

Qu No.	Solution	Page ... of
Marks		
5	<p data-bbox="368 327 1342 427"><i>Please note that both exams have identical solutions, however the level of detail expected in ES is less, and the questions are phrased to provide more guidance on how to provide the solution.</i></p> <p data-bbox="225 450 1334 562">1a An Ethernet frame (represented below in hexadecimal) was recorded by network monitor. Explain how this frame may be decoded to show the transport protocol that was used and the expected application.</p> <p data-bbox="296 595 1350 663">The answer is that the packet has a protocol ID of 6 (SAP to Transport Layer), indicating TCP.</p> <p data-bbox="296 696 520 730">Hexadecimal data:</p> <pre data-bbox="296 763 1054 943"> 001e c2be 4c73 001a 2f18 9790 0800 4510 0040 c72f 4000 4006 4aa3 0a0a 0ac1 0a0a 0a01 c059 0017 404b be89 0000 0000 b002 ffff 54ff 0000 0204 05b4 0103 0303 0101 080a 2863 d181 0000 0000 0402 0000 </pre> <p data-bbox="296 976 448 1010">Full decode:</p> <p data-bbox="296 1077 759 1111">Decode at MAC layer (fixed 14 bytes):</p> <p data-bbox="296 1144 895 1290">Destination 00:1e:c2:be:4c:73 (unciastr) Source 00:1a:2f:18:97:90 Ethertype IPv4 (0x0800) --- SAP to network layer Length 78</p> <p data-bbox="296 1312 903 1346">Decode at IP layer (variable, in this case 20 bytes):</p> <p data-bbox="296 1379 1318 1559">Src IP: 10.10.10.193.49241 Dest IP: 10.10.10.1 IP (tos 0x10, ttl 64, id 50991, offset 0, flags [DF], proto TCP (6) --- SAP to Transport Layer Length 64)</p> <p data-bbox="296 1581 552 1615">Decode at TCP layer:</p> <p data-bbox="296 1648 775 1794">Dest port: telnet TCP Flags: S, cksum 0x54ff (correct), 1078705801:1078705801(0) win 65535</p> <p data-bbox="296 1816 1270 1850">TCP data <mss 1460,nop,wscale 3,nop,nop,timestamp 677630337 0,sackOK,eol></p>	

Qu No.1	Solution	Page ... of
Marks 8	<p>1b 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.</p> <p>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.</p> <p>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.</p> <p>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: $\{0, 1, 2, 3, 4, 5, \dots L\}$ where L is $([2 \text{ to the power } (K)]-1)$ and where $K=N$; $K \leq 10$; A random value R is picked from this set, and the transmitter waits (defers) for a period $R \times (\text{slot time})$ i.e. $R \times 51.2 \text{ Micro Seconds}$</p> <p>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 $\{0,1,\dots,1023\}$ is a large set of numbers, there is very little advantage from further increasing the set size.</p> <p>Each transmitter also limits the maximum number of retransmissions of a single frame to 15 (N=15). After this number of attempts, the transmitter gives up transmission and discards it logging an error. In practice, a network that is not overloaded should never discard frames can effectively share the available capacity of an Ethernet segment.</p>	

Qu No.1	Solution	Page ... of
Marks 2	1c There are two issues that are important: •Some nodes in the network may be hidden from other nodes. This can be alleviated by use of a CTS/RTS handshake. • The WiFi base station can become a dominant source of data, and hence methods may be needed to prevent CSMA/CD capture. Answers must at least identify the first and may also indentify the second point.	

1d Key issues faced - Cross talk and restricted bandwidth, requiring new PHY technology.

100BASE-T provides 100 Mbit/s Ethernet in either half-duplex (using CSMA/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 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, thus 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.

6

- 1e Gigabit Ethernet utilises five levels and 8b/10b encoding, to provide even more efficient use of the limited cable bandwidth, sending 1 Gbps within approx 100 MHz of bandwidth (i.e. the capacity of a UTP Cat5e cable)

Using 8b/10b encoding each byte of data is assigned a 10 bit code. The byte is split up into the 3 most significant bits and the 5 least significant bits. This is then represented as two decimal numbers with the least significant bits first e.g. for the octet 101 00110 the result is the decimal 6.5. 10 bits are used to create this code and the naming convention follows the format /D6.5/. There are also 12 special codes which follow the naming convention /Kx.y/.

The 10 bit code must contain either five ones and five zeros, or four ones and six zeros, or six ones and four zeros. This prevents a sequence of too many consecutive ones and zeros, assisting clock synchronisation. Two 'commas' are used to aid in bit synchronisation, these 'commas' are the 7 bit patterns 0011111 (+comma) and 1100000 (-comma).

To maintain a DC balance, a calculation called the Running Disparity calculation is used to keep the number of '0's transmitted the same as the number of '1's transmitted. The method uses 10 bits for each 8 bits of data (byte) and therefore increases the rate required to send the data.

A 1Gbps the line speed results in a transmission rate of $10/8 \times 1 = 1.25\text{Gbps}$. In Gigabit Ethernet this rate is then reduced using PAM-5 a 5-level code (achieving less bandwidth than possible with a 3 level code).

The interface encoded byte of data generates a 10 bit code that is scrambled and converted into a physical layer signal using a 5-level Pulse-Amplitude-Modulation (PAM).

Appropriate diagrams should be supplied.

10

2a

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”. 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).

Summary:

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

3

2b

Diagrams should be used to illustrate the answer.

What *forwarding rate* would be required by a bridge port to fully utilise a 1 Gbps Ethernet output interface?

Method:

Calculate the minimum frame size from the following:

- * Preamble – 8 bytes
- * 64 byte minimum size

= 72 bytes or 576 bits

Maximum forwarding rate is $1000/576 \times 1000\ 000$ pps
~ 2 million packets per second

This ignores inter-frame gap effects.

4 2ci Preamble :
A node starts transmission by sending an 8 byte (64 bit) preamble sequence. This consists of 62 alternating 1's and 0's followed by the pattern 11. 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 5 MHz square wave. When the special pattern (11), is received, the Ethernet receive interface starts collecting the bits into bytes for processing by the MAC layer.

The purpose of the start of frame delimiter is to allow time for the receiver in each node to achieve lock of the receiver Digital Phase Lock Loop which is used to synchronise the receive data clock to the transmit data clock. At the point when the first bit of the preamble is received, each receiver may be in an arbitrary state (i.e. have an arbitrary phase for its local clock). During the course of the preamble it learns the correct phase, but in so doing it may miss (or gain) a number of bits.

4 2cii A Content-addressable memory (CAM) is a type of computer memory used in high speed searching applications, where the a data word is applied to the input (in this case a MAC address) and the CAM searches its entire memory to see if that data word is stored anywhere in it. If the data word is found, the CAM returns a list of one or more storage addresses where the word was found (in an Ethernet MAC CAM, it also returns an associated piece of data, that indicates the port, timeout, and can be used to flag other information – such as filter/trap information).

The use of CAMs dramatically speeds-up address lookup, essential for high-speed forward

4 2c iii A Filter Table may be used in a router or managed switch 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 by discarding packets to specified destination addresses.

In sophisticated devices, filter tables may also be used above the MAC layer, e.g. to perform VLAN functions or to set / police QoS codepoints at the Ethernet or IP layers.

Filter table entries may also be used to trigger management functions, such as traps – to alert a network operator to specific traffic patterns.

8

3a The set of fields that are modified in the Layer 2 and Layer 3 protocol headers as a router forwards an IP packet are:

ETHER: Destination - replaced by next hop Layer 2 address
ETHER: Source - replaced by routers outgoing interface address
IP: Total length = 40 bytes (unchanged – unless options were updated)
IP: Flags = (unchanged – unless DF=0 and fragmentation applied)
IP: Fragment offset = (unchanged – unless fragmentation applied)
IP: Time to live = reduced by one (unless 1 already, then discarded)
IP: Header checksum = always updated , since TTL decremented

A more complete analysis is given below:

ETHER: ----- [Ether Header](#) -----

ETHER:

ETHER: Destination - replaced by next hop Layer 2 address

ETHER: Source - replaced by routers outgoing interface address

ETHER: Ethertype - 0800 (IP) – as in received frame

ETHER:

IP: ----- [IP Header](#) -----

IP:

IP: Version = 4 (unchanged)

IP: Header length = 20 bytes (unchanged)

IP: Type of service = 0x00 (unchanged, unless ECN marked)

IP: Total length = 40 bytes (unchanged – unless options were updated)

IP: Identification = 43773 (unchanged)

IP: Flags = (unchanged – unless DF=0 and fragmentation applied)

IP: Fragment offset = (unchanged – unless fragmentation applied)

IP: Time to live = reduced by one (unless 1 already, then discarded)

IP: Protocol = (unchanged)

IP: Header checksum = always updated , since TTL decremented

IP: Source address = (unchanged)

IP: Destination address = (unchanged)

IP: Options not normally present.

Transport: (unchanged, operates end-to-end)

6 3b Utilisation is a measure of the capacity used in the physical layer. It includes all protocol header information added by the layer itself and the layers below. It is usually measured as a percentage.

No frame may have less than 46 bytes of payload - i.e. 64 B in total.

Smallest Ethernet Frame Size = 64B

Utilisation = $(164 \times 8) \times 100 \times 100,000 / 100,000,000 \%$
= 51%

(N.B. Ignoring Interframe gap).

5 3c The field is called a Cyclic Redundancy Check (also known as a Frame Check Sequence). The link layer CRC protects the frame from corruption while being transmitted over the physical medium (cable). The CRC is removed by routers - as part of the processing. A new CRC is added if the packet is forwarded by the router on another Ethernet link. While the packet is being processed by the router the packet data is protected by the CRC. Router processing errors may otherwise pass undetected.

The 32-bit CRC 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.

The CRC is also used to verify correct bit-timing, i.e. to identify the DPLL has not missed/

6 2d The “traceroute” program 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.

Traceroute 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).

Diagrams may be provided to show the effect of TTL scoping.

8

4a

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". 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).

Summary:

- MAC Source address observed for learning
- Associated with a port in the address table (each frame)
- MAC Destination address observed for forwarding
- Learned addresses -> forward only to specified port
- Discard frames to received port (don't forward back to own interface)
- Flood frames with unknown addresses to all ports (As broadcast L2)
- Aging required and re-learning when computers change the port they are connected

Important note the above topology contains a loop (ring). This topology is illegal in a L2 network, since it would result in continuous looping of packets and there is no layer 2 mechanism to prevent this. Complete answers must note this!

An addition to Ethernet, the Spanning Tree Algorithm, can allow rings to be formed, by logically breaking the ring at a point to form a tree.

- 4b In the case of IP routers, forwarding is performed at the network layer, rather than the link layer. The key points here are that routers in the network have a perspective of the correct path to use to reach a destination. This is acquired via a L3 routing protocol (or by static configuration). This allows routers to take advantage of alternate paths. A topology including loops and parallel links is common, and traffic can flow over any link available to reach the required destination.

A router is an Intermediate System (IS) which operates at the network layer of the OSI reference model. Routers may be used to connect two or more IP networks, or an IP network to an internet connection. A router consists of a computer with at least two network interface cards supporting the IP protocol. The router receives packets from each interface via a network interface and forwards the received packets to an appropriate output network interface. Received packets have all link protocol headers removed, and transmitted packets have a new link header added prior to transmission.

The router queue of packets waiting to be sent also introduces a potential cause of packet loss. Since the router has a finite amount of buffer memory to hold the queue, a router which receives packets at too high a rate may experience a full queue. In this case, the router has no other option than to simply discard excess packets. If required, these may later be retransmitted by a transport protocol. The buffers are passed (drained) into the input to the **forwarding engine**. This takes each buffer, one at a time, and removes it from the interface receiver. The packet is then forwarded to an appropriate output interface, corresponding to the "best" path to the destination specified in the destination address of the IP packet header. The router uses the information held in the network layer header (i.e. IP header) to decide whether to forward each received packet, and which network interface to use to send the packet. Most packets are forwarded based on the packet's IP destination address, along with routing information in the routing table.

The routing and filter tables resemble similar tables in link layer bridges and switches. Except, that instead of specifying link hardware addresses (MAC addresses), the router tables specify network (IP addresses). The routing table lists known IP destination addresses with the appropriate network interface to be used to reach that destination. A default entry may be specified to be used for all addresses not explicitly defined in the table. 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 computers by discarding packets to specified destination addresses.

All IP packets contain a TTL value that determines the number of router hops that a packet may be routed, this is decremented by most routers (although routers can be configured to reduce the value by more than 1) It is used to prevent routing loops, and ensures topologies with loops do not result in packets that circulate indefinitely.

At the output interface, the packet (together with a new link layer header) is placed into a transmit queue until the **link layer processor** is ready to transmit the packet. This, like the receive queue, is a FIFO queue, and usually also takes the form of a ring of memory buffers. Each out-going packet requires a new link layer protocol header to be added (encapsulation) with the destination address set to the next system to receive the packet. The link protocol controller also maintains the **hardware address table** associated with the interface. This usually involves using the Address Resolution Protocol (arp) to find out the hardware (Medium Access Control) addresses of other computers or routers directly connected to the same cable (or LAN). The packet is finally sent using the media interface with the hardware address set to the next hop system. When complete, the buffer (memory) allocated to the frame, is "freed", that is, it is returned as an empty buffer to the receive queue, where it may be used to store a new received packet.

6 4c 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.


To reduce the number of address resolution requests, a client normally caches resolved addresses for a (short) period of time. The arp cache is of a finite size, and would become full of incomplete and obsolete entries for computers that are not in use if it was allowed to grow without check. The arp cache is therefore periodically flushed of all entries. This deletes unused entries and frees space in the cache. It also removes any unsuccessful attempts to contact computers which are not currently running.

Reasons for fewer ARP responses than requests:

- 1) Link corruption (least likely explanation)
- 2) Ethernet switches forward broadcast request packets to all parts of the LAN, but responses are unicast to the requester. So number of responses seen is a function of location of the probe monitoring the connection.
- 3) The respondent may not be connected to the network!

6 4d Determine packet headers:
 Ethernet Frame Header (14B); IP Header (20B); UDP Header (8B); UDP Message (1400 B); Ethernet Trailer (4B)
 = 1446 bytes (11568 bits)
 The Ethernet preamble adds a further 8B (including SFD)

Total = 1454 bytes (11632 bits)



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

8

- 5a 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 quantity, the net mask. 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 forwarding method compares the source and destination network IDs at the received interface.

If the two match, use ARP to find a MAC address, and send directly.

If they do not match, send the packet to a router (this may require a route-lookup - if routing is used, otherwise the default route is used). If the router's MAC address is not in the ARP cache, then ARP 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.

Hence in this example:

Destination B:

Dest MAC = 00:01:00:00:02:00
Dest IP = 192.0.2.3

Destination C:

Dest MAC = 00:02:00:00:01:00 – note identify which interface is the local one
Dest IP = 208.77.188.166

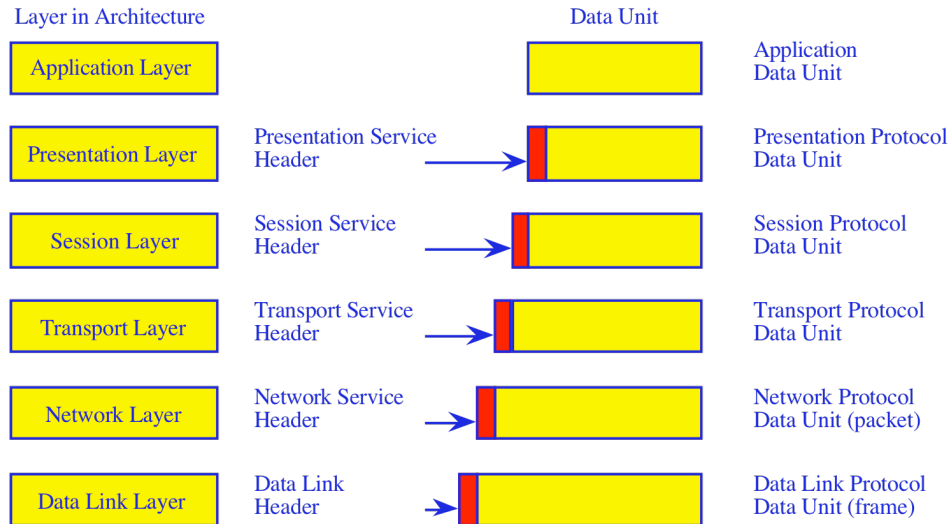
2

5b 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 marks for the diagram showing OSI Layers, correctly labeled., additional marks for detailed explanation.



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 identified - the role of the network layer, it's hop-by-hop forwarding and the way in which it routes traffic between systems.

5cii Circuit switching is the most familiar technique used to build a communications network. It is used for ordinary telephone calls. It allows communications equipment and circuits, to be shared among users. Each user has sole access to a circuit (functionally equivalent to a pair of copper wires) during network use. Network use is initiated by a connection phase, during which a circuit is set up between source and destination, and terminated by a disconnect phase. After completion of the connection, a signal confirming circuit establishment (a connect signal in the diagram) is returned; this flows directly back to node A with no search delays since the circuit has been established. Transfer of the data in the message then begins. After data transfer, the circuit is disconnected; a simple disconnect phase is included after the end of the data transmission. Delays for setting up a circuit connection can be high, especially if ordinary telephone equipment is used. Call setup time with conventional equipment is typically on the order of 5 to 25 seconds after completion of dialling.

Message switching sends a complete message, forwarding it hop-by-hop through a network of nodes with internal storage. At each node the message is forwarded as the outbound link becomes available. Since the message may be competing with other messages for access to facilities, a queuing delay may be incurred while waiting for the link to become available. It repeats this process until it reaches its destination.

Circuit setup delays are replaced by queuing delays. Considerable extra delay may result from storage at individual nodes. A delay for putting the message on the communications link (message length in bits divided by link speed in bps) is also incurred at each node en route. Message lengths are slightly longer than they are in circuit switching, after establishment of the circuit, since header information must be included with each message; the header includes information identifying the destination as well as other types of information.

Packet switching is similar to message switching using short messages. Any message exceeding a network-defined maximum length is broken up into shorter units, known as packets, for transmission; the packets, each with an associated header, are then transmitted individually through the network. The fundamental difference in packet communication is that the data is formed into packets with a pre-defined [header format \(i.e. PCI\)](#), and well-known "idle" patterns which are used to occupy the link when there is no data to be communicated.

A packet network equipment discards the "idle" patterns between packets and processes the entire packet as one piece of data. The equipment examines the packet header information (PCI) and then either removes the header (in an end system) or forwards the packet to another system. If the out-going link is not available, then the packet is placed in a queue until the link becomes free. A packet network is formed by links which connect packet network equipment.

There are two important benefits from packet switching:

The first and most important benefit is that since packets are short, the communication links between the nodes are only allocated to transferring a single message for a short period of time while transmitting each packet. Longer messages require a series of packets to be sent, but do not require the link to be dedicated between the transmission of each packet. The implication is that packets belonging to other messages may be sent between the packets of the message being sent from A to D. This provides a much fairer sharing of the resources of each of the links.

Another benefit of packet switching is known as "pipelining". Pipelining is visible in the figure above. At the time packet 1 is sent from B to C, packet 2 is sent from A to B; packet 1 is sent from C to D while packet 2 is sent from B to C, and packet 3 is sent from A to B, and so forth. This simultaneous use of communications links represents a gain in efficiency, the total delay for transmission across a packet network may be considerably less than for message switching, despite the inclusion of a header in each packet rather than in each message.

Note: additional diagrams for all answers available on course web site.

Course co-ordinator

G Fairhurst

Scrutineer

J Watson