Worked solutions for EG/S 3567 2006/2007

UNIVERSITY OF ABERDEEN **SESSION 2006-07**

Degree Examination in EG 3567 Communications Engineering 1A
Degree Examination in ES 3567 Communications Engineering 1B

EG/ES 3567 Worked Solutions.

Please note that both exams have identical solutions, however the level of detail expected in the 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

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

1. (a) Provide a diagram that illustrates the seven layers of the *Open Systems Interconnection (OSI)* layers, and *Protocol Data Units (PDUs)* are sent using this model. [6 marks]

2 marks for the diagram showing OSI Layers, correctly labelled.

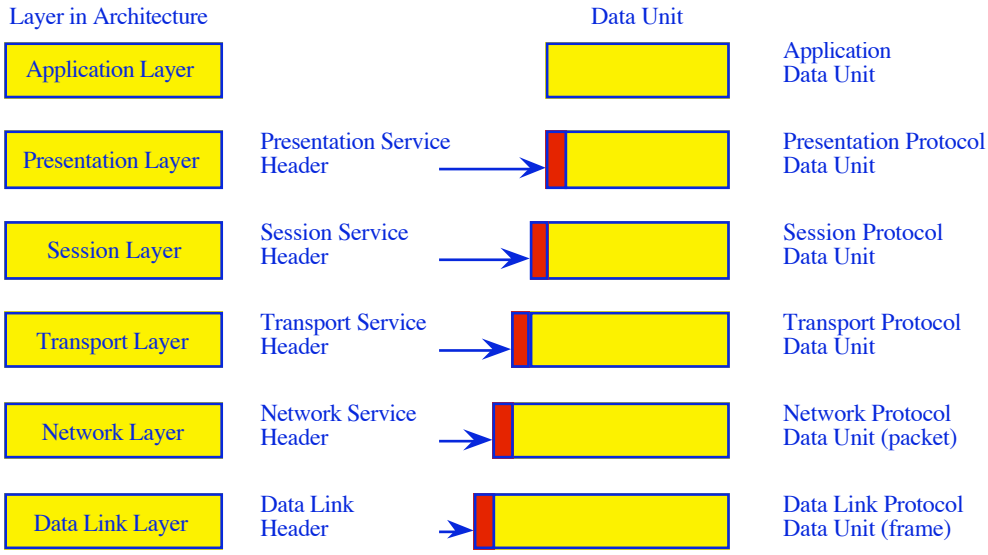

When data (PDU) is passed by the application process to a layer, the layer processes it. It then prefixes a protocol header (PCI) and passes the packet (together with the new header) to the layer below (as an SDU). This layer in turn processes the data and prefixes an additional new header. Each layer treats the assemblage of information from higher layers as data, and does not worry about its contents. This process continues until the packet reaches the Physical Layer, where they then become a sequence of bits that can be sent across a specific Physical Medium (cable, wireless link, etc).

The physical layer serialises the packet (i.e. converts it to a series of bits) and sends it across a cable or communications circuit to the destination. At the receiver, the remote system reassembles the series of bits to form a frame and forwards the frame for processing by the link layer. This removes the link layer header, and passes it (up) to the next layer. The processing continues layer by layer as the packet rises up the remote protocol stack, until finally the original packet data is sent to the remote application. The upper three layers, or *middleware* (application, presentation, and session) are primarily concerned with ensuring that information is delivered in correct and understandable form. The transport layer forms an interface between these groupings.

Key issues to be identfied - the addition of PCI to the PDU to form an SDU, and the transmission via the SAP to the layer below (sending) or above (receiving) - 4 marks.

(b) Describe the services provided by the *Transport Layer*. [2 marks]

The transport layer provides transparent transfer of data between systems, relieving the user from concern with providing reliable and cost effective data transfer; it provides end-to-end *control* and *information interchange* with the *quality of service* needed by the application program; it is the first true end-to-end layer. It is also responsible for reliability and congestion-control.

(c) Explain how *routers* can redirect traffic around a failed link. [6 marks]

A router forwards packets from one IP network to another IP network. The router uses the information held in the network layer header (i.e. IP header address, DSCP, perhaps source address) to decide whether to forward each received packet, and which network interface to use to send the packet (via the routing/forwarding table). Most packets are forwarded based on the packet's IP destination address, along with routing information held within the router in a routing table. The router also monitors the status of each out-going link, ensuring that it does not choose to use links that are

not operating (a routing protocol or link protocol may inform this decision). Routing packets received from other routers help a router discover alternate links that may be used to replace a failed link, allowing the router to divert traffic along a different network path. This enables the routing table to be updated to accurately reflect the current interner connectivity. In the absence of detailed routing information, systems may use a default route (common in edge/acess networks). Answer should identify that this action is performed at the Network Layer (L3). Could refer to hierarchy of networks or use of trace-route, and issues from changes in route (but not required).

(d) An Ethernet frame (represented below in hexadecimal) is recorded by a network monitor.

0x0000: 0001 0800 0604 0001 000a 95f8 81a1 0a0a
0x0010: 0a29 0000 0000 0000 0a0a 0a64

Explain how this frame may be decoded to show which protocols were used. [6 marks]

Students are expected to know and explain how to decode packets using the packet header chart provided (although will not need it in this specific example). The chart provides the format of the protocol control information (PCI) for a set of well-known protocols. Solution should explain decoding of PCI and SAP to determine next-layer PCI, etc

MAC-Layer dest address 0001 0800 0604 (unicast) <- should note it is unicast.

MAC-Layer src address 0001 000a 95f8

MAC-Layer type field (81a1) <- should note that it is not a standard network protocol (in fact it is a L2 protocol, but students are not expected to know this).

Length 27 (Solution could note padding was omitted, because too short).

Normally would provide diagram showing above.

Should explain how type field is used in general to identify L3 protocol (although not expected to utilise this information. May provide diagrams showing demultiplexing. Could use header chart to provide examples of payload formats that could be decoded, e.g. IP, ARP. (Much more complex decodes were performed in class).

2. (a) Describe the development of Ethernet which led to the IEEE Ethernet standard. [6 marks]

1972: Original design in Xerox, Palo Alto using 75 Ohm cable offering 2 Mbps, and used originally to share expensive printers between a workgroup.

1980: DIX standard emerged when Xerox was joined by Digital and Intel, leading to blue-book Ethernet (1982). This used 50 Ohm cable and worked at 10 Mbps. Bob Metcalfe's 3COM network interface cards followed, forming the basis for a succesful industry.

The 10B5 cabling system uses thick (low loss) coaxial cable which forms a shared bus. Up to 100 transceivers may be used to connect computers to the bus. The cable is difficult to install (due to its weight, large 0.5" diameter, and constraints on minimum curvature). It's key advantage is the extended transmission distance and the fact that it has good noise immunity. Modern 10B5 cabling is used for backbone connections and it is now fairly uncommon to find this type of cabling using to connect user's workstations. Since this type of cabling is relatively difficult to install, the cable is normally installed in a communications duct or along a corridor in an office building. An AUI drop cable is then run from the cable to an AUI port on the network interface card of each computer to be connected.

In 1985, the Institute of Electrical and Electronic Engineers (IEEE) in the United States of America, produced a series of standards for Local Area Networks (LANs) called the IEEE 802 standards, of which 802.3 specifies a series of standards for Ethernet. The official IEEE name for this cable is 10 Baseband5 (10B5), indicating that it is specified for baseband communications (i.e. not modulated) at 10 Mbps over distances up to 500m. Answers may also discuss following technologies, briding, repeaters, wireless, etc as a part of the IEEE standards work.

(b) What challenges were faced when Fast Ethernet was introduced over twisted pair cable technology and how were these overcome? [6 marks]

Key issues faced - Cross talk and restricted bandwidth, requiring new PHY technology.

100BASE-T is the provides 100 Mbit/s Ethernet in either half-duplex (using CSM/CD) or full-duplex forms. 100BASE-TX runs over two pairs of wires in category 5 unshielded twisted pair cable. Like 10BASE-T, the normal pairs are coloured orange and green pairs (using pins 1, 2, 3 and 6 of the RJ-45 connector). This cable has a bandwidth of less than 100 MHz.

A manchester encoded waveform would require 200-400 MHz of bandwidth, far in excess of that offered by the cable. A scheme using 4B5B binary encoding therefore generates a series of 0 and 1 bits clocked at 125 MHz; the 4B5B encoding provides DC equalisation and spectrum shaping. 4B5B works by mapping each group of four bits (a 1/2 of a byte) to one group of 5 bits.

Since there are (2^5) 32 possible combinations of 5 bits, and there are only (2^4) 16 combinations of 4 bits one half the patterns are unused. The chosen set of 16 5-bit patterns are those with the most transitions, this ensures clocking information is present in the signal (for locking the receiver DPLL). This results in a bandwidth increased of 25%.

Cross-Talk requirements / RF Emission led to the need for a scrambler. The data is finally sent as a 3-level physical waveform known as MLT-3. MLT-3 cycles through a set of voltage levels {-1, 0, +1}, to indicate a 1-bit. The signal stays the same when transmitting a 0 bit. It takes four 1 bits to generate a complete cycle, this the maximum fundamental frequency is reduced to one fourth of the baud rate.

This scheme of 45/5b with MLT-3 encoding leads to a waveform of 31.25 MHz, well within the specification for Unshielded Twisted Pair Cabling.

No attempt was made to support the older co-axial cable installations, instead these must be replaced by UTP or Fibre.

(c) Some modern Ethernet links support a larger Maximum Transmission Unit (MTU), how do end systems know whether they should use the larger MTU? [6 marks]

The MTU 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 is a mutually agreed value (i.e. both ends of a link agree to use the same specific value) for most WAN links. The size of MTU may vary greatly between different links

The Path MTU Discovery algorithm operates at the sender at the boundary of the Transport Layer and Network Layers of the OSI Reference Model, generating probe messages and responding to ICMP error reports. In intermediate Systems (IS), i.e. Routers, this operates in the Network Layer, returning CGMP error messages based on the link-layer configuration. The way in which the end system finds out this packet size, is to send a large packet (up to the MTU of the link to which it is connected). The packet is sent with the Don't Fragment (DF) flag set in the IP protocol header. If a router finds that the MTU of the next link exceeds the packet size, the DF flag tells the router not to segment the packet, but instead to discard the packet. An CGMP message is returned by the router (R1 in the example below) to the sender (H0), with a code saying the packet has been discarded, but IMPORTANTLY, also saying the reason and indicating the maximum MTU allowed (in this case the MTU of the link between R1 and R2).

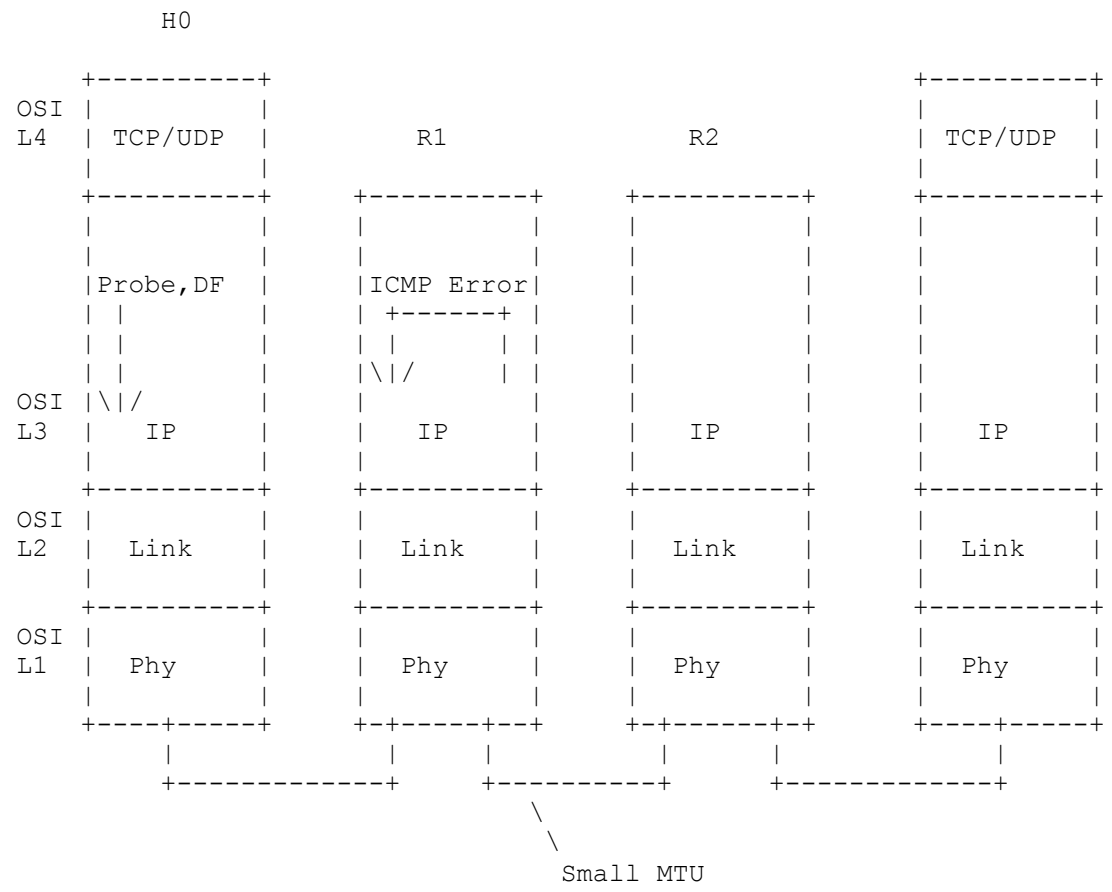

If the end system receives an CGMP message saying a packet is too large, it sets a variable called the PATH-MTU (P-MTU) to the appropriate maximum size and then itself fragments the packet to make sure it will not be discarded next time. The end system keeps a set of P-MTU values for each IP address in use.

When there are a series of links along the path, each with smaller MTU's, the above process may take place a number of times, before the sender finally determines the minimum value of the P-MTU. Once the P-MTU has been found, all packets are sent segmented to this new value. Routers do not therefore have to do any additional processing for these packets. Occasionally the end system will generate a large packet, just to see if a new Internet path has been found (i.e. a different route). The new path may allow a larger P-MTU.

(d) What is the purpose of the *Link Layer CRC*? [2 marks]

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.

CRC is used to detect transmission errors over the physical medium.
This is also used to verify correct bit-timing, i.e. to identify the DPLL has not missed/added bits.

3 (a) Provide a detailed description of how a computer sends a frame from an *Ethernet Network Interface Card (NIC)* connected to a shared Ethernet segment [10 marks]

The Ethernet network may be used to provide **shared** access by a group of attached nodes to the physical medium which connects the nodes. These nodes are said to form a Collision Domain. All frames sent on the medium are physically received by all receivers, however the Medium Access Control (MAC) header contains a MAC destination address which ensure only the specified destination actually forwards the received frame (the other computers all discard the frames which are not addressed to them).

Ethernet uses a refinement of ALOHA, known as Carrier Sense Multiple Access (CSMA), which improves performance when there is a higher medium utilisation. When a NIC has data to transmit, the NIC 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 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. Any NIC which does not need to transmit, listens to see if other NICs have started to transmit information to it.

However, this alone is unable to prevent **two NICs transmitting at the same time**. If two NICs 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 NIC is currently using the medium. In this case, both will then decide to transmit and a collision will occur. The collision will result in the corruption of the frame being sent (which will subsequently be discarded by the receiver since a corrupted Ethernet frame will (with a very high probability) not have a valid 32-bit MAC CRC at the end.

A second element to the Ethernet access protocol is therefore used to detect when a **collision** occurs. When there is data waiting to be sent, each transmitting NIC also monitors its own transmission. If it observes a collision (excess current above what it is generating, i.e. > 24 mA for coaxial Ethernet), 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 all NICs start to receive a frame before the transmitting NIC 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**, corresponding to 512 bit times at 10 Mbps.

After one complete round trip propagation time (twice the one way propagation delay), both NICs are aware of any collision. Finally the cable becomes idle.

An overview of the transmit procedure is given below. The transmitter initialises the number of transmissions of the current frame (n) to zero, and starts listening to the cable (using the carrier sense logic (CS) - e.g., by observing the Rx signal at transceiver to see if any bits are being sent). If the cable is not idle, it waits (defers) until the cable is idle. It then waits for a small Inter-Frame Gap (IFG) (e.g., 9.6 microseconds) to allow to time for all receiving nodes to return to prepare themselves for the next transmission.

Answers must explain the use of:

- The retry counter (upto 15 attempts) - noting possible loss.
- The backoff set (upto 2^10)
- The random retry interval (incl slot-time)

The preamble/DPLL function may be mentioned (1 additional mark)

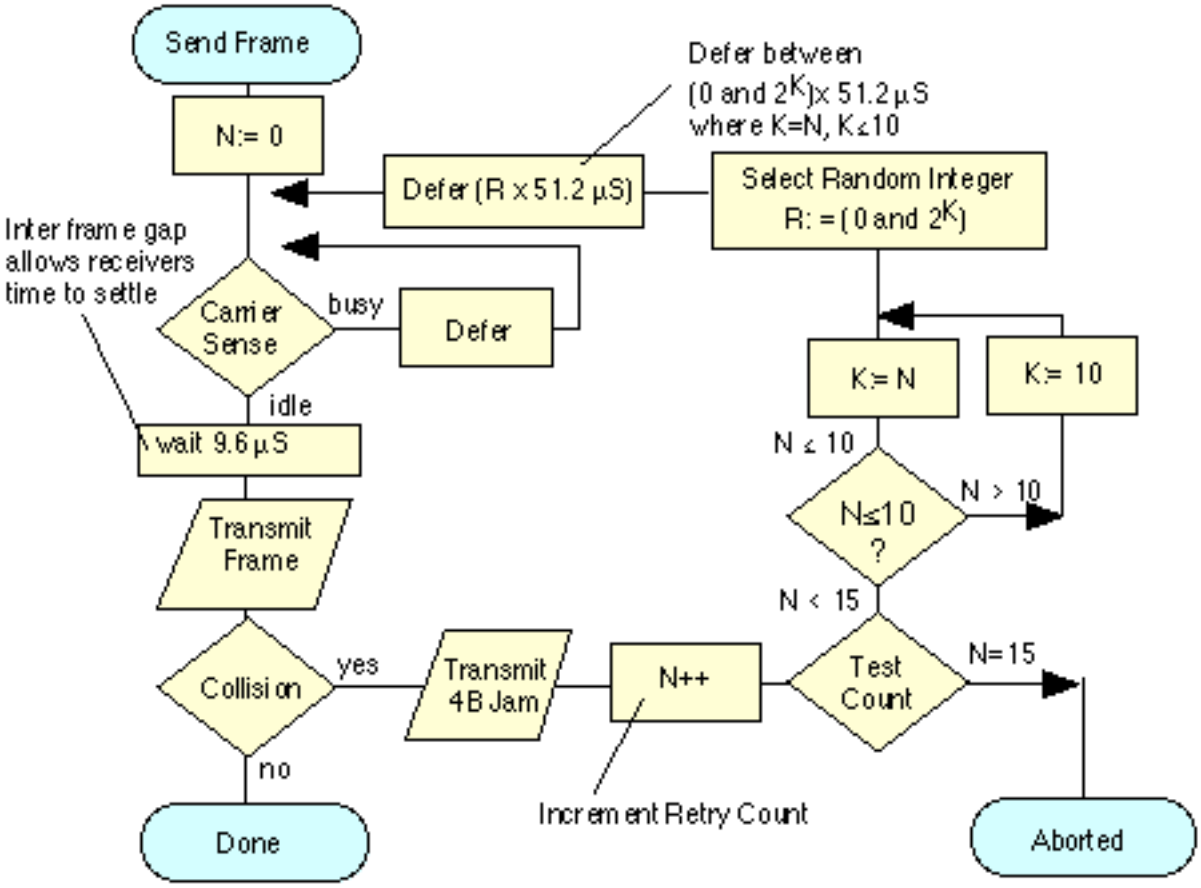

On a busy network, a retransmission may still collide with another retransmission (or possibly new frames being sent for the first time by another NIC). The protocol therefore counts the number of retransmission attempts (using a variable N in the above figure) and attempts to retransmit the same frame up to 15 times.

For each retransmission, the transmitter constructs a set of numbers:

{0, 1, 2, 3, 4, 5, ... L} where L is ([2 to the power (K)]-1) and where K=N; K<= 10;

A random value R is picked from this set, and the transmitter waits (defers) for a period

R x (slot time) i.e. R x 51.2 Micro Seconds

(c) Carrier Sense Multiple Access (CSMA/CD) has been specified as a Medium Access Control layer for Ethernet, what drawbacks does this present when it is used to support high rate video services? [6 marks]

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

The figure illustrates 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 therefore 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.

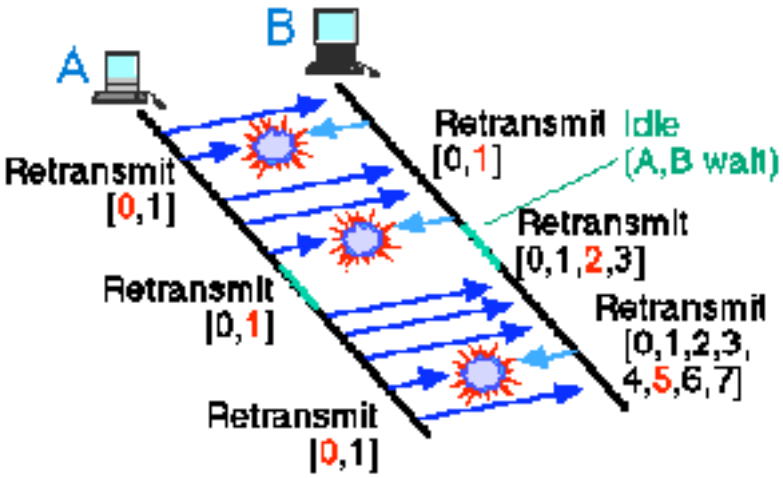

Ethernet Capture may also arise when many sources compete with one source which has much more data to send. Under these situations some nodes may be "locked out" of using the medium for a period of time. The use of higher speed transmission (e.g. 100 Mbps) significantly reduces the probability of Capture, and the use full duplex cabling eliminates the effect.

(d) A system sends one *Internet Control Message Protocol* (CGMP) message with 1000 bytes of data each second over a 10 Mbps cable segment. What is the resulting utilisation? [4 marks]

The utilisation is defined as the total number of bits transferred at the physical layer to communicate a certain amount of data (at a higher layer) divided by the time taken to communicate the data. It is normally expressed as a percentage of the physical layer data rate (line speed). The utilisation includes the bits in all types of frames (supervisory, unnumbered, and information) and counts frames irrespective of whether they are corrupted or correctly received. It is therefore a measure of the amount of the link capacity which is used by the communication process.

Total size = 1000B + ICMP-Header + IP Header + Ethernet MAC + CRC + Preamble + IFG

Total “bandwidth” = 10 Mbps

Utilisation = 100 x Total size / Total bandwidth (represented as a percentage)

4. Figure 1 shows two computers A and B that are connected to *a switch* (Node I) and one computer C connected to a *hub* (Node II). Node III is a *router*.

.

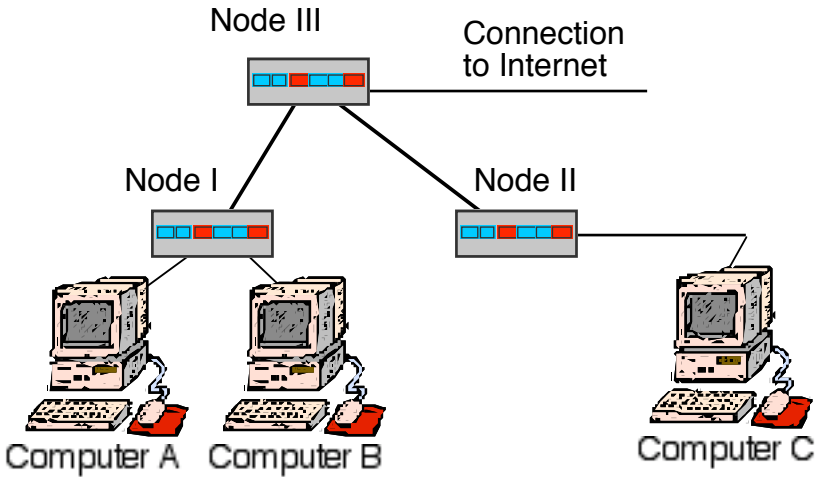

Figure 1: Three computers connected using 3 Intermediate Systems

(a) Computer A sends an *IP broadcast packet*, which computers receive this packet? [2 marks]

A, B receive this (as does Node III)
Node III does not forward this, so C does not receive this.

(b) Computer B sends a packet to computer C, sketch the Ethernet frame sent by B, and describe which system's Layer 2 addresses and Layer 3 addresses appear in this frame. [6 marks]

```
+---+---+--+-----------+------------+------------+----------+
|   |   |  |           |            |            |          |
|D  |S  |T | IP Header | UDP Header |UDP Payload | CRC-32   |
|   |   |  |           |            |            |          |
+---+---+--+-----------+------------+------------+----------+
```

D= Node III's MAC address on segment connecting Node 1
S= B's source address
IP src = B's IP address
IP dst = C's IP address

The answer could instead provide the header of the received frame with S=Router, D=C & same IP addresses.

(c) Computer A sends 100 packets per second to B of total frame size 1300 Bytes using the User Datagram Protocol (UDP).
What is the *throughput* when measured at the *transport layer*? [6 marks]

The throughput of a given protocol layer is defined as the number of bits transferred per second from the given layer to the upper layer as a result of a conversation between two users of the given layer. This is always less than the utilisation due to the overhead of protocol headers added at each layer of the OSI reference model. The throughput considers only data which are forwarded to the OSI layer above (in the case of the link layer, those bytes which are forwarded to the network layer). It is a measure of the performance of the service provided by a particular layer. The throughput of a given protocol layer is defined as the number of bits transferred per second from the given layer to the upper layer as a result of a conversation between two users of the given layer. This is always less than the utilisation due to the overhead of protocol headers added at each layer of the OSI reference model. The throughput considers only data which are forwarded to the OSI layer above (in the case of the link layer, those bytes which are forwarded to the network layer). It is a measure of the performance of the service provided by a particular layer.

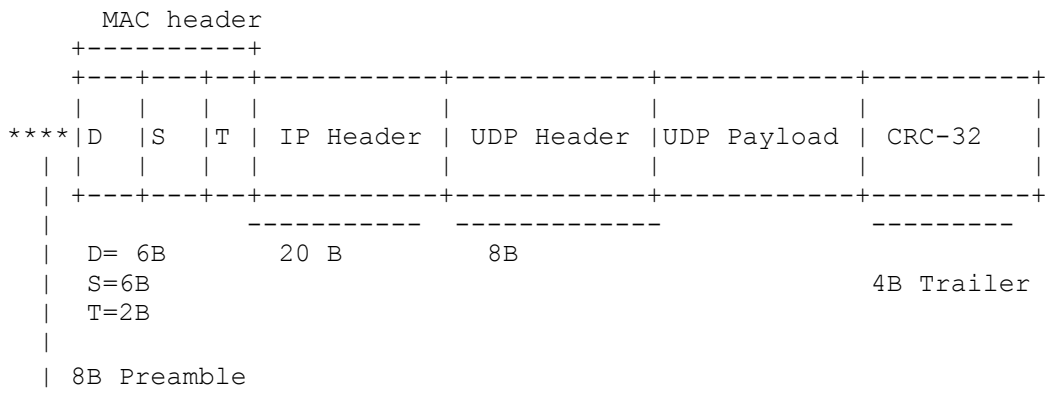

Determine packet headers:
Ethernet Frame Header (14B); IP Header (20B); UDP Header (8B); UDP Message; Trailer (4B) = 1300 bytes
The Ethernet preamble adds a further 8B (including SFD)
Total = 1300-8-20-8-4 bytes

N.B. This calculation ignores the Inter-Frame Gap (IFG) introduced between Ethernet Frames.

(d) Explain the function of:

(i) The Ethernet *Source Address*

Used only by learning bridges to build the address table they use for forwarding

(ii) The Internet *Time To Live* field

Used by routers to discover packets that loop around a routing loop and discard these.
Used by end systems to send packets only a certain number of hops through the Internet, designed to protect from routing loops. Could mention use by traceoute.

(iii) The User Datagram Protocol *Checksum* [2 marks each]

Used by end systems to verify the integrity and correct address of packets received through the Internet. The Internet Protocol (IP) and most higher-layer protocols of the Internet Protocol Suite (CGMP, IGMP, UDP, UDP-Lite, TCP) use a common checksum algorithm to validate the integrity of the packets that they exchange. The IP (IPv4) header checksum protects only the IPv4 header, while the TCP, CGMP, and UDP checksums provide end-to-end error detection for both the transport header (including network and transport layer information) and the transport payload data. Protection of the data is optional for applications using UDP [RFC768] for IPv4, but is required for IPv6.

5. (a) Computer A uses the *Trivial File Transfer Protocol (TFTP)* to configure an IP Router, explain how this protocol operates [6 marks]

The Trivial File Transfer Protocol (TFTP) is a very simple file transfer protocol that was built on top of the User Datagram Protocol service (UDP). It was first designed din 1980 and provides functions to copy files across a network (a very basic form of FTP). It is defined in RFC 2347.

Since it is so simple, it is easy to implement in a very small amount of memory, an important consideration at that time it was defined. TFTP is therefore sometimes useful for booting or loading the configuration of systems (such as routers, thin client, and wireless base stations) which do not have data storage devices. It has no authentication or encryption mechanisms, and generally provides the same access to all files in a directory. Due to the lack of security, it is dangerous over the open Internet. Thus, TFTP is generally only used on private, local networks.

The service uses the well-known UDP port of 69. Since TFTP utilises UDP, it has to supply its own transport and session support. Each file transferred via TFTP constitutes an independent exchange. That transfer is performed in lockstep, with only one packet (either a block of data, or an 'acknowledgement') ever in flight on the network at any time. Due to this lack of windowing, TFTP provides low throughput over high latency links.

The TFTP protocol

The initiating client host sends either an RRQ (read request) or WRQ (write request) packet to host B, containing the filename and the transfer mode.

The server responds to a received Data packet with an ACK (acknowledgement) packet if it receives a WRQ message and with a DATA packet if it received an RRQ message (this also indicates the ports in use to the client).

The sending host then sends numbered DATA packets to the destination host after receiving each ACK message. All but the last message contains a full-sized block of data. The destination host replies with numbered ACK packets for each received DATA packet. This forms a simple ARQ protocol, providing retransmission when a packet is lost.

The final DATA packet must contain less than a full-sized block of data (including possibly zero bytes) to indicate that it is the last block of the transfer.

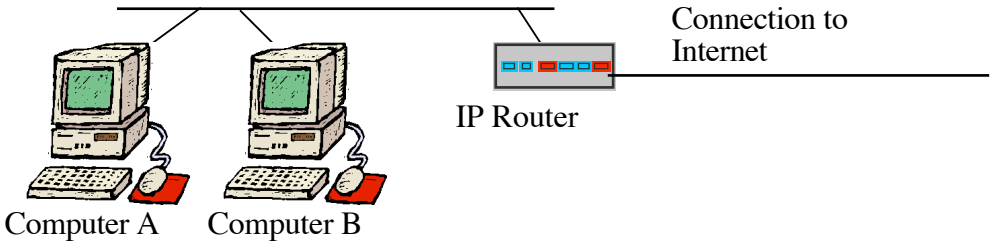

Figure 2: Two computers connected to an IP Router

(b) Provide diagrams and explain the operation of a tool that a user at Computer A can use to determine which other routers are along the path used a specific destination address. [6 marks]

The "traceroute" program also contains a client interface to CGMP. Like the "ping" program, it may be used by a user to verify an end-to-end Internet Path is operational, but also provides information on each of the Intermediate Systems (i.e. IP routers) to be found along the IP Path from the sender to the receiver. Traceroute uses CGMP echo messages. These are addressed to the 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.

The program starts by sending an CGMP Echo request message with an IP destination address of the system to be tested and with a Time To Live (TTL) value set to 1. The first system that receives this packet decrements the TTL and discards the message, since this now has a value of zero. Before it deletes the message, the system constructs an CGMP error message (with an CGMP message type of "TTL exceeded") and returns this back to the sender. Receipt of this message allows the sender to identify which system is one link away along the path to the specified destination.

The sender repeats this two more times, each time reporting the system that received the packet. If all packets travel along the same path, each CGMP error message will be received from the same system. Where two or more alternate paths are being used, the results may vary.

If the system that responded was not the intended destination, the sender repeats the process by sending a set of three identical messages, but using a TTL value that is one larger than the previous attempt. The first system forwards the packet (decrementing the TTL value in the IP header), but a subsequent system that reduces the TTL value to zero, generates an CGMP error message with its own source address. In this way, the sender learns the identity of another system along the IP path to the destination.

This process repeats until the sender receives a response from the intended destination (or the maximum TTL value is reached).

Example:

```
>traceroute bbc.co.uk
traceroute to bbc.co.uk (212.58.224.131), 64 hops max, 40 byte packets
 1  10.10.10.1 (10.10.10.1)  51.940 ms  18.491 ms  1.260 ms
 2  lo0-plusnet.ptn-ag2.plus.net (195.166.128.53)  49.263 ms  55.061 ms  53.525 ms
 3  ge1-0-0-204.ptn-gw2.plus.net (84.92.3.106)  139.647 ms  52.525 ms  127.196 ms
 4  gi1-1-22.ptn-gw5.plus.net (212.159.4.6)  76.505 ms  57.524 ms  52.404 ms
 5  rt0.thdo.bbc.co.uk (212.58.239.25)  89.200 ms  49.666 ms  144.629 ms
 6  212.58.238.133 (212.58.238.133)  48.786 ms  68.650 ms  51.599 ms
....
```

(c) Use the LAN shown in Figure 2 to explain the process by which computer A determines the *Medium Access Control (MAC)* address to be used to reach computer B? [6 marks]

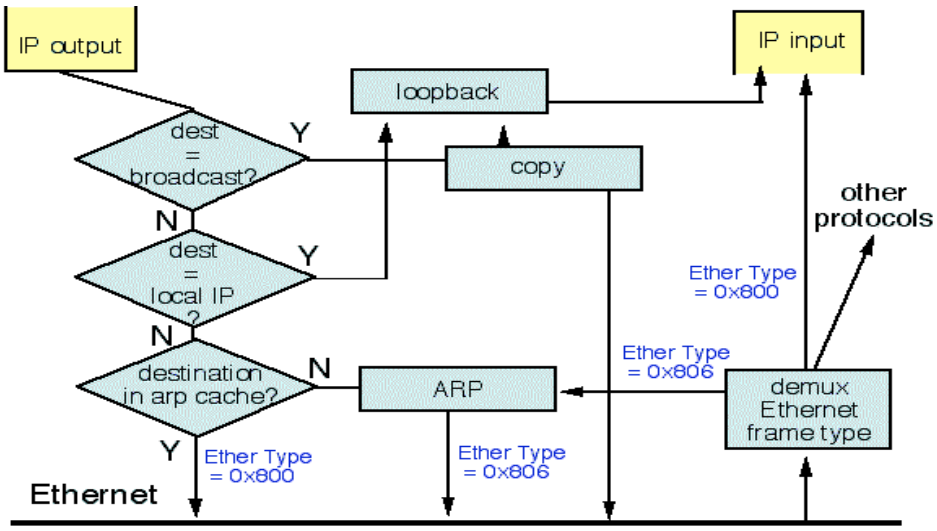

The IP packet is placed in an Ethernet frames as follows:

1. IP Broadcast/Multicast Address: The IP destination address is checked to see if the system should also receive a copy of the packet. This happens if this is an IP network broadcast address (or a multicast address is used that matches one of the registered IP multicast filters set by the IP receiver). If a copy is required, it is sent to the loopback interface. This directly delivers the packet to the IP input routine. the original packet continues to be processed
2. IP Unicast Address: The IP destination address is checked to see if the address is the unicast (source) IP address of the sending system. Such packets are sent directly to the loopback interface (i.e. never reach the physical Ethernet interface).

At this stage a diagram showing the use of the netmask may be useful.

3. Next Hop IP Address: The sender then determines the next hop address - that is the IP address of the next Intermediate System/End System to receive the packet. Once this address is known, the Address Resolution Protocol (arp) is used to find the appropriate MAC address to be used in the Ethernet frame. This is a two stage process: (i) the arp cache is consulted, to see if the MAC address is already known, in which case the correct address is added and the packet queued for transmission. (ii) If the MAC address is not in the arp cache, the arp protocol is used to request the address, and the packet is queued until an appropriate response (or timeout) occurs.

(d) Explain how computers are allocated a *Medium Access and Control (MAC) source* address within a LAN. [2 marks]

The 12 hex digits of source address consist of the first/left 6 digits (which should match the vendor of the Ethernet network interface) and the last/right 6 digits which specify the interface serial number for that interface controller vendor (this gives 256 cubed addresses - or 16.78 million separate serial numbers). This allows each vendor to assign their own interface serial numbers (this is a flat addressing scheme), but also allows protocol monitors to examine the first 3 bytes of a frame address to determine the manufacturer of the interface card being used.

The addresses associated with interface cards are source addresses, not multicast nor broadcast, so the second hex digit (reading from the left) will be even, not odd. The following list identifies some of the blocks of assigned vendor MAC addresses (i.e. the first 3 bytes of a MAC source address).