Marks	Quest	Solution
4	**1a**	(c) **Explain how a system is assigned a Medium Access Control (MAC) address for use on an Ethernet Local Area Network (LAN). [3 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. Vendor MAC addresses (i.e. the first 3 bytes of a MAC source address, the OUI) are purchased from the IEEE.
4	**1b**	**Ethernet Address are sometimes described as belonging to a "Flat Address Space". What does this mean?** The address space is unstructured, in comparison to a hierarchical method. In this scheme blocks of addresses are allocated, and usually used on a first-come first-served basis. Addresses do not denote specific positions within the network topology.
6	**1c**	**Describe the process of *Manchester Encoding*, explaining why this was introduced and illustrating your answer by sketching the waveforms when the following sequence of binary data is transmitted over 10BT cabling:** 0 0 1 1 0**.** Manchester encoding is a synchronous clock encoding technique used by the OSI physical layer to encode the clock and data of a synchronous bit stream. In this technique, the actual binary data to be transmitted over the cable are not sent as a sequence of logic 1's and 0's (known technically as Non Return to Zero (NRZ)). Instead, the bits are translated into a slightly different format that has a number of advantages over using straight binary encoding (i.e. NRZ). In the Manchester encoding, a logic 0 is indicated by a 0 to 1 transition at the centre of the bit and a logic 1 is indicated by a 1 to 0 transition at the centre of the bit. Note that signal transitions do not always occur at the 'bit boundaries' (the division between one bit and another), but that there is always a transition at the centre of each bit. (N.B. since most line driver electronics actually inverts the bits prior to transmission, you may observe the opposite on an oscilloscope connected to a cable). Manchester encoding may be alternatively viewed as a phase encoding where each bit is encoded by a positive 90 degree phase transition, or a negative 90 degree phase transition. The Manchester code is therefore sometimes known as a Biphase Code. Original Data Value Sent Logic 0 0 to 1 (upward transition at bit centre) Logic 1 1 to 0 (downward transition at bit centre) 0 0 1 1 0 therefore becomes: 0101101001 - answer must show timing for the 10 Mbps

Required waveform looks like:

```
 +-+ +----+ +-+ +-+    +-
 |  |  |     |  |  |  |  |     |
-+ +--+     +--+ +--+  +----+
-+  +--+    +--+  +--+ +----+
 |  |  |     |  |  |  |  |     |
 +-+ + +----++ +-+  +-+ +-
```

Note that in some cases you will see the encoding reversed, with 0 being represented as a 0 to 1 transition. This occurs because when an inverting line driver is used to convert the binary digits into an electrical signal, the signal on the wire is the exact opposite of that output by the encoder.

Differential physical layer transmission, (e.g. 10BT) transmits both a positive and negative version of the signals. This does not suffer this inversion, since the polarity can be determined from the SFD.

Additional, not required:

The bi-phase Manchester encoding can consume up to approximately twice the bandwidth of the original signal (20 MHz). While this was of little concern in coaxial cable transmission, the limited bandwidth of CAT5e cable necessitated a more efficient encoding method for 100 Mbps transmission using a 4b/5b MLT code. This uses three signal levels (instead of the two levels used in Manchester encoding) and therefore allows a 100 Mbps signal to occupy only 31 MHz of bandwidth. Gigabit Ethernet utilises five levels and 8b/10b encoding, to provide even more efficient use of the limited cable bandwidth, sending 1 Gbps within 100 MHz of bandwidth.

An Ethernet controller 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. Strictly speaking the last byte which finished with the '11' is known as the "Start of Frame Delimiter". When encoded using Manchester encoding, at 10 Mbps, the 62 alternating bits produce a square wave.

Marks	Quest	Solution
6	**1d**	**Sketch the *Ethernet Preamble* waveform and explain the purpose of the *Ethernet Preamble* at the start of each frame.** An Ethernet controller 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. Strictly speaking the last byte which finished with the '11' is known as the "Start of Frame Delimiter". When encoded using Manchester encoding, at 10 Mbps, the 62 alternating bits produce a square wave. 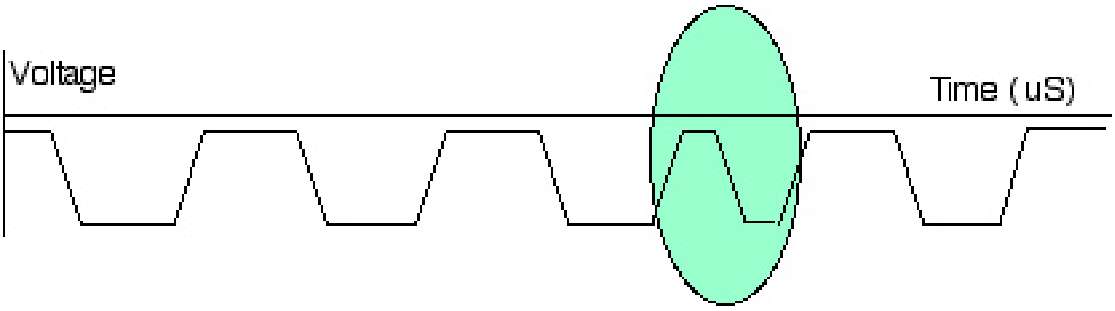 The preamble allows time for the receiver in each node to achieve lock of the receiver Digital Phase Lock Loop which is used to synchronise the receive data clock to the transmit data clock. At the point when the first bit of the preamble is received, each receiver may be in an arbitrary state (i.e. have an arbitrary phase for its local clock). During the preamble the receiver learns the correct phase, but in so doing it may miss (or gain) a number of bits. Hence the SFD is needed to align the first byte of the MAC header.

Marks	Quest	Solution
8	**2a**	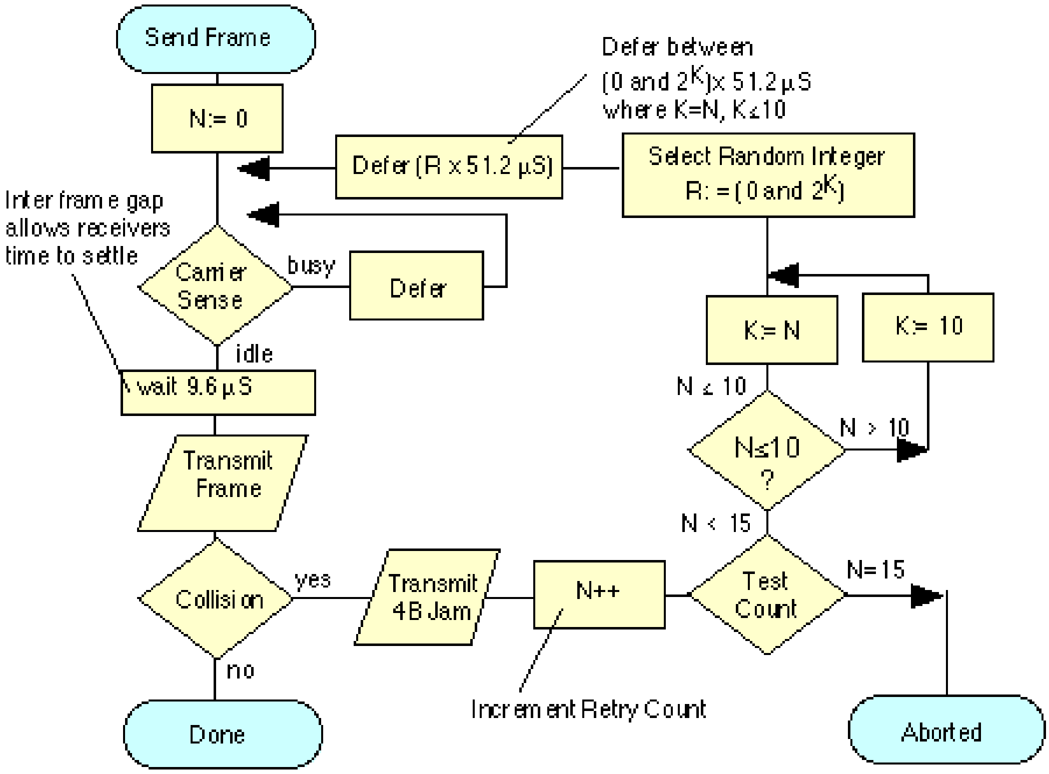 Send Frame N:= 0 Defer between (0 and 2K)x 51.2 µS where K=N, K≤10 Defer (R x 51.2 µS) Select Random Integer R: = (0 and 2K) Inter frame gap allows receivers time to settle Carrier Sense busy Defer idle wait 9.6 µS Transmit Frame K:= N K:= 10 N ≤ 10 N≤10 ? N > 10 N < 15 Collision yes Transmit 4B Jam N++ Test Count N=15 no Increment Retry Count Done Aborted Students should provide details in diagram or describe operation using text. 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. Transmission then starts with the preamble, followed by the frame data and finally the CRC-32. After this time, the transceiver transmit logic is turned off and the transceiver returns to passively monitoring the cable for other transmissions. During this process, a transmitter must also continuously monitor the collision detection logic (CD) in the transceiver to detect if a collision occurs. If it does, the transmitter aborts the transmission (stops sending bits) within a few bit periods, and starts the collision procedure, by sending a Jam Signal to the transceiver transmit logic. It then calculates a retransmission time. If all nodes attempted to retransmit immediately following a collision, then this would certainly result in another collision. Therefore a procedure is required to ensure that there is only a low probability of simultaneous retransmission. The scheme adopted by Ethernet uses a random back-off period, where each node selects a random number, multiplies this by the slot time (minimum frame period, 51.2 µS) and waits for this random period before attempting retransmission. The small Inter-Frame Gap (IFG) (e.g., 9.6 microseconds) is also added. On a busy network, a retransmission may still collide with another retransmission (or possibly new data being sent for the first time by another node). 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. The scaling is performed by multiplication and is known as exponential back-off. This is what lets CSMA/CD scale to large numbers of nodes - even when collisions may occur. The first ten times, the back-off waiting time for the transmitter suffering collision is scaled to a larger value. The algorithm includes a threshold of 1024. The reasoning is that the more attempts that are required, the more greater the number of computers which are trying to send at the same time, and therefore the longer the period which needs to be deferred. Since a set of numbers {0,1,...,1023} is a large set of numbers, there is very little advantage from further increasing the set size. Each transmitter also limits the maximum number of retransmissions of a single frame to 16 attempts (N=15). After this number of attempts, the transmitter gives up transmission and discards the frame, logging an error. In practice, a network that is not overloaded should never discard frames in this way and can effectively share the available capacity of an Ethernet segment.

Marks	Quest	Solution
8	2b	***Unshielded Twisted Pair* (UTP) cabling was originally used as the physical layer for 10BT LANs. What challenges were faced when using this links operating at 100 Mbps?**

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. Diagrams showing frequency response of typical cable may be helpful.

The bi-phase Manchester encoding can consume up to approximately twice the bandwidth of the original signal (20 MHz). While this was of little concern in coaxial cable transmission, the limited bandwidth of necessitated a more efficient encoding method 100 Mbps using a 4b/5b MLT code. 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.

4B/5B encoding is a type of 'Block coding'. This processes groups of bits rather than outputting a signal for each individual bit (as in Manchester encoding). A group of 4 bits is encoded so that an extra 5th bit is added. Since the input data is taken 4-bits at a time, there are 2^4, or 16 different bit patterns. The encoded bits use 5-bit, and hence have 2^5 or 32 different bit patterns. As a result, the 5-bit patterns can always have two '1's in them even if the data is all '0's a translation occurs to another of the bit patterns. This enables clock synchronisation, required for reliable data transfer.

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 combined scheme of 4b/5b with MLT-3 encoding leads to a waveform of 31.25 MHz, well within the specification for Unshielded Twisted Pair Cabling. Diagrams may be useful.

Answer could also highlight the issue of inter-frame gaps, identifying how the 10B2 specification adapted in later technologies. Answer may also wish to discuss half and full duplex operation and design constraints for 100BT and 1000BT. No attempt was made to support the older co-axial cable installations, instead these must be replaced by UTP or Fibre.

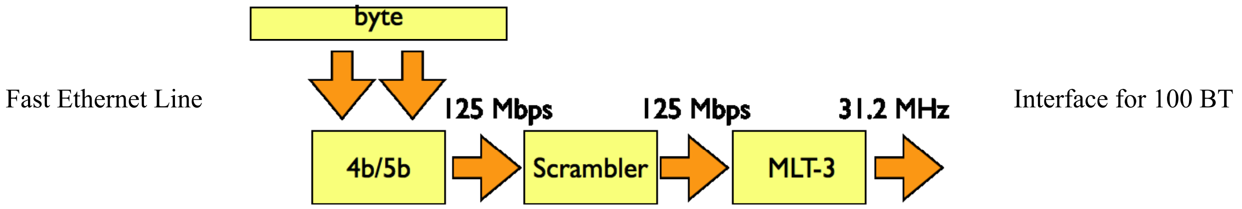

<table>
<tr><th>Marks</th><th>Quest</th><th>Solution</th></tr>
<tr><td><b>4</b></td><td><b>2c</b></td><td>
<p><b>In the context of Fast Ethernet explain how the sequence of bits {1 0 0 1 1 1 } are encoded using Multi-level threshold, MLT-3 line encoding .</b></p>
<p>The line encoding using 3-level MLT encoding should be shown. The student should note the three levels used (+, -, 0) to reduce the required bandwidth to 31.25 MHz, suitable for transmission over the installed based of twisted pair cables.</p>
<p>Encoded as: +++0-- (assuming a zero start and positive waveform)</p>
<pre>
+ve   +-+-+-+
      | |
0   -+   +-+
              |
-ve   +-+-+
</pre>
<p>Appropriate diagrams similar to above should be provided.</p>
</td></tr>
</table>

Marks	Quest	Solution
8	3a	**Describe the operation of a switch. Illustrate your answer by showing how the switch in figure 1 (above) uses address learning to determine whether frames from computer A are to be forwarded from network X to network Y.** A bridge is a LAN interconnection device which may be used to join two LAN segments (X,y), 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 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 A,B are observed to be on network X, while the address of computer C will be observed to be on network Y. 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 override the normal forwarding by inserting entries in a filter table to inhibit forwarding between different workgroups (for example to provide security). Packets with a source of A and destination of B are received and discarded, since the computer B is directly connected to the LAN X, whereas packets from A with a destination of C are forwarded to the LAN Y by the bridge. Summary A learning bridge identifies which addresses are remotte and local by observing the source address. Filtering based on destination MAC address and may provide security filtering 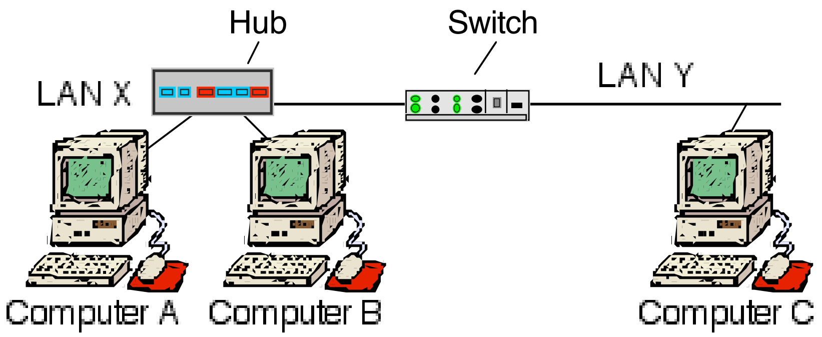 Hub Switch LAN X LAN Y Computer A Computer B Computer C *Figure 1: An Ethernet LAN*

Marks	Quest	Solution
2	3b	**Why should simple Ethernet switches not be connected in a "loop"?** Layer 2 switches connected in a loop will lead to **looping** of packets (amplification). Unmanaged bridges must form a tree, and not a ring. That is, there must be only one path between any two computers. If more than one parallel path were to exist, a loop would be formed, resulting in endless circulation of frames over the loop. This would soon result in overload of the network.

Marks	Quest	Solution
5	3c	**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 the computer B.** 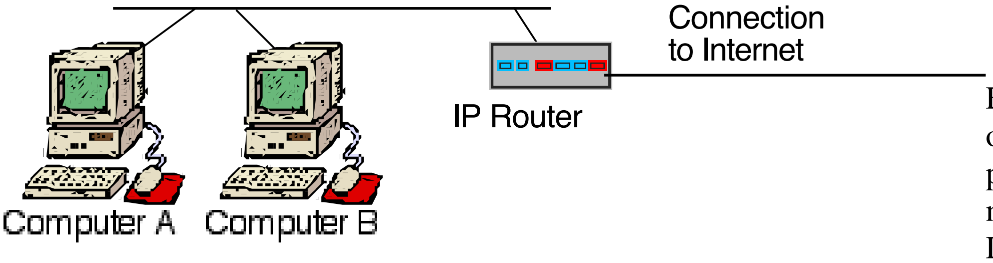 Connection to Internet IP Router Computer A Computer B Each IP address consists of two parts, the network part (identifying the network number, or LAN collision domain, to which the computer is attached) and the host part (which identifies the host within the local network). The IP network ID is identified as the bit-wise logical AND of the 32-bit IP address with another 32-bit quanity, the netmask. All systems with the same network number share the same netmask (sometimes called a "subnet mask"). This has a bit with a logical '1' for each bit that is a part of the network number, and a logical '0' for each bit which is a part of the host number. The first part of the method compares source and destination network IDs. If the two match, use ARP to find a MAC address, and sends directly (posisblky invoking arp) If they do not match the system sets the mac address to be the default router's MAC address. If this is not known, the address is discovered using the arp protocol with the IP address of the default router.
3	3d	**Identify the protocol layers associated with *each* of the following three protocols:** *User Datagram Protocol (UDP)*, - Transport *Carrier Sense Multiple Access* (CSMA/CD) protocol, - Link *Internet Protocol* (IP). - Network
2	3e	**IP is specified as a *Best Effort* service. What does this term mean?** Packets are not necessarily delivered - there may be loss Packets may be delivered out of order Some packets may be duplicated

Marks	Quest	Solution
4	4a	***The User Datagram Protocol (UDP) is a simple transport protocol supported by the Internet Protocol (IP) suite. Explain the function of each of the protocol headers. [5 marks]*** The UDP header consists of four fields each of 2 bytes in length: Source Port (UDP packets from a client use this as a service access point (SAP) to indicate which session on the local client originated the packet. UDP packets from a server carry the server SAP in this field) Destination Port (UDP packets from a client use this as a service access point (SAP) to indicate which service is required from the remote server. UDP packets from a server carry the client SAP in this field) UDP length (The number of bytes of data) UDP Checksum (A checksum to verify that the end to end data has not been corrupted by routers or bridges in the network or by the processing in an end system. If this check is not required, the value of 0x0000 is placed in this field, in which case the data is not checked by the receiver.)

Marks	Quest	Solution
9	4b	**How does the Domain Name Service (DNS) resolver in the network stack use a remote DNS server to resolve the name that corresponds to a specific IP network address?** A name is a human-readable label assigned to a system. An address is the basic routing identifier used to locate a system in the network. - 2 marks each for clarity of the above definitions DNS Service: Mapping between the two is performed using the domain name service. This is an example of a client/server system which is used by the Internet Protocol (IP) Suite to resolve the logical names of nodes in an IP network to an IP address (see also arp - which is used to resolve Ethernet addresses to IP addresses). The address resolution procedure is completed when the client receives a response from the server containing the required address. this is then used as the IP destination address. Next hop resolution provides a MAC address based on this IP address. Resolution query: 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 client resolver must be pre-configured with the IP address of the DNS server. Resolution response: 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. Resolution may require recursive lookup on one or more DNS servers to finally receive an authoritative answer. Recursion involves searching multiple databases until the result is retrieved or it is concluded the name is unknown. Cache: A key point to be noted is that the system is requested by an application and the results are cached - so that the lookup does not need to be performed for every single use. In DNS the information provider determines the cache time - vastly different values are used for different applications (small where there is a churn of addresses, large for main infrastructure stability where change is not envisaged).

Marks	**Quest**	**Solution**

7	**4c**	**The figure below show the first few bytes of a frame captured using a network monitor:**
		`00e0 f726 3fe9 0800 2086 354b 0800 4500`
		Figure 3
		(i) What is the Ethernet *Destination Address*? [2 marks]
		`00e0 f726 3fe9`
		(ii) Is this a multicast frame? [1 mark]
		No - lsb of first byte is 0
		(iii) The interface was set to promiscuous mode, what does this mean? [3 marks]
		It means all frames are captured irrespective of the source address (or other addresses configured as L2 adddress filters). This mode is used in bridging (but not as a host or router). It can also be used for network analysis
		(iv) What sort of packet is contained in the frame payload?
		IP - see sheet at end of exam paper for Ethertypes. **[1 mark]**

Marks	Quest	Solution
5	**5a**	The figure below show the first few bytes of a frame captured using a network monitor: 001a 2f18 9790 001e c2be 4c73 **0800** **4**5c0 0028 d0b5 0000 ff**06** c184 0a0a 0a01 0a0a 0ac1 0017 c059 4435 64b2 404b becb 5010 0fdf 0eb1 0000 0000 0000 0000 Figure 4: Captured Ethernet Frame **Explain how this frame may be decoded to find the *Service Access Point* added by each protocol layer, and hence determine the set of protocols that were used.** 0800 EtherType - indicates type of packet payload, e.g. IP, IUPv6, ARP, LLC, etc [1 mark] Network Layer Type = 4 (IPv4) [2 marks] Network transport (protocol) = 6 = TCP [2 marks]
5	**5b**	**An *End System* uses the "ping" program to send a message with a payload of 100B. What is the total size of the Ethernet frame sent?** Determine packet headers: Ethernet Frame Header (14B); IP Header (20B); ICMP Mesage (108 B); Ethernet Trailer (4B) ***Size = 1188B.*** N.B. This calculation ignores the Inter-Frame Gap (IFG) introduced between Ethernet Frames. = 8+14+20+8+100+4 =1 54 bytes = 1232b

Marks	Quest	Solution
10	**5c Option 1**	**By comparing the operation of the "ping" program and the "traceroute" program describe the key differences between these two applications.** The "ping" program contains a client interface to ICMP. It may be used by a user to verify an end-to-end Internet Path is operational. The ping program also collects performance statistics (i.e. the measured round trip time and the number of times the remote server fails to reply. Each time an ICMP echo reply message is received, the ping program displays a single line of text. The text printed by ping shows the received sequence number, and the measured round trip time (in milliseconds). Each ICMP Echo message contains a sequence number (starting at 0) that is incremented after each transmission, and a timestamp value indicating the transmission time. The "traceroute" program also contains a client interface to ICMP. 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 ICMP 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 ICMP 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 ICMP error message (with an ICMP 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 ICMP 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 ICMP 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). Some Routers are configured to discard ICMP messages, while others process them but do not return ICMP Error Messages. Such routers hide the "topology" of the network, but also can impact correct operation of protocols. Some routers will process the ICMP Messages, providing that they do not impose a significant load on the routers, such routers do not always respond to ICMP messages. When "traceroute" encounters a router that does not respond, it prints a "*" character.

Marks	**Quest**	
10	**5c Option 2**	**Explain how an IP host using Ethernet can automatically determine the Medium Access Control address for another IP host connected to the same Ethernet broadcast domain. You answer should include relevant diagrams to illustrate how this works when three consecutive packets are sent to a new host that has just been added to the Ethernet network.** A sender consults the arp cache before sending.If the IP address is in the cache it may be used to set the destination MAC address of a frame. If the MAC address is not in the ARP cache, then an ARP request packet is first used. The address resolution protocol is used by the Internet Protocol (IP) to map IP network destination addresses to the hardware addresses used by a link protocol. The protocol operates below the network layer as a part of the OSI link layer, and is used when IP is used over Ethernet. The system queues a packet for which an ARP request is being sent, hoping to receive an arp response indicating the target MAC address bound to the target IP address. When ARP response is received, the address is entered in the cache. If the address resolves the address needed for a queued packet - this is then sent using the new information to set the MAC destination address. Note any queries can also be use to populate the cache. The cache entries are expired after a period of time - to prevent the cache holding out of date information. relevant diagrams showing the protocol exchange are desirable. These should clearly show the request is broadcast and the response is unicast back to the querier.