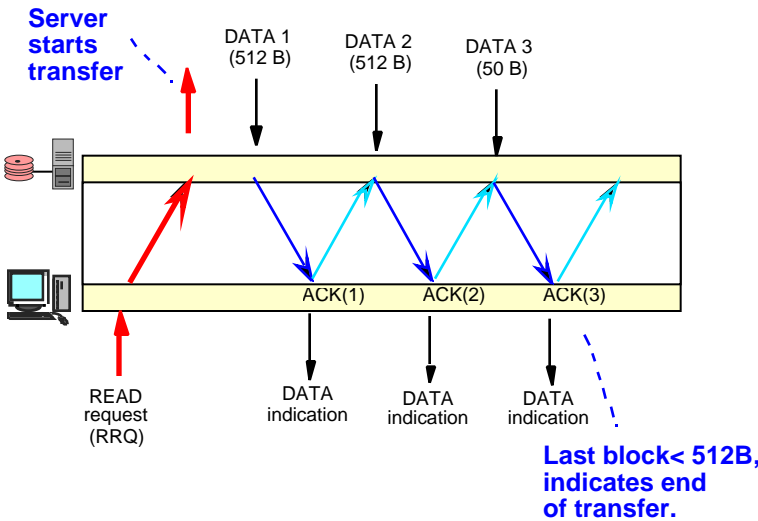
Reliable delivery has been succinctly defined as "Data is accepted at one end of a link in the same order as was transmitted at the other end, without loss and without duplicates." This implies four constraints:

- (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) A frame must be delivered within a reasonable period

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

4

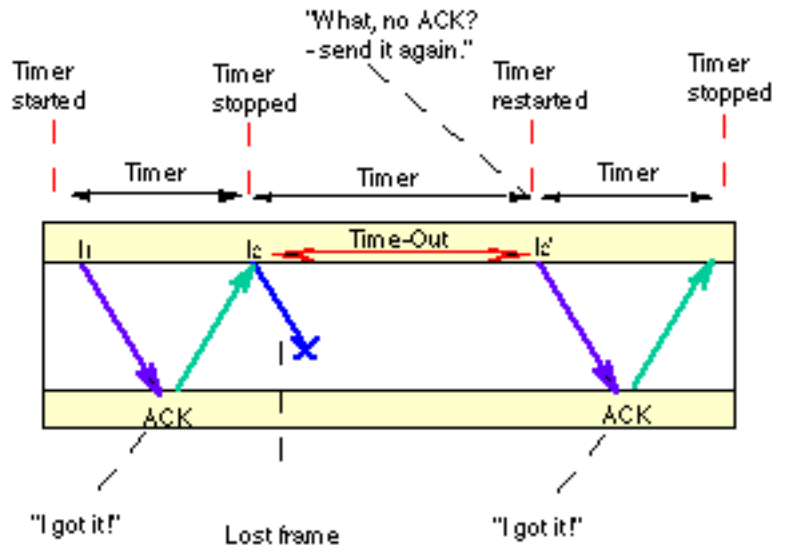
(b) The Trivial File Transfer Protocol (TFTP) may be used to provide a reliable service over an IP network. Explain in detail (using appropriate diagrams) how TFTP may receiver from missing IP packets. [8 marks]



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 sending the 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 timer expires and the block is retransmitted. The end of the file is signalled by reception of an incomplete (not full) block. The throughput of tftp is limited because 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 (full duplex or half duplex) 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 generation of the corresponding ACK.

When PDUs are lost, the receiver will not normally be able to identify the loss (most receivers will not receive anything, not even

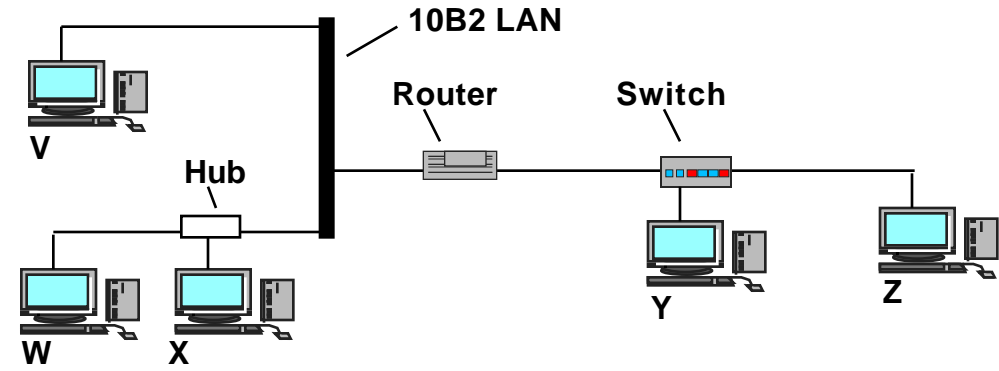


Question Number	4	Solution	Page of 12
Mark		<p>an indication that something has been corrupted). The transmitter must then rely upon a timer to detect the lack of a response.</p> <p>In the diagram, the second PDU of Data is corrupted during transmission. The link layer receiver discards the corrupted data (by noting that it is followed by an invalid data checksum). The sender is unaware of this 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.</p>	
8		<p>Appropriate diagrams may replace the above text, <i>if</i> they are explained and labelled. Key points to receive marks - How receiver knows a packet is missing; need for ACKs.; need for Timer; numbering of packets. 2 marks for each properly explained</p> <p>(c) An end system sends 50 packets per second using the User Datagram Protocol (UDP) over a 10B2 Ethernet LAN. Each packet consists 1500B of MAC payload data. What is the throughput, when measured at the UDP layer? [6 marks]</p> <p>Throughput is defined as " the number of bits per second transferred by a protocol layer using the service of the layer below". In this case the layer is the transport layer, and the protocol is UDP.</p> <p>Inclusion of a definition is desirable.</p> <p>Packet has the following headers:</p> <p>Ethernet / MAC Header MAC payload (1500B) Ethernet CRC-32</p> <p>The MAC payload is: IP header (20B) UDP header (8B) UDP payload.</p> <p>A sketch of the above is <i>desirable</i>.</p> <p>Total header at or above the IP layer in each packet = 28B (IP+UDP)</p> <p>Total UDP payload data is therefore 1500-28 = 1472B.</p> <p>Note 8 bits per byte. 50 packets per second.</p> <p>Total bits sent per second = 1472 x 8 x 50 = 588800 bps</p>	
6		<p>or 588 kbps.</p>	

Question Number	4	Solution	Page of 12
Mark		<p>(a) Before an end system may communicate over a Local Area Network (LAN) it must first perform name resolution, and hardware address resolution. Explain the frames / packets exchanged when performing:</p> <p>(i) Name resolution [6 marks]</p> <p>Once there were only a few computers connected to the first internet, called the ARPANET, at that time everyone knew each others IP address, so communication was easy, one simply typed the appropriate sequence of digits representing the IP number for each destination. After a while, the number of computers started to grow, and people began to forget the strange numeric IP numbers. So IP names came into being, and each computer held a table of names and their associated addresses, which had to be updated as new computers were connected to the network.</p> <p>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. 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 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.</p>	
6		<p>(ii) Hardware address resolution [6 marks]</p> <p>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 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.</p> <p>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 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. 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.</p> <p>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.</p> <p>To reduce the number of address resolution requests, a client normally caches resolved addresses</p>	
6		<p>Diagrams may be included in answer and will receive marks</p>	

Question Number	4	Solution	Page of 12
Mark		<p>(b) An end system uses the Transmission Control Protocol (TCP) to communicate with another end system. Both are connected to the same 10 Mbps Ethernet Local Area Network (LAN). The sender transmits 10 packets per second with 120 B of TCP data, and receives 5 packets per second of Acknowledgements (ACKs) with no data. Calculate the utilisation of the network.</p> <p>Students <i>should define utilisation</i>.</p> <p>Utilisation = total bits sent in one second / channel bit rate, expressed as a percentage.</p> <p>Note: This includes all overheads added to the frame as it is sent on the wire/fibre.</p> <p>TCP header = 20 B - taken from PDU Header Sheet IP header = 20B - taken from PDU Header Sheet Ethernet MAC Header = 14 B - students should know this. Ethernet MAC trailer = 32bit CRC - students should know this. Ethernet Preamble = 8B - students should know this.</p> <p>The students may ignore (or choose to include) the Inter Frame Gap . They SHOULD say that it is present, even if they ignore its effect on the calculation.</p> <p>Students <i>should sketch the above frame format</i>.</p> <p>120 B TCP payload, 20 B TCP header, 20B IP Header, 14 B MAC Header, 4 B Mac Trailer, 8 B preamble. Ignore IFG</p> <p>Transmit total frame size = $120+20+20+14+4+8 = 186$ B Given 8 bits per byte, total number of bits is 1488 b</p> <p>Receive total frame size = $60+14+4+8 = 86$ B = 688 b</p> <p>0 B TCP payload, 20 B TCP header, 20B IP Header, 14 B MAC Header, 4 B Mac Trailer, 8 B preamble.</p> <p>But note minimum Ethernet frame size is 60B payload. Ignore IFG</p> <p>Total bits sent per second = $(1488 \times 10) + (688 \times 5) = 14880 + 3440 = 14880 + 3440 = 18320$ bits</p> <p>Utilistaion = $(18320 \times 100) / (10\ 000\ 000) \% = \mathbf{0.18\%}$</p> <p>This is much less than 20%, and we can therefore ignore the impact of collisions.</p> <p>NOTE EG3567 students must note the following with their answer:</p> <p>Presence of Inter-frame gap Minimum frame size May ignore impact of collisions</p> <p>Above assumptions carry 2 marks. (not required for ES answer)</p>	

8

Question Number	5	Solution	Page of 12
Mark	3	<p>5. (a) Which of the following cable technologies may be used to support a 100 Mbps Ethernet Local Area Network? (i) Coaxial cable (ii) Fibre optic cable (iii) Unshielded twisted pair cable [3 marks]</p> <p>(i) is only suited to 10 Mbps operation (ii) and (iii) may be used at 10 Mbps and also at 100 Mbps and 1 Gbps.</p> <p>(b) Consider the following network shown in figure 3:</p>  <p>(i) If W sends an IP packet to X which systems receive this packet? [3 marks]</p> <p>V- receive & discard W - sender X - receiver Y - Z - Hub - Forward Router -receive & discard receive & discard (not for a remote IP network) Switch</p> <p>(ii) If X sends an IP packet to Y, which system's Medium Access Control (MAC) address will be inserted in the source address field of the packet which is received by Y. [2 marks]</p> <p>At X, X inserts its own address, but the router removes this and replaces it by the MAC interface address of the Router Ethernet interface connected to the switch.</p> <p>(iii) If V sends an IP broadcast packet, which End Systems will receive it? [2 marks]</p> <p>V- sender W - receiver X - receiver Y - Z - Router will not forward an IP broadcast packet (denoted by a host part of all zero, or all one bits)</p> <p>(iv) Explain how the switch determines whether to forward a packet received from Z and destined for V. [5 marks]</p> <p>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</p>	
3			
3			
2			
2			

Question Number	5	Solution	Page of 12
Mark	5	<p>two LANs and may enforce a security policy separating different work groups located on each of the LANs.</p> <p>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.</p> <p>The bridge learns which MAC 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 Y, Z, and the router are observed to be on each of the associated interface ports.</p> <p>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 packets, 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 override the normal forwarding by inserting entries in a filter table to inhibit forwarding between different workgroups (for example to provide security).</p> <p>To reach V, the computer inserts the router address in the destination address field. Packets with a source of Z and destination of the router are received and forwarded only to the interface port to the router is connected.</p> <p>A learning bridge identifies which addresses are remote and local by observing the MAC <i>source</i> address. Filtering based on destination MAC address and may provide security filtering</p>	
		5	<p>(c) Explain how the router determines whether to forward a packet received from Z and destined for V. [5 marks]</p> <p>A router consists of a computer with at least two network interface cards supporting the IP protocol. The router receives packets from each interface. Input packets are first checked for validity, and then the link layer header is removed, leaving just the IP network layer header.</p> <p>The router checks the IP header integrity by first checking the Time To Live (TTL) value. If this is less than 1 the packet is discarded, if it is greater, the TTL value in the header is decremented, and the IP header checksum value is updated to reflect the change. The router then inspects the IP destination address in each received packet header. When the TTL=0 the router may send an ICMP message back to the sender.</p> <p>The address is checked to see if the packet is sent to the router itself. If so, the router acts as an end-host and passes the received packet to the transport layer within the router.</p> <p>The address is first checked using the interface netmask (applying a logical and to the IP destination address and the interface IP address, and comparing the result). If the packet is local to the interface on which it was received (two matching network numbers) it is discarded.</p> <p>The router then decides if it knows how to forward the packet using the routing information held within the routing table. If an appropriate entry is found, or no appropriate entry, but a default route is found, the router then determines the destination output interface for the packet.</p> <p>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. Packets which are not filtered are therefore forwarded to an appropriate output interface</p> <p>Lectures in 2000 showed the algorithm used by the router - this may be drawn here.</p>