These notes are intended to help demonstrators answer questions during tutorials - they do not provide full/complete answers, but do illustrate the key points of a valid answer.

TUTORIAL 1: Origins of Ethernet

(a) The IEEE 802.x family of LAN specifications support many physical media, explain with the aid of diagrams the differences between 10B5 and 10B2 media.

Both cable technologies were standardised by the IEEE and both use a coaxial cable as the physical medium.

10B5 cabling was used for backbone connections and is now fairly uncommon. It uses thick (low loss) coaxial cable to form a shared bus. Up to 100 transceivers may be connected to a single bus. This type of cable can be difficult to install (due to its weight, large 0.5 diameter, and constraints on minimum curvature). Its key advantage is the extended transmission distance and high noise immunity. A NIC is normally connected via an in-line transceiver with N-type connectors or a vampire tap that pierces the cable insulation.

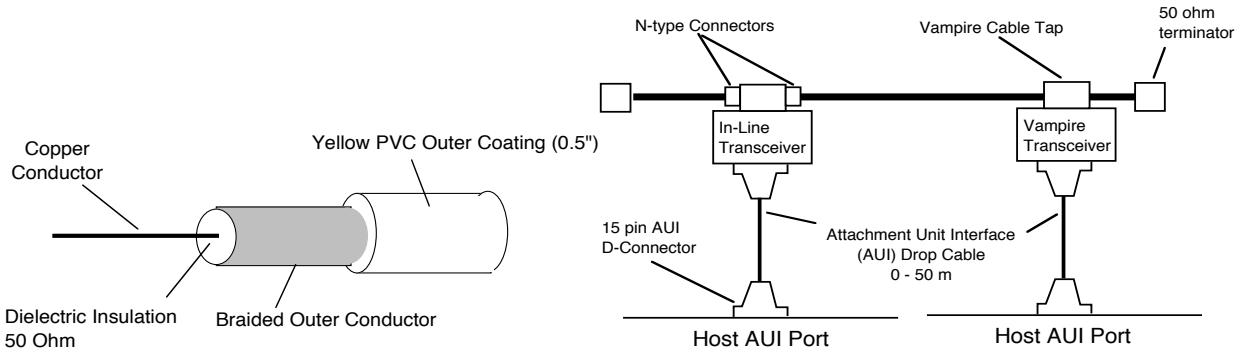

The later 10B2 specification in contrast was used for office or building cabling and was restricted to 185m of RG58/U cable. Like 10B5, this uses a coaxial cable to form a shared bus, but uses the BNC connector in place of the higher specification N-series connector. Up to 30 devices (including repeaters) can be connected to a cable segment.

The two key advantages of 10B2 are cheaper transceivers (A NIC is normally connected directly using a T-piece connector attached to the NIC but could be remotely connected by an AUI drop cable and external transceiver.) and the ease of installation. 10B2 cables can easily be unplugged and additional equipment added as needed.

TUTORIAL 2: Ethernet Frames

(a) How does a manufacturer decide which address to use?

The MAC address space architecture is flat, composed of two parts - the manufacturer part (at the start), called the Organisationally Unique ID, OUI and the system ID (at the end). Addresses are purchased from the IEEE Registry in blocks, normally one full OUI - the purchaser can then allocate an address which starts with their allocated OUI.

see: http://www.erg.abdn.ac.uk/users/gorry/course/lan-pages/mac-vendor-codes.html

Note: Normally manufacturers choose not to start with 00001, so that don't give away that this is the first Ethernet board that they have every made!

(b) Why is the first bit never set in an Ethernet source address?

The source address is a 6-byte source address, which is set to the sender's globally unique MAC address. It is sometimes called the NIC hardware address. The source address field of a frame may be used by to identify the sender in a network analyser, but usually other mechanisms are used by the network layer of computers (e.g., the arp protocol - covered in the IP part of this course).

Note: Its function is to allow address-learning by a bridge. The first bit (i.e. least significant bit of the first byte) indicates that an address is a broadcast/multicast frame. As a destination address this means the frame may be received by multiple recipients. A broadcast/multicast address (that sets the lsb of the first byte an Ethernet address) is hence not a valid identifier for the sender of a frame!

http://www.erg.abdn.ac.uk/users/gorry/course/lan-pages/mac-vendor-codes.html

(c) What is the Carrier Sense Circuit in the transceiver used for? Explain the "CS" and "CD" functions.

Ethernet uses CSMA. When a node has data to transmit, the node first listens to the cable (using a transceiver) to see if a carrier (signal) is being transmitted by another node. This may be achieved by monitoring whether a current is flowing in the cable (each bit corresponds to 18-20 milliAmps (mA) in 10B5). The Ethernet transceiver contains the electronics to perform this detection (labelled CS in the figure).

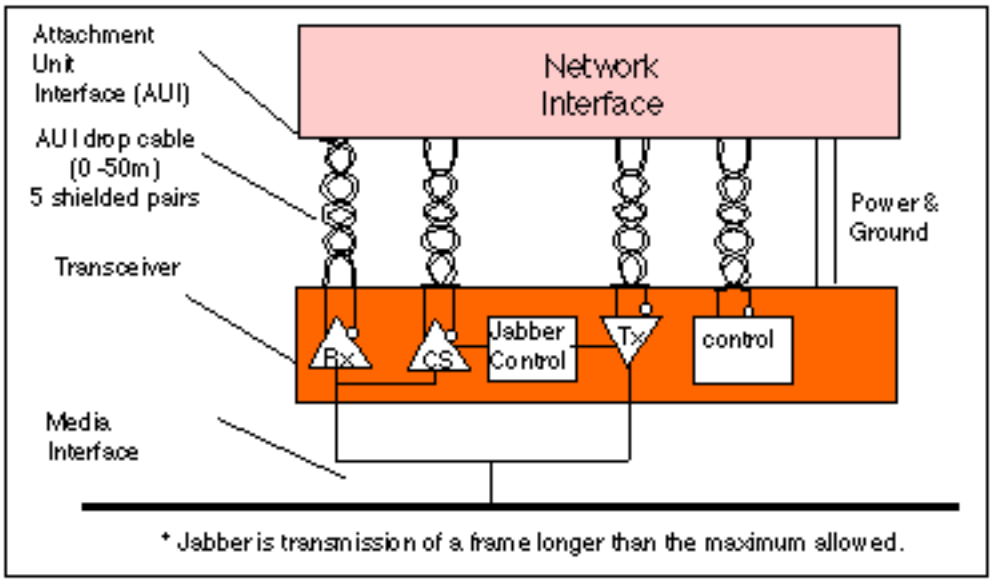

The individual bits are sent by encoding them with a 10 (or 100 MHz for fast Ethernet) clock to generate a pair of Manchester-encoded bauds. Frames are only started when no carrier is observed (i.e., no current is detected by the CS circuit) and the physical medium is therefore idle.

Note: Carrier Sense alone is unable to prevent two nodes transmitting at the same time. If two nodes try transmit within 51.2 µS, then both could see an idle physical medium (i.e., neither would see the other's carrier signal), and both will conclude that no other node is currently using the segment. In this case, both will then decide to transmit and a collision will occur. The collision will result in the corruption of the data being sent, which will subsequently be discarded by all receivers, since a corrupted Ethernet frame will not carry a valid 32-bit MAC CRC at the end.

This circuit is also used for collision detection.

Once a node starts transmission, it continues to monitor the cable during the frame transmission (using the CS circuit). If it observes an excess current above what it is generating, i.e. > 24 mA, it stops transmission immediately. This is called collision detection. Once a collision is detected the node continues instead to transmit a 32-bit jam sequence. The purpose of this sequence is to ensure that any other node which may currently be receiving this frame will receive the jam signal in place of the correct 32-bit MAC CRC, this causes the other receivers to discard the frame due to a CRC error.

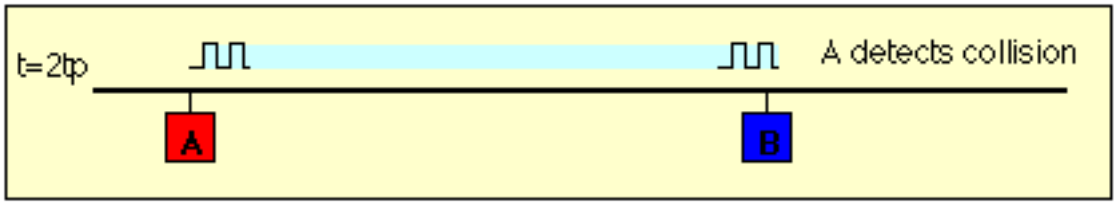

To ensure that no node may completely receive a frame before the transmitting node has finished sending it, Ethernet defines a minimum frame size (i.e. no frame may have less than 46 bytes of payload). The minimum frame size is related to the distance that cable may span, the type of media being used 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.

When two or more transmitters each detect a corruption of their own data (i.e. a collision), each responds in the same way by transmitting the jam sequence. At time t=0, a frame is sent on the idle medium by computer A. A short time later, computer B also transmits. (In this case, the medium, as observed by the computer at B happens to be idle too). After a period, equal to the propagation delay of the network, the computer B detects the other transmission from A, and is aware of a collision, but computer A has not yet observed that computer B was also transmitting. B continues to transmit, sending the Ethernet Jam sequence (32 bits).

Note: After one complete round trip propagation time (twice the one way propagation delay), both computers are aware of the collision. B will shortly cease transmission of the Jam Sequence, however A will continue to transmit a complete Jam Sequence. Finally the cable becomes idle.

see: http://www.erg.abdn.ac.uk/users/gorry/course/lan-pages/csma-cd.html

(d) What is Ethernet Capture?

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

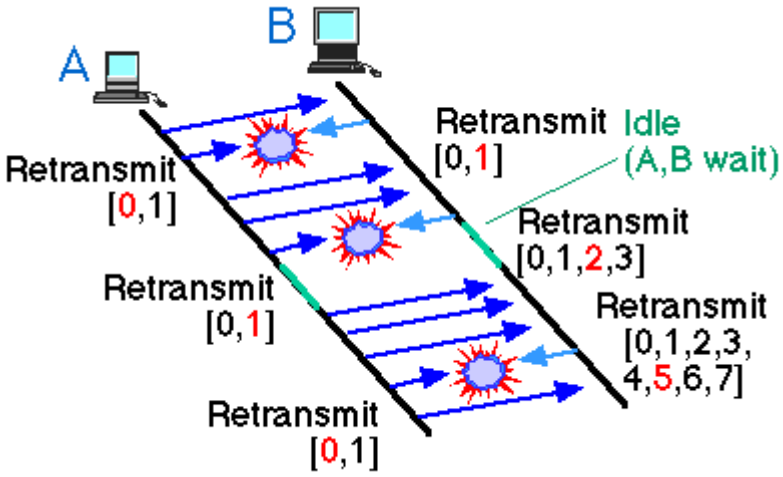

Computer A dominates computer B. Originally both computers have data to transmit. A transmits first. A and B then both simultaneously try to transmit. B picks a larger retransmission interval than A (shown in red) and defers. A sends, then sends again. There is a short pause, and then both A and B attempt to resume transmission. A and B both back-off, however, since B was already in back-off (it failed to retransmit), it chooses from a larger range of back-off times (using the exponential back-off algorithm). A is therefore more likely to succeed, which it does in the example. The next pause in transmission, A and B both attempt to send. This fails in this case, node B further increases its back-off and is now unable to fairly compete with node A. A similar situation may arise when many sources compete with one source which has much more data to send. Under these situations nodes may be "locked out" of using the medium for a period of time.

Note: A way to eliminate the problem is for a busy sender to pause from time to time, allowing other nodes free access to the medium. The use of full duplex cabling eliminates the problem in 100BT and GBE - discussed later in the course.

TUTORIAL 3: Ethernet Transmission

(a) Sketch the format of an Ethernet frame carrying 1000 B of data. Your answer should show all the protocol fields (headers) present in the frame.

A frame is formed by the Ethernet MAC layer PDU is prefixed by a 14B MAC header -> A 4B CRC is added -> An 8 B preamble added -> Finally an IFG delay precedes the start (9.5-10.5 μ S)

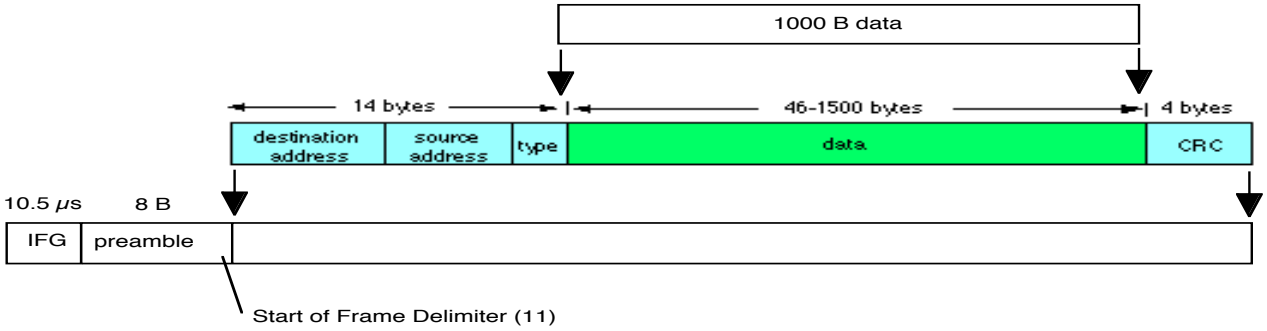

1000 B of data (PDU)
Total size = IFG+8+14+1000+4 in bytes (with 8 bits/byte).
- An answer needs to note the presence of the IFG, and could also assume a sensible value, e.g. 10.5 μS.

http://www.erg.abdn.ac.uk/users/gorry/course/lan-pages/enet-calc.html

Note: Minimum frame carries 46 B of data, i.e. a total length of 64B (excluding preamble).
Maximum frame carries 1500 B of data, i.e. a total length of 1518 B (excluding preamble).

(b) Sketch an outline block diagram of the process by which a byte is Manchester encoded by an Ethernet processor prior to transmission by the physical layer. Use the example of sending a byte with the hexadecimal value of 0x57.

The transmit process is:

Each Byte is serialised (i.e. shifted out of a register), least significant bit (lsb) first.
Each bit is then encoded into two bauds using Manchester encoding: a 1 becomes 01; and a 0 becomes 10.
The stream of bauds is transmitted using the internal transmit clock within each NIC.
This stream of bauds is passed to the transceiver (e.g., via an AUI cable), where an inverting amplifier generates the cable signal.

The data sent is:

Hexadecimal value 0x57 is 0101 0111.
Serialises to (lsb first): 1110 1010
Encodes (in Manchester encoding): 01 01 01 10 01 10 01 10
Appears on the cable (as a voltage/current inverted): 10 10 10 01 10 01 10 01.

The baud period is 0.5 x 10-7 (2 Manchester encoded bauds per data bit).

see: http://www.erg.abdn.ac.uk/users/gorry/course/phy-pages/man.html

Note: In 10B5 each baud is actually represented as a current along the 50 Ohm coaxial cable.

(c) Explain the function of the Ethernet Preamble

The preamble provides a simple signal before the start of the MAC header that serves to synchronise the receivers. The preamble when encoded using Manchester encoding generates a square wave - simplifying the process of synchronisation. This provides a signal for the Receiver DPLL to acquire phase synchronisation. This is necessary, since a receiver is unaware which NIC originates each frame, and each NIC has a separate and independent clock that it uses to send the Manchester encoded data.

see: http://www.erg.abdn.ac.uk/users/gorry/course/lan-pages/mac.html

(d) Plot the waveform which you would observe on an oscilloscope when the first byte of the preamble is sent along an Ethernet coaxial cable.

The preamble sequence is 10101010 (alternating one and zero data bits)
Serialises to (lsb first): 0101 0101
Encodes to (in Manchester encoding): 10 01 10 01 10 01 10 01 (two bauds/bit)
This appears on the cable (as a voltage/current inverted): 01 10 01 10 01 10 01 10

In 10B5 each baud is actually represented as a current along the 50 Ohm coaxial cable.

The result is a square wave of period 1 x 10-7.

A valid answer must show a vertical axis as Volts or Amps; and a horizontal axis with time.

see: http://www.erg.abdn.ac.uk/users/gorry/course/lan-pages/mac.html

(e) Explain the function of the Ethernet Cyclic Redundancy Check (CRC)

The CRC is a mathematical hash/signature that is constructed by the NIC by encoding all the bits within a MAC frame to be transmitted, including the MAC header, but excluding the preamble. This is transmitted as the last four bytes of each Ethernet frame.

At a receiving NIC this is used to determine if a received frame is identical to the frame that was sent. The receiver does this by separately encoding the all the bits within a received MAC frame, including the MAC header, but excluding the frame CRC. The receiver-calculated CRC is compared to the CRC within the frame. If the two match, the frame is assumed to be correct, i.e., it has not been modified in transmission. This is necessary to ensure truncated frames (generated by CSMA/CD Collisions) are discarded. It also protects against timing errors from the DPLL and bit errors (inversions) during transmission.

see: http://www.erg.abdn.ac.uk/users/gorry/course/dl-pages/crc.html

NOTE: The network layer (IP) - covered in part 2 - also includes an integrity check, known as a checksum. Although this has a similar role in detecting errors, the way the checksum is calculated is different, since this is used to protect from different errors.

TUTORIAL 4: Collision Domains and Repeaters

(a) The Ethernet standard states that each frame should carry two addresses: What is the Ethernet destination address used for?

The Destination Address is a 6B number that uniquely identifies the intended recipient (node) within a LAN. This address may take three forms : A unicast address is globally unique (within the whole world). The all 1's broadcast address indicate a frame to be received by all NICs. A multicast address indicates a frame that will be carried to all parts of a LAN and carries an address to indicate the multicast group being used. Receivers always receive frames with a unicast address matching their own source address, or the broadcast address. Receivers may optionally register to receive frames with specified multicast addresses.

(b) Explain with the aid of diagrams the key features of 10BT cables and suggest two situations where optical fibre is preferable to copper cabling.

10BT cabling uses a RJ-45 connector and 100 Ohm unshielded twisted pair cabling. This connects a NIC directly to a hub that acts as a repeater. The maximum distance of a 10BT link is 100 m. It is normally used to connect work groups of users, sometimes by wiring an entire floor with outlets to a work area.

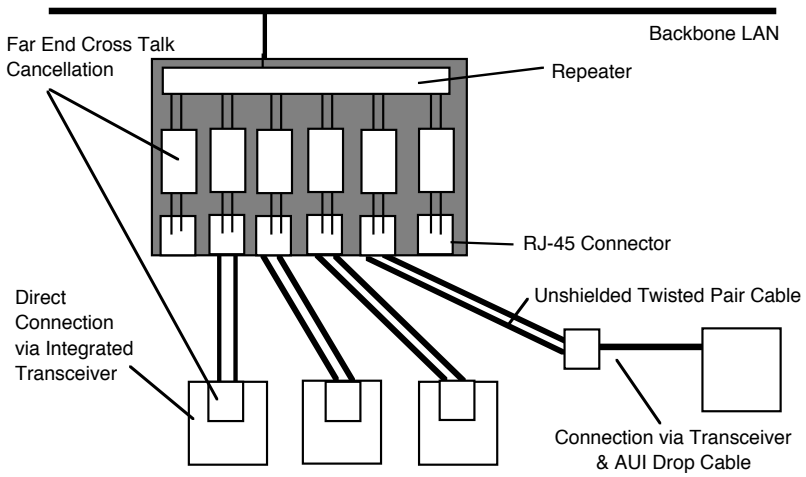

Fibre LAN segments may be used to connect much greater distances compared to their copper counter parts. Fibre also has the added advantages that it does not provide an electrical path between the sender and receiver. It may therefore be used to connect buildings which could potentially operate on different mains phases, provides immunity form lightning strike, and is suitable for use in (industrial) environments where there is a high level of background noise. In some applications, fibre is used to provide a very high level of security - i.e., because it is very difficult to monitor or tap the cable to observe the frames being carried.

(c) What is the 5-4-3 repeater rule?

Not more than 5 segments in series
Not more than 4 repeaters
Not more than 3 active segments

The rules are governed by the end-to-end transmission delay < 2 x slot time. The longest path can have no more than 5 maximal length segments - each adds propagation delay (up to the maximum permitted cable length). This implies the need for 4 repeaters - which adds processing delay for the preamble. However, a directly connected transmitter can retain frequency lock and never experiences collisions, this is known as a passive segment - the delay through a repeater on a passive segment (e.g. 10BT, 10BF) is therefore much shorter.

(d) Here are 5 networks, can you explain whether they pass the 5-4-3 Repeater rule and if not, why not?

Discuss your answer in the tutorial

1

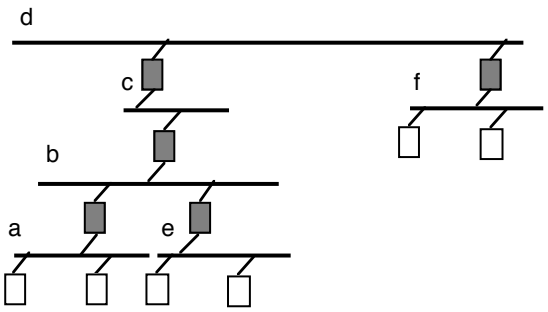

2

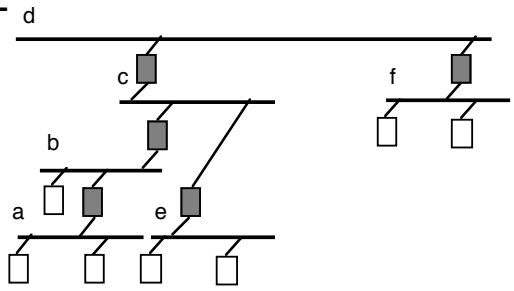

3

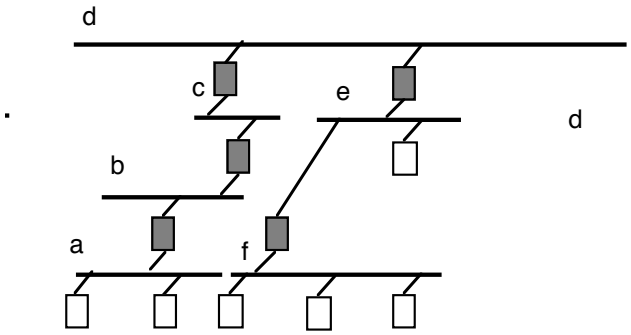

4

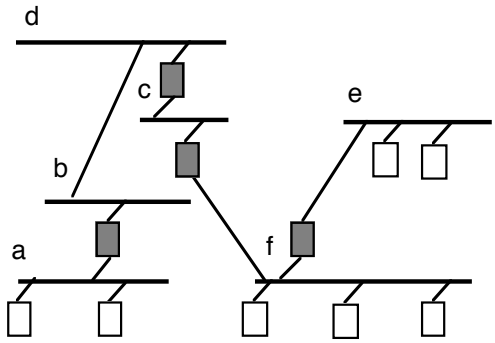

5

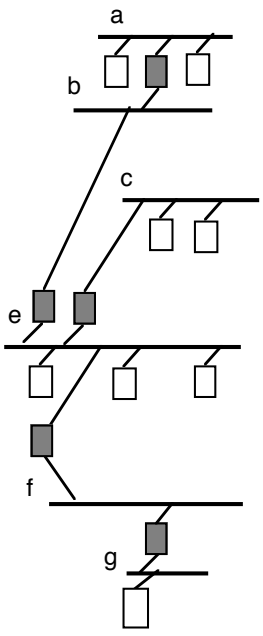

TUTORIAL 5: Bridges and Switches

(a) Explain in detail the operation of an Ethernet bridge, and describe how the bridge filters frames which need not be forwarded between Ethernet segments.

A bridge is a LAN interconnection device that operates at the (data) link layer. It is used to join two LAN segments (A,B), 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. The Ethernet frame format consists of two 6 byte addresses and a one byte protocol ID / IEEE length field. The destination address allows a frame to be sent to single or groups of systems.

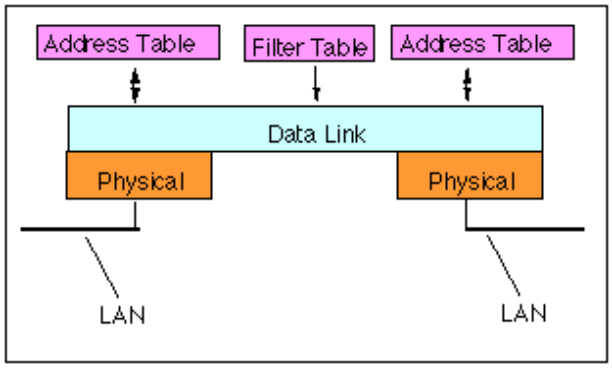

The bridge learns the addresses that belong to the NICs, by observing the source address values (in the MAC headers) of Ethernet frames received at each bridge port. This is called "learning". For example, consider three systems with source addresses X, Y that are connected to port A on the left of the figure, while one system has an address of Z which is associated with port B, on the right of the figure.

The bridge discovers these addresses by observing which port has received a frame and stores the source address and the associated interface (port) in the address table.

It then uses this address information to determine which frames need to be forwarded by the bridge by looking up the destination address of all frames, and forwards the frame to the port associated with the destination in the address table, except for the case were the port entry associated with address is the same as the port on which the frame was received. In the example, frames received on port A with a source of X and destination of Y will be received and then discarded, since the computer Y is directly connected to the Port A, whereas frames from X with a destination of Z are forwarded by the bridge onto Port B. A system administrator can override the normal forwarding by inserting entries in a filter table to inhibit forwarding between workgroups (e.g., to provide security).

(b) What is the role of the Ethernet Source Address?

The frame source address is set by each NIC prior to sending the frame, as a copy of the unicast source address stored in the NIC. Bridges can use this information to learn the set of NICs that may be reached via a bridge port, and thereby build a dynamic address table to use to intelligently forward frames.

(c) Explain the concept of a Collision Domain

A collision domain is formed from a cable segment and may include other cable segments connected via hubs/repeaters. It defines that part of a LAN that is governed by the CSMA/CD MAC protocol. If two nodes in a collision domain send within one Ethernet Slot Time of each other, then a collision will occur, and both nodes will need to retransmit the lost frames. A bridge/switch places all ports in different collision domains and can split larger networks into isolated areas each governed by the CSMA/CD protocol.

TUTORIAL 5.1 Bridge v Repeaters

(a) **Four computers (W,X,Y,Z) are connected by 3 Ethernet segments (A,B,C) using a Repeater (R) and a Bridge/Switch (B).**

Each computer has the ability to send frames to unicast, broadcast, and multicast addresses. For example, W can send a frame to X (W -> X) using the repeater (R) which connects segments A and B. In this case, the bridge (B) does not forward the frame, so it can not be received by computers Y or Z which are connected to LAN C.

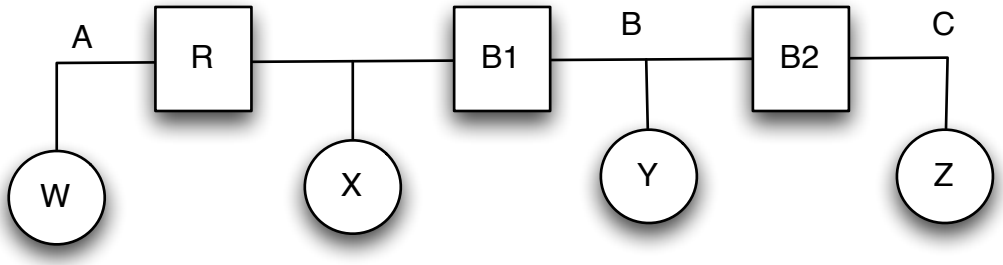

Which computers receive (at the network level) the following frames? (Make sure your answer also shows which LAN segments carry each frame.)

Frame	Received				LAN			
	W	X	Y	Z	A	B	C	
W-> BCast	Y	Y	Y	Y	Y	Y	Y	- bcast, sent to all segments
X-> Z	-	-	-	Y	Y	Y	Y	
Y-> Z	N	N	-	Y	N	Y	Y	
Y->Bcast	Y	Y	Y	Y	Y	Y	Y	- bcast, sent to all segments

(items marked "-" are received at MAC level, and filtered because destination address does not match)

(b) The computers W and X have MAC addresses of W = 0x00102030 and X = 0x00102040. They are both also members of the multicast group 0x23. Sketch the MAC header for a multicast frame sent from X to the multicast group 0x23.

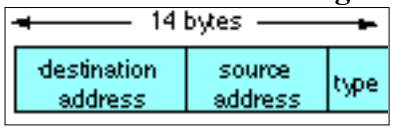

Destination = 0x 01 00 00 00 23; Source = 0x 00 10 20 40; Type = two bytes (e.g., 0x0800 indicates IP).
Note: the least significant bit of destination address indicates broadcast or multicast.

(c) Which LAN segments carry a multicast frame?

All segments normally carry multicast traffic. A simple bridge treats multicast and broadcast frames the same. It is unable to identify which computers require the multicast packets, because it is unaware of the multicast addresses that have been registered in each attached NIC. A more powerful enterprise switch can include additional logic and processing to adaptively forward multicast frames.

(d) Which of the following may be used to extend a LAN to allow computers to be connected by more than 5 cable segments?

(i) Transceiver -NO (ii) Repeater -NO (iii) Hub -NO (iv) Bridge -YES (v) Switch -YES
NO = Because forms one collision domain. YES = Each bridge/switch port forms one collision domain.

TUTORIAL 5.3: Bridges and Switches - Supplementary Question

A small Local Area Network (LAN) is shown in the figure:

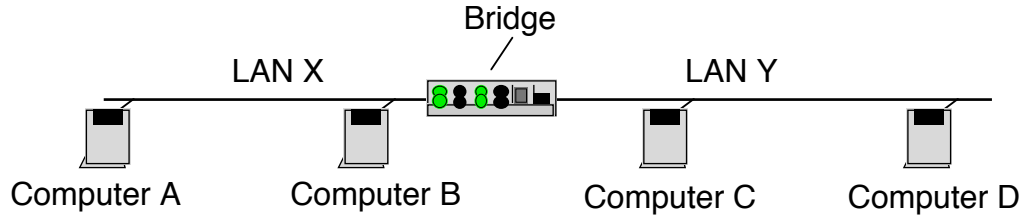

Four computers, A, B, C and D connected to a LAN. The LAN is formed from two shared Ethernet segments joined via a bridge.

(a) The computer A sends three simultaneous Unicast file transfer packets. each to computers B, C, and D. Calculate the size of a frame, given that it carries 1032 B of IP network payload data. Using this information, calculate the Utilisation of LAN X, assuming that the transmission continues at 50 packets per second to each of the three destinations.

The utilisation is defined aas 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.

All packets travel on LAN X.

Each packet has the following protocol headers (PCI):
MAC-Preamble (8 bytes) + MAC Header (14 bytes) + IP Header (20 bytes) + UDP(8bytes) + UDP Payload (1024 bytes) + CRC-32 (4 bytes)

Assume 10 Mbps Ethernet operation. The inter-frame gap (number of bits equivalent to 10.4 μS at 10 Mbps) may also be considered as overhead, which will yield a slightly higher answer.

Total size= (8+14+20+8+1024+4)x8 = 8624 bits

50 messages sent per second to 3 computers = 150 UDP messages/second

Total utilisation = % bits per second/transmission rate = (8624 x 150 / (10 x 1000 000))x 100 =13%

(b) What is the utilisation on LAN Y?

This is different, since we consider unicast transmission, the bridge will not forward packets from A to B. It will forward packets from A to B and C.

The utilisation on LAN Y is therefore: 2/3 of 13%, i.e. 9%.

(c) How does Multicast transmission differ from Unicast transmission?

Unicast is the term used to describe communication where a piece of information is sent from one point to another point. In this case there is just one sender, and one receiver. Unicast frames are sent from a single source to a specified destination, is still the predominant form of transmission on LANs and within the Internet. All LANs (e.g. Ethernet) and IP networks support the unicast transfer mode, and most users are familiar with the standard unicast applications (e.g. http, smtp, ftp and telnet) which employ the TCP transport protocol. http://www.erg.abdn.ac.uk/users/gorry/course/intro-pages/uni-b-mcast.html

Using unicast, simultaneous delivery of high quality video to each of a large number of receivers will exhaust the capability of even a high bandwidth network with a powerful server. This poses a major scalability issue for applications that require sustained high bandwidth.

One common example of an application that can use multicast is a video server sending out networked IPTV channels. Multicast can significantly ease scaling to larger groups of clients is to employ multicast networking. Multicast is the term used to describe communication where a piece of information is sent from one or more points to a set of other points. In this case there is may be one or more senders, and the information is distributed to a set of receivers (there may be no receivers, or any other number of receivers - the sender simply does not know). Where there is a common need for the same data by a group of clients, multicast may save bandwidth (using 1/N of the capacity compared to N separate unicast frames).

At the network layer, the format of IP multicast packets is identical to that of unicast packets and is distinguished only by the use of a special class of IP destination address that denotes a specific multicast group (http://www.erg.abdn.ac.uk/users/gorry/course/inet-pages/ip-address.html). Unlike broadcast transmission, multicast receivers receive a stream of packets only if they have previously choose to do so (by joining the specific multicast group address). Membership of a group is dynamic, controlled by the applications. IP multicast packets are sent in link-layer frames with an Ethernet multicast destination address.

At the link layer, Ethernet multicast uses a destination address with the lab of the first byte set to 1 to indicate a broadcast/multicast address. The destination address (all 1 s) may also identify a broadcast packet (to be sent to all connected computers). The remainder of a multicast address indicates the group ID to which the message is sent. A NIC registers each corresponding multicast MAC address required with the NIC to receive multicast frames with this address.

Note: The hardware address is also known as the Medium Access Control (MAC) address. Each computer network interface card is allocated a globally unique 6 byte source address when the factory manufactures the card (stored in a (EE)PROM), used for all frames it creates, and receives all packets which match this hardware address, the LAN broadcast address, or one of the registered multicast addresses.

(d) Calculate the utilisation for LAN Y when the file is sent using multicast packets instead of the unicast packets used in section (a).

Bridges always forward multicast packets. Only one multicast packet is sent to each destination.

The utilisation is therefore:
Total utilisation = 8624 x 50 / (10 x 1000 000 x 100) =4.3%

TUTORIAL 6: Faster Ethernet

(a) In the context of Fast Ethernet explain how the following sequence of bits {1 0 0 1 1 0 } are encoded using Multi-Level Threshold, MLT-3 line encoding.

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.

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 have the most transitions, this ensures clocking information is present in the signal (for locking the receiver DPLL). This results in a bandwidth increase 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.

Fast Ethernet Line Interface for 100 BT

{1 0 0 1 1 0 } are encoded as: +++0-- (assuming a zero start and positive waveform)

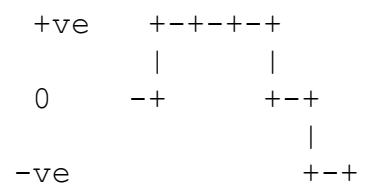

(b) What is the purpose of 4b/5b encoding?

This primary purpose of this encoding is to ensure sufficient transitions in the physical layer that the receiver DPLL can recover the clock associated with the data. This also is used to encode the start and end of each frame.

(c) Explain the operation of the physical layer used by a Gigabit Ethernet Network Interface Card (NIC).

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 x 1 = 1.25Gbps. 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 by mapping pairs of bits using a 5-level Pulse-Amplitude-Modulation (PAM).

TUORIAL A.1: IP

(a) Define the following terms:
(i) The IP Version Number

Informs receiver which version of IP is being used. In today's Internet this is nearly always v4 - but in the future you may also expect to see packets with a different header corresponding to IPv6. See: http://www.erg.abdn.ac.uk/users/gorry/course/inet-pages/ip-packet.html

(ii) The IP Network Number

An address is a data structure understood by a network which uniquely identifies the recipient within the network. Each system in the Internet is assigned a a 32 bit value (IP address) consisting of two parts, the network part (identifying the network to which the computer is attached) and the host part (which identifies the host within the local network). Systems are also assigned a 32-bit mask value (called the "netmask" or the "subnet mask") which is used to separate the two parts of the address. This mask contains a logical 1 in each bit position of the network part of the address. The IP network is identified as the bit-wise logical AND of the netmask and the 32-bit IP address. By using the logical AND operation on another IP packet destination address, a system may determine whether a packet to be sent is destined for the same network (and can be sent directly) or a remote network (hence requiring it to first be sent to an intermediate router).
See: http://www.erg.abdn.ac.uk/users/gorry/course/inet-pages/ip-address.html

(iii) The Network Layer Address

An internet address is a unique 32-bit value described by four decimal numbers (one for each byte) separated by dots. The numbers have the form: 139.133.204.99 where the left part indicates the network number (to which the host is attached) and the right part the individual address of the host (in the network). See: http://www.erg.abdn.ac.uk/users/gorry/course/inet-pages/ip-address.html

(b) Describe in detail how a computer connected to an Ethernet LAN determines its own Ethernet address and how it determines the Ethernet address of other computers with which it needs to communicate.

Systems (e.g. router, computers) are connected to an Ethernet LAN using a NIC. The NIC contains an Ethernet processor which implements the MAC protocol used to access the Ethernet medium (cable) and the digital logic required to transmit Ethernet frames using a transceiver. The Ethernet MAC header uses two hardware addresses to identify the source and destination of each frame sent by the Ethernet. The destination address (all 1s) may also identify a broadcast or a multicast packet. Each computer network interface card is allocated a globally unique 6 byte address when the factory manufactures the card (stored in a PROM or EEPROM). The local address read from the PROM is the normal source address used by an interface. A computer sends all packets that it creates using this MAC source address, and receives all packets which match this address or the broadcast address (and also any that match a registered multicast address when these are configured)

When first started, a computer is unaware of the MAC addresses of other computers and routers connected to a LAN. It must use address resolution to find the destination address of any IP network devices with which it wishes to communicate. The remote MAC address is resolved using a protocol that sends a piece of information from the local computer to a server process executing on a remote (target) computer. The information received by the target allows it to uniquely identify the IP address for which the MAC address was required. Address resolution is completed when the client receives a response from the target containing the required address. To reduce the number of address resolution requests, a client normally caches resolved addresses for a (short) period of time. When a frame is to be sent to a destination address contained in the

cache, the computer does not need to query to find the required address, instead it copies the value directly from the cache. This is important, since it prevents an address resolution query prior to each transmission.

The Ethernet address is a link address and is dependent on the NIC that is used. IP operates at the network layer and is not concerned with the addresses of individual NICs. The arp client and server processes operate on all computers using IP over Ethernet.

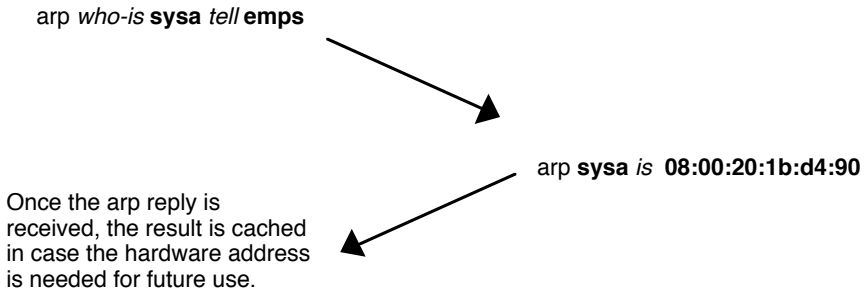

The above exchange shows increasing time on the vertical axis, and a sender (in this case emps) on the left, asking for the hardware address of a computer (sysa). The response by the destination (sysa) is shown by the arrow from right to left.

See: http://www.erg.abdn.ac.uk/users/gorry/course/inet-pages/arp.html

(c) Two computers, A and B, are connected by an Ethernet LAN. A has not previously sent any packets to the LAN. It sends 20B of data to B in a single IP datagram.

(i) What is the type of the first Ethernet frame sent by A?

(i) arp who-is B tell A

(ii) To which MAC address is the first frame sent?

(ii) arp requests are sent to the Layer 2 (MAC) broadcast address (ff:ff:ff:ff:ff:ff)

(iii) What is the IP source address of the first IP packet received by B?

(iii) A

See: http://www.erg.abdn.ac.uk/users/gorry/course/inet-pages/arp.html

TUTORIAL A.3: ARP - Supplementary Question

(a) An Ethernet Local Area Network (LAN) connects two workstations (A and B). The following information is provided about the IP interfaces of the computers connected to the LAN, also giving the hardware address (ha, or MAC address) for each interface.

A: IP = 139.101.1.7; h.a.= 08:00:20:02:b7:f9
B: IP = 133.101.1.63; h.a.= 08:00:20:ff:00:00

The computer A has not previously sent any packets to the Ethernet LAN. The user at A runs a program which sends one IP packet and has a destination address of B. Explain the hardware address values in the first frame which is sent by the computer A.

The user at A sends 20 bytes of data to B in a single IP packet. A has not previously sent any packets to the LAN. This implies that the arp cache is empty and that A will need to send an arp request prior to the transmission of any unicast packet to the Ethernet. The Ethernet interface software driver of computer A stores the packet in a temporary buffer (it has no MAC destination address to use) and broadcasts an arp request to on the Ethernet network, requesting that B replies with its own hardware address.

The Ethernet frame sent by A is: arp who-has B 133.101.1.63 (?) tell A 139.101.1.7 (08:00:20:02:b7:f9) This is sent to the broadcast address (ff:ff:ff:ff:ff:ff)

The hardware source address for the arp request is the NIC's own Ethernet MAC address (08:00:20:02:b7:f9), the destination MAC address is the Ethernet broadcast (ff:ff:ff:ff:ff:ff). The interface to computer B receives this request, and recognise the target IP address. It generates a reply to the originator of the arp request (A). When the reply is received, it indicates Bs hardware address (08:00:20:ff:00:00). It then forwards the stored IP packet (original message) using an Ethernet frame directly addressed to the MAC address of computer B.

(b) What are the hardware source and destination addresses and the source and destination IP addresses of the first network layer packet which is received by B?

Every computer interface is allocated a unique IP address. This identifies the computer in the network (the Internet). Each IP packet carries a pair of addresses in the 20 byte network layer header. When a computer sends an IP packet, it inserts its own 4 byte address in the IP source address field. The IP destination address indicates the recipient end system and is examined by routers within a network to determine a route to this destination. An end system receives all packets which match its own IP address.

Ethernet uses link frames that carry a 14 byte header. The MAC header identifies the sender and receiver within the local LAN (formed from Ethernet segments, repeaters and bridges). The addresses relate to NICs directly connected to the LAN. The first Ethernet frame received by B from A is an arp request, but this does not contain an IP packet. The first IP network packet containing information from a higher layer protocol is the 20 B message that follows. This packet has a source hardware address of 08:00:20:02:b7:f9 and a destination hardware address of 08:00:20:ff:00:00. The packet has a source IP address of 139.101.1.7 and a destination IP address of 139.101.1.63.

(c) Calculate the time taken to transmit one ARP packet using an Ethernet LAN operating at 10 Mbps. Your answer should first calculate the total size of the ARP packet including ALL overhead introduced by each layer during the transmission process.

MAC+ ARP request + CRC-32 = 14 + 28 + PAD + 4 = 64 B
Note minimum Ethernet PDU is 60 B (including MAC header, excluding CRC-32)
Total time to send = 10.5 micro secs + ((64+8)x8/10 000 000) = 10.5+57.6 micro secs
= 1068.1 microsecs.

TUTORIAL B.1: IP

(a) An end to end connection may be checked in an IP internet using the ping program which uses the Internet Control Management Protocol, ICMP. Describe the operation of the ICMP echo request and ICMP echo reply to perform this check, and how this may measure the round trip delay across the network.

The Internet Control Message Protocol (ICMP) protocol is classic example of a client server application. The ICMP server executes on all IP end system computers and all IP intermediate systems (i.e routers). The protocol is used to report problems with delivery of IP packets within an IP network. It can be sued to show when a particular End System (ES) is not responding, when an IP network is not reachable, when a node is overloaded, when an error occurs in the IP header information, etc. The protocol is also frequently used by Internet managers to verify correct operations of End Systems (ES) and to check that routers are correctly routing packets to the specified destination address.

An ICMP message includes an 8-bit type code which identifies the types of message. In this case two types of message are involved the ECHO request (sent by the client) and the ECHO reply (the response by the server). Each message may contain some optional data. When data are sent by a server, the server returns the data in the reply which is generated. ICMP packets are encapsulated in IP for transmission across an internet. The "ping" program contains a client interface to ICMP. This may be used by a user to verify an end to end connection is operational. The -s option of "ping" also collects some performance statistics (i.e. the measured round trip time and the number of times the remote server fails to reply. Each time an echo reply packet is received a single line of text is displayed. Each echo request packet contains a sequence number (starting at 0) which is incremented after each transmission, and a timestamp value indicating the transmission time. The text printed by ping shows the received sequence number, and the measured round trip time.

(b) Provide a step by step explanation using diagrams to show the way an Ethernet network interface card and the network layer protocol process an IP frame received from an Ethernet transceiver.

The Ethernet controller verifies that the frame is:

- Not less than the minimum frame length & Not greater than the maximum length
- Contains a valid CRC at the end
- Does not contain a residue (i.e. extra bits which do not form a byte)

The frame is then filtered and accepted only if it is a broadcast frame or a multicast frame to a joined group address or is a unicast frame to the nodes own MAC address

The frame is then demultiplexed using the MAC packet type It is passed to the appropriate protocol layer (e.g.ARP, IP) Packets destined for IP have a type field of 0x0800.

IP processing checks the packet header:
Checking the protocol type =4 (i.e. current IP)
By verifying the header checksum
and checking the header packet length
The destination address is then checked and accepted if:

- It matches an IP address of the node then it is accepted
- It is a network broadcast packet
- It is a multicast packet to an IP multicast address (to a group being used)

If it is none of these, it is forwarded using the routing table (if the system is a router)
or discarded (if a an end system)

The IP protocol field is checked.
The complete packet is passed to the transport layer protocol.

(c) For each packet, specify which computers addresses are used for the source and destination address at both the link layer AND the network layer.

The address resolution protocol is a protocol used by the Internet Protocol (IP) network layer protocol to map IP network addresses to the hardware addresses used by a link protocol. The protocol operates below the network layer, and is used when IP is used over Ethernet. It is considered to below the network layer, because it is concerned with the details of a specific type of link - in this case Ethernet, whereas the network layer functions (e.g. fragmentation, ICMP) apply to all types of link.

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.

An Ethernet network uses two hardware addresses. The destination address may also identify a broadcast packet (to be sent to all connected computers) or a multicast packet (lsb of first byte set to 1) (to be sent only to a selected group of computers). The hardware address is also known as the Medium Access Control (MAC) address. Each computer network interface card is allocated a globally unique 6 byte address when the factory manufactures the card. A computer sends all packets which it creates with its own hardware source address, and receives all packets which match its hardware address or the broadcast address. When configured to use multicast, a selection of multicast addresses may also be received.

The Ethernet address is a link layer address and is dependent on the interface card which is used. IP operates at the network layer and is not concerned with the network addresses of individual nodes which are to be used. A protocol known as address resolution protocol (arp) is therefore used to translate between the two types of address. The arp client and server processes operate on all computers using IP over Ethernet. The processes are normally implemented as part of the software driver for the network interface card.

Frames:

MAC+ ARP request + CRC-32
src: percy-enet (x)
dst: Broadcast

MAC + ARP reply + CRC-32
src: james-enet (y)
dst: percy-enet (x)

MAC + IP + ICMP ECHO request + DATA + CRC-32
src: percy-enet (x)
dst: james-enet (y)

MAC + IP + ICMP ECHO reply + DATA + CRC-32
src: james-enet (y)
dst: percy-enet (x)

TUTORIAL C: Transport Layer

(a) What is the transport service provided by the transport layer?

The transport layer provides transparent transfer of data between end systems using the services of the network layer (e.g. IP) below to move data between the two communicating systems. It is said to perform "peer to peer" communication, with the remote (peer) transport entity.

The data communicated by the transport layer is encapsulated in a transport layer data and sent in a network layer packet. The network layer routers (i.e.transfer the packets, without decoding or modifying the content of the payload). In this way, only the peer transport entities actually communicate using the payload of the transport protocol.

The transport layer relieves the upper layers from any concern with providing reliable data transfer. It provides end-to-end control and information transfer with the quality of service needed by the application program. It is the first true end-to-end layer, implemented in all computers).

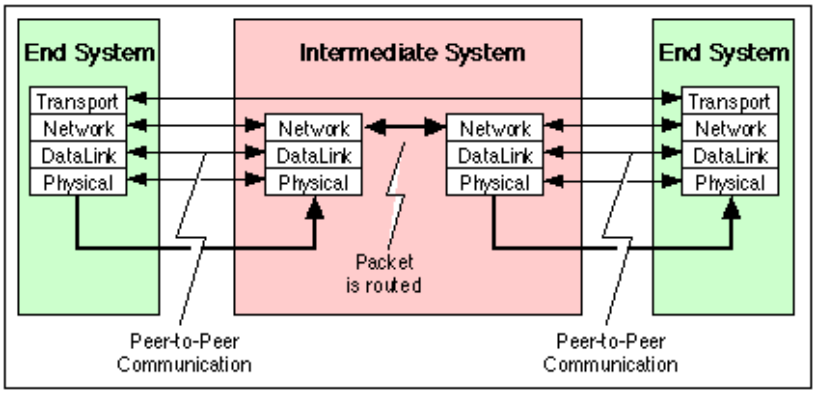

Figure showing peer-to-per communication at the transport layer

The Internet Protocol (IP) provides two transport layer protocols:
The User Datagram Protocol (Best Effort Datagram Service)
The Transmission Control Protocol (Reliable Service)

(b) Explain the function of each of the fields in the UDP packet header.

The UDP header consists of 8 bytes of data. 16 bits allocated to the source port ID which identifies the client process. The next 16 bits identify the destination port, in the case of packets from a client, this identifies the server process. A 16 bit value identifies the length of the payload carried by the message and the final 16 bits provide an optional checksum over the payload data. The checksum is normally used to validate the correct end-to-end transmission of the payload data.

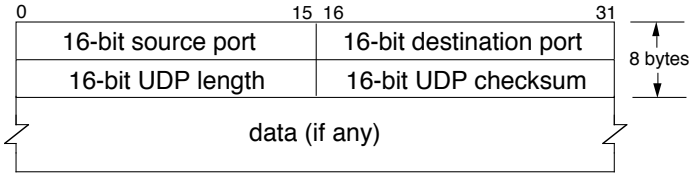

Figure showing the UDP protocol header at the transport layer

Port number (service) 2 bytes source + 2 bytes dest.
Checksum of Payload using 16 bit value - 2 bytes
Size of Payload 2 bytes

(c) A UDP packet containing 50 B of payload data is transmitted using IP over an Ethernet LAN. What is the total size of the frame transmitted on the Ethernet LAN

This is the sequence of headers:

Phy	Link	Network	Transport	Payload	Link	Physical
Preamble	MAC	IP	UDP	Data	CRC-32	IFG

Layer	Protocol	Size of PCI
Physical Layer	Ethernet (Phy)	Preamble: 8 including SFD *
Link Layer	Ethernet (MAC)	Frame Header: 6 B + 6 B + 2B
Network Layer	IP	20 B
Transport Layer	UDP	8 B
Payload	Data	50 B
Link Layer	Ethernet CRC-32	Trailer: 4 B
Total		96 B

Total size = 14+20+8+50+4 B = 96 B

TUTORIAL D: Packet Decodes

(a) Define the following terms

(i) IP Header Checksum

A 2's complement checksum inserted by the sender and updated whenever the packet header is modified by a router - Used to detect processing errors introduced into the packet inside a router or bridge where the packet is not protected by a link layer cyclic redundancy check. Packets with an invalid checksum are discarded by all nodes in an IP network.

(ii) UDP Checksum

A 16-bit 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.

(b) An Ethernet protocol analyser observes the following frame:

08 00 20 00 70 DF 08 00 20 01 62 F0 08 00 45 00 00 1E 4A 02 00 00 3C 17 84 53 8B 85 CC 16 8B 85 CC 13 06 1B 04 25 00 0A 00 00 42 42 00 00 00 00 ...

By decoding the hexadecimal bytes of this frame using the header chart supplied, determine the values for the MAC source address, the EtherType field; the IP and UDP Checksums.

(i) Ethernet Source Address = 08 00 20 01 62 f0

(ii) MAC Service Access Point = 08 00 (IP)

(ii) IP Header Checksum = 84 53

(iii) UDP Checksum = 00 00 (this specific value indicates that the checksum is not being used)

Note - not required - here is a full decode:

```
08 00 20 00 70 df      : MAC Destination address
08 00 20 01 62 f0      : MAC Source address
08 00                  : MAC Type (08 00 IP)

------ The above MAC header is for an Ethernet packet -----

45 00                  : IP Version (4) IHL (5)
                       : and Type of service (00)
00 1E                  : IP Total length
4a 02                  : IP Identification
00 00                  : IP Flags and Fragment offset
3c 17                  : IP Time to live(3c)
                       : and Protocol (17)
84 53                  : IP Header checksum
8b 85 cc 16    : IP Source address
8b 85 cc 13    : IP Destination address

—— The above header is the IPv4 packet —-
—- The following header is for an IP payload carrying UDP ------

06 1b                  : UDP Source Port
04 25                  : UDP Destination address
 00 0A                 : UDP Length (10 B)
00 00                  : UDP Checksum (disabled)
42 42                  : UDP Data  (2B)
                       : followed by padding
```

Remember there are lots more packet decodes on the course web site!

Note Decoding, Traceroute and DNS will be topics of assessment in a “practical” exercise!