Session 1998-99 Exam 1

EG / ES 3561 Worked Solutions.

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.

Dr Gorry Fairhurst

G.Fairhurst@eng.abdn.ac.uk

Mark

1 (a) The Ethernet Local Area Network (LAN) uses Carrier Sense Multiple Access wit Collision Detection (CSMA/CD) to share the transmission medium. In the context c CSMA/CD, what is meant by the following terms?

(i) Carrier Sense [3 marks]

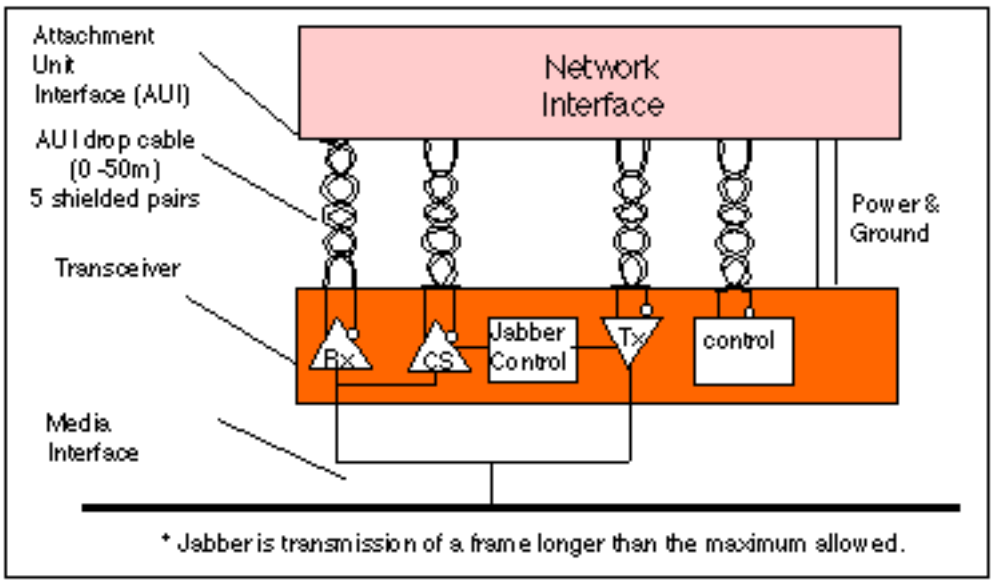

Ethernet uses a refinemer of ALOHA, known a CSMA, which improve performance when there is higher medium utilisatio When a node has data t transmit, the node first lis tens to the cable (using transceiver) to see if a carr er (signal) is being transmi ted by another node. Thi may be achieved by mon toring whether a current flowing in the cable (eac bit corresponds to 18-2 milliAmps (mA)). Th Ethernet transceiver cor tains the electronics to pe form this detection (labelle CS in the figure).

The individual bits are sent by encoding them with a 10 (or 100 MHz for fast Ethernet) clock using Manches ter encoding. Data is only sent when no carrier is observed (i.e. no current present) and the physical mediur is therefore idle.

However, this alone is unable to prevent two nodes transmitting at the same time. If two noes simultaneousl try transmit, then both could see an idle physical medium (i.e. neither will see the other's carrier signal), an both will conclude that no other node is currently using the network. In this case, both will then decide t transmit and a collision will occur. The collision will result in the corruption of the data being sent, which wi subsequently be discarded by the receiver since a corrupted Ethernet frame will not have a valid 32-bit MA CRC at the end.

3

(ii) Collision Detection [3 marks]

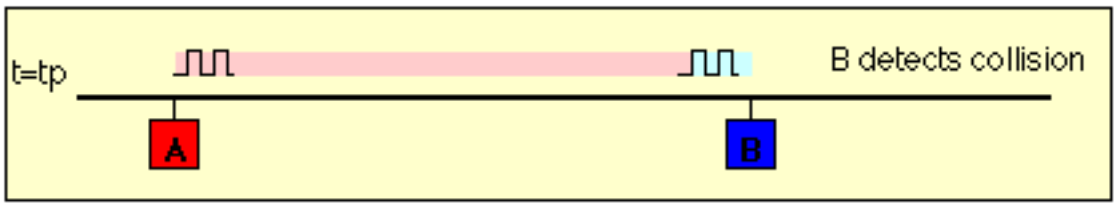

A second element t the Ethernet acces protocol is used t detect when a coll sion occurs. Eac transmitting nod monitors its ow transmission, and it observes a collision (i.e. excess current above what it is generating, i.e. > 24 mA) it stops transmission im mediately and instead transmits a 32-bit jam sequence. The purpose of this sequence is to ensure that any otl er node which may currently be receiving this frame will receive the jam signal in place of the correct 32-b 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 i Ethernet defines a minimum frame size (i.e. no frame may have less than 46 bytes of payload). The minimur frame size is related to the distance which the network spans, the type of media being used and the number c

Mark

repeaters which the signal may have to pass through to reach the furthest part of the LAN. Together these de fine 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 i the same way by transmitting the jam sequence. At time t=0, a frame is sent on the idle medium by compute A.

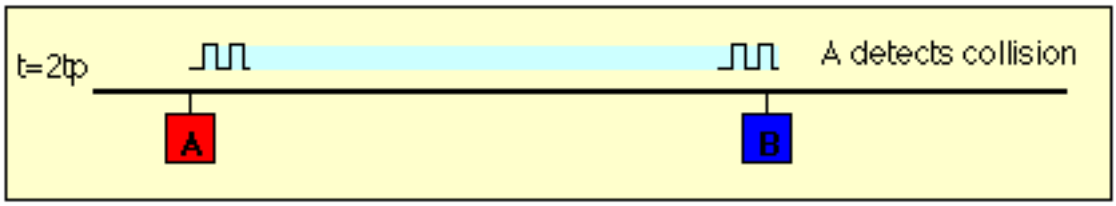

A short time later computer B als transmits. (In thi case, the mediun as observed by th computer at B hap pens to be idle too After a perioc equal to the propagation delay of the network, the computer B detects the other transmission from A, and i aware of a collision, but computer A has not yet observed that computer B was also transmitting. B continue to transmit, sending the Ethernet Jam sequence (32 bits).

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

3

(iii) Collision Domain [3 marks]

Traditional Ethernet uses a bus architecture in which all the computers connected to the cable share the capaci ty of the medium using CSMA/CD. In practice, most Ethernet networks employ hubs and repeaters, bu these do not change thebasic rules of sharing. A network of repeaters and hubs is therefore called a "Share Ethernet" or a "Collision Domain". The various systems sharing the Ethernet all compete for access using th CSMA/CD access protocol. This means that only one system is allowed to transmit within the Collision Do main at any one time. Each system has to share a proportion of the available network bandwidth.

In contrast, the use of bridges, switches and routers separates each cable segment into an independent colis sion domain.

3

(b) Describe the phenomenon of Ethernet Capture [5 marks]

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

Computer A dominates computer B. Originally both computers have data to transmit. A transmits first. A an B then both simultaneously try to transmit. B picks a larger retransmission interval than A (shown in red) an defers. A sends, then sends again. There is a short pause, and then both A and B attempt to resume transmis sion. A and B both back-off, however, since B was already in back-off (it failed to retransmit), it choose from a larger range of back-off times (using the exponential back-off algorithm). A is teherfore more likely t succeed, which it does in the example. The next pause in transmission, A and B both attempt to send, howev er, since this fails in this case, B further increases its back-off and is now unable to fairly compete with A.

A similar situation may arise when many sources compete with one source which has much more data t send. Under these situations some nodes may be "locked out" of using the medium for a period of time.
The use of full duplex cabling or higher speed transmission (e.g. 100 Mbps Ethernet) eliminates this problen

Mark

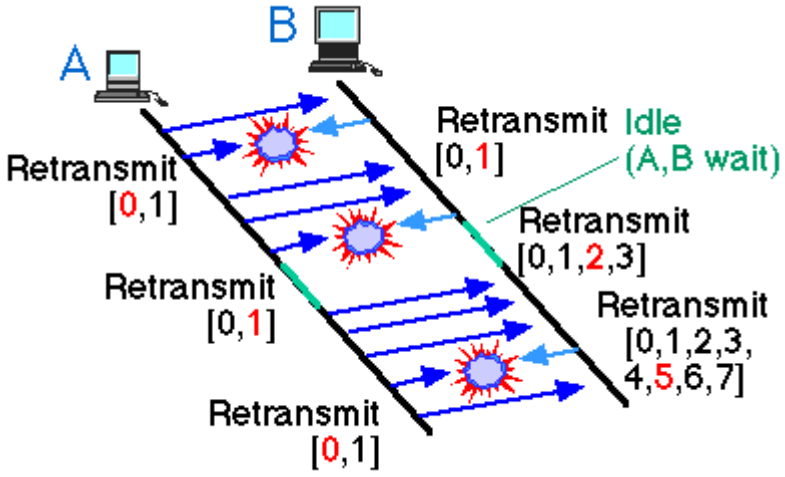

5

(c) A TCP session sends 10 packets per second over an Ethernet Local Area Networ (LAN). Each packet has a total size of 1480 B (excluding the preamble and cyclic redun dancy check (CRC)). Calculate the size of the headers, and hence the TCP payload data. What therefore is the TCP throughput of the session? [6 marks]

First determine theprotocol headers which contribute to the PDU size:

MAC Header (14 bytes) + IP Header (20 bytes) + TCP(20 bytes) + TCP Payload (?bytes)

Next determine the size of the payload:

Payload = 1480 - (14+20+20) = 1426 B

Throughput = number of useful (data) bits transferred by a layer using the services of the layer below.

= 1426 x 8 x 10 = 114 kbps.

6

Mark

2. A small Local Area Network (LAN) has four computers, A, B, C and D connected in th following topology:

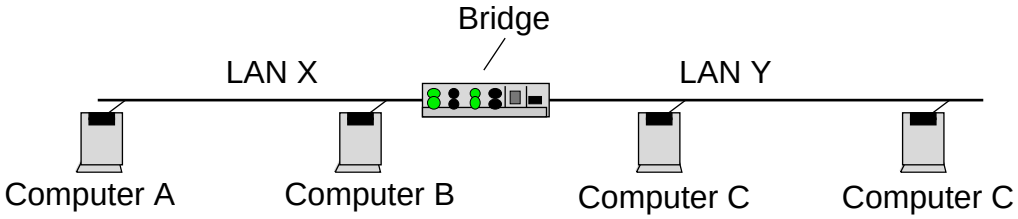

(a) The computer A sends a graphics file of size 10 MB simultaneously to computers B, C and D using Unicast packets constructed by the Universal Datagram Protocol (UDP). Cal culate the utilisation of LAN X, given that each frame carries 1024 B of UDP payloa data, and transmission is at 50 packets per second to each destination. [10marks]

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 Payloa (1024 bytes) + CRC-32 (4 bytes)

+4 The inter-frame gap may also be considered as overhead, which will yield a slightly higher answer.

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

+2

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

+2 Assume 10 Mbps Ethernet operation.

=8 Total utilisation = 8624 x 150 / (10 x 1000 000 x 100) =13%

(b) What is the utilisation on LAN Y? [2 marks]

This is different, since this is unicast transmisison, the bridge will not forward packets from A to B. It wi however forward packets from A to B and C. The utilisation on LAN Y is therefore:

2 2/3 of 13%, i.e. 9%.

(c) How does Multicast transmission differ from Unicast transmission? [6 marks]

Unicast is the term used to describe communication where a piece of information is sent from one point to an other point. In this case there is just one sender, and one receiver.
Unicast transmission, in which a packet is sent from a single source to a specified destination, is still the pre dominant form of transmission on LANs and within the Internet. All LANs (e.g. Ethernet) and IP network 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.

The hardware address is also known as the Medium Access Control (MAC) address, in reference to the IEE 802.x series of standards which define Ethernet. Each computer network interface card is allocated a globall unique 6 byte address when the factory manufactures the card (stored in a PROM). This is the normal sourc address used by an interface. A computer sends all packets which it creates with its own hardware source ac

Mark

dress, and receives all packets which match its hardware address or the broadcast address.

Multicast is the term used to describe communication where a piece of information is sent from one or mor points to a set of other points. In this case there is may be one or more senders, and the information is distrib uted to a set of receivers (theer may be no receivers, or any other number of receivers).

One example of an application which may use multicast is a video server sending out networked TV channels Simultaneous delivery of high quality video to each of a large number of delivery platforms will exhaust th capability of even a high bandwidth network with a powerful video clip server. This poses a major salabilit issue for applications which required sustained high bandwidth. One way to significantly ease scaling to larg er groups of clients is to employ multicast networking.

Multicasting is the networking technique of delivering the same packet simultaneously to a group of clients IP multicast provides dynamic many-to-many connectivity between a set of senders (at least 1) and a group o receivers. The format of IP multicast packets is identical to that of unicast packets and is distinguished only b the use of a special class of destination address (class D IP address) which denotes a specific multicast group Since TCP supports only the unicast mode, multicast applications must use the UDP transport protocol.
Unlike broadcast transmission (which is used on some local area networks), multicast clients receive a strean of packets only if they have previously elect to do so (by joining the specific multicast group address). Mem bership of a group is dynamic and controlled by the receivers (in turn informed by the local client applica tions). The routers in a multicast network learn which sub-networks have active clients for each multicas group and attempt to minimise the transmission of packets across parts of the network for which there are n active clients.

The multicast mode is useful if a group of clients require a common set of data at the same time, or when th clients are able to receive and store (cache) common data until needed. Where there is a common need for th same data required by a group of clients, multicast transmission may provide significant bandwidth saving (up to 1/N of the bandwidth compared to N separate unicast clients).

The Ethernet network uses two hardware addresses which identify the source and destination of each fram sent by the Ethernet. The destination address (all 1 s) may also identify a broadcast packet (to be sent to a connected computers) or a multicast packet (msb=1) (to be sent only to a selected group of computers).6

The appearance of a multicast address on the cable (in this case an IP multicast address, with group set to th bit pattern 0xxx xxxx xxxx xxxx xxxx xxxx) is therefore as shown below (bits transmitted from left to right)

```
 0                   23 IP Multicast Address Group  47
 |                   | <-------------------------->|
 1000 0000 0000 0000 0111 1010 xxxx xxx0 xxxx xxxx xxxx xxxx
 |                         |
 Multicast Bit                 0 = Internet Multicast
                             1 = Assigned for other uses
```

8

(d) Calculate the utilisation for LAN Y when the file is sent using multicast packets instea of the unicast packets used in section (a). [2 marks]

Bridges always forward multicast packets.

Only one multicast packet is sent to each destination.

The utilisation is therefore:

2

Total utilisation = 8624 x 50 / (10 x 1000 000 x 100) =4.3%

Mark

3. (a) Summarise the functions of the lowest three layers of the Open System Interconnection (OSI) reference model. Ensure your answer includes a sketch of the model with each of the three layers labelled. [6 marks]

The two lowest layers operate between adjacent systems connected via the physical link and are said to work "hop by hop". The protocol control information is removed after each "hop" across a link (i.e. by each System) and a suitable new header added each time the information is sent on a subsequent hop.

The network layer (layer 3) operates network-wide and is present in all systems and responsible for overall co-ordination of all systems along the communications path. The OSI layers may be summarised by:

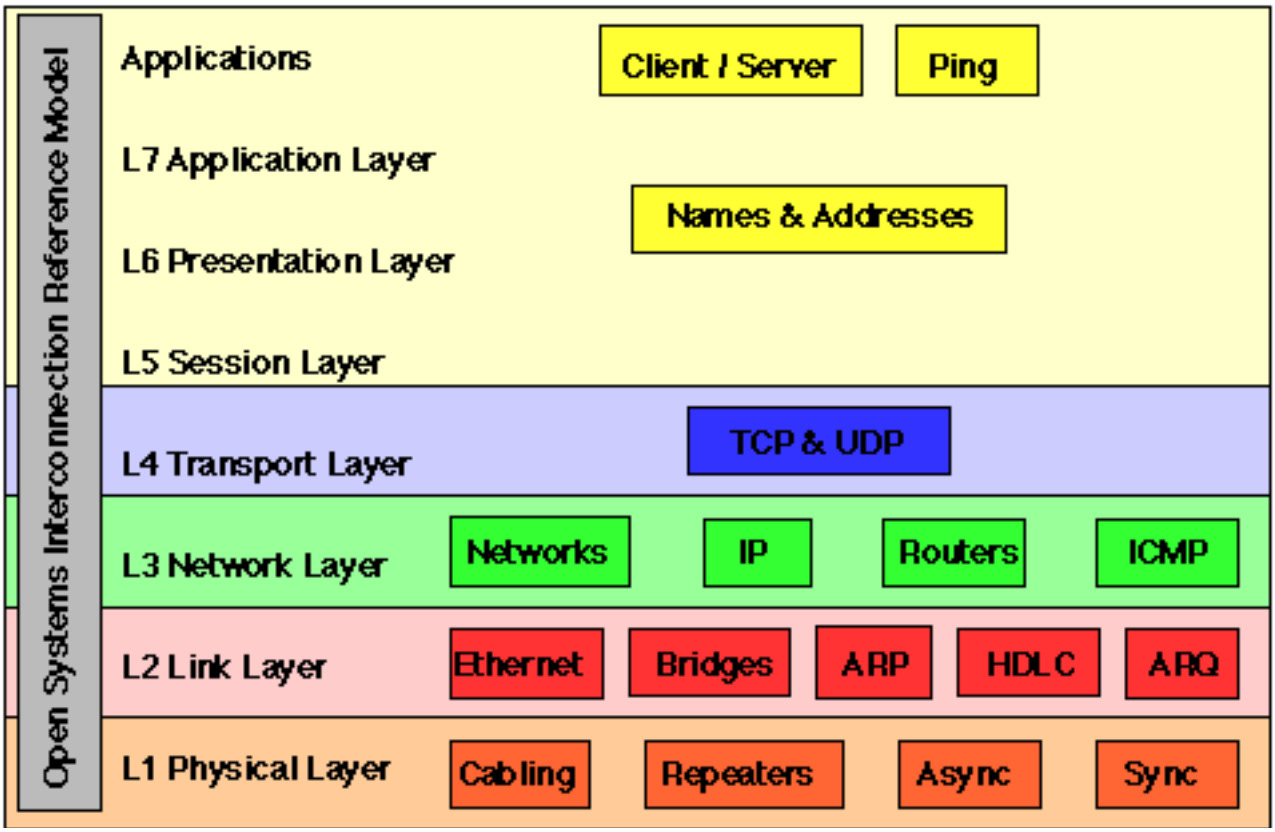

Physical layer: Provides electrical, functional, and procedural characteristics to activate, maintain, and deactivate physical links that transparently send the bit stream; only recognises individual bits, not characters or multicharacter frames.

Data link layer: Provides functional and procedural means to transfer data between network entities and (possibly) correct transmission errors; provides for activation, maintenance, and deactivation of data link connections, grouping of bits into characters and message frames, character and frame synchronisation, error control, media access control, and flow control (examples include HDLC and Ethernet)

Network layer: Provides independence from data transfer technology and relaying and routing considerations; masks peculiarities of data transfer medium from higher layers and provides switching and routing functions to establish, maintain, and terminate network layer connections and transfer data between users.7

The layers above layer 3 operate end-to-end and are only used in the End Systems (ES) which are communicating. The Layer 4 - 7 protocol control information is therefore unchanged by the IS in the network and is delivered to the corresponding ES in its original form. Layers 4-7 (if present) in Intermediate Systems (IS) play no part in the end-to-end communication.

6

Mark

(b) With reference to the communication between layers in the Open System Interconnection (OSI) reference model describe the terms:

(i) Hop-by-hop [2 marks]

Protocol layers may be defined in such a way that the communications within a layer is independent of the operation of the layer being being used. This is known as "peer-to-peer" communication and is an important goal of the OSI reference model. The communication takes place with the peer data link layer protocol in the next directly connected system (either an Intermediate System or an End System). Communications between an ES and an IS or between an IS and another IS is always hop-by-hop. Using the services of teh link layer which joind the two systems. This is also true of two directly connected End Systems, although usually the two end systems will not directly communicate over a wide area network.

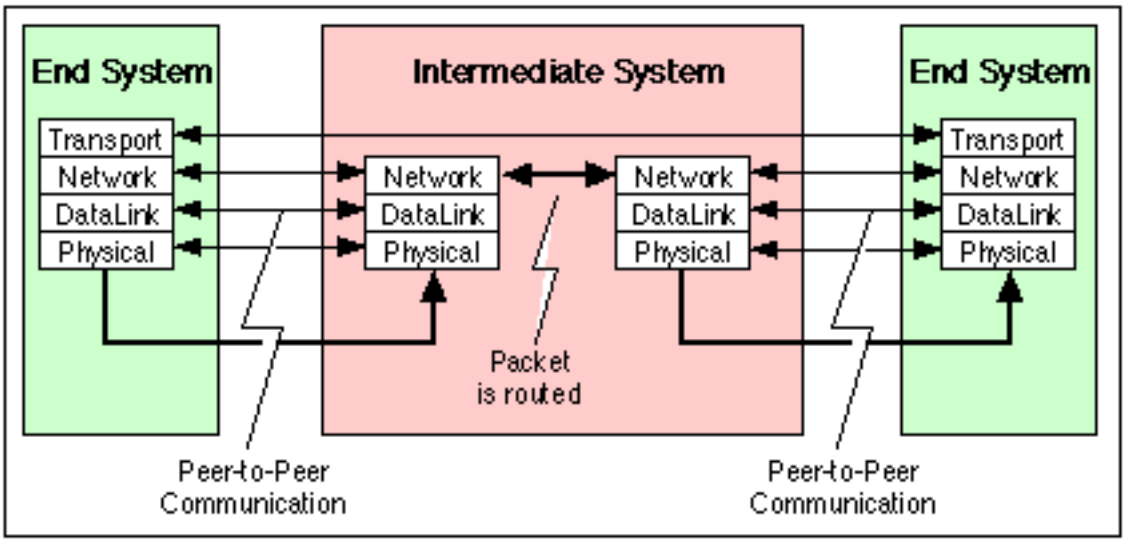

2

(ii) End-to-End [2 marks]

The figure above provides an example of the OSI reference model supporting peer-to-peer communication between two End Systems (ES). In this case, the transport protocol communicates end-to-end using the service of the network layer below. The peer-to-peer communication takes place between the end systems using the transport protocol (e.g. TCP) which using the services of the network layer (e.g. IP).

2

(c) What are the four requirements for reliable data transfer? [4 marks]

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

(i) No loss (at least one copy of each frame is sent)
(ii) No duplication (no more than one copy is sent)
(iii) FIFO delivery (the frames are forwarded in the original order)
(iv) A frame must be delivered within a reasonable period

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

4

(d) Which layer provides reliability in the TCP/IP protocol suite? [1 mark]

TCP normally provides a realiable transport protocol at level 4 of the OSI reference model. UDP also provides an alternate datagram service.

1

Mark

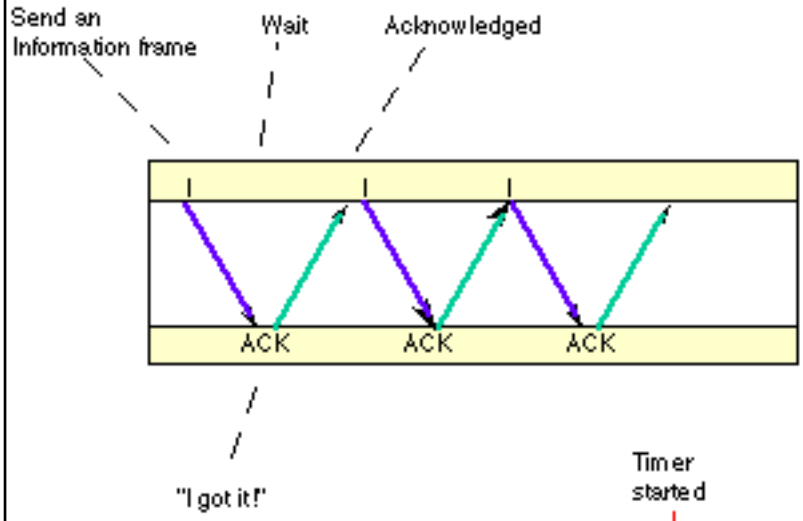

(e) Explain the operation of Automat ic Repeat Request (ARQ) protocols and illustrate your drawing by show ing how a Stop and Wait ARQ proto col may retransmit a single packe which was corrupted and discarde within a network. [5 marks]

Stop and Wait transmission is the simples reliability technique and is adequate for very simple communications protocol. A

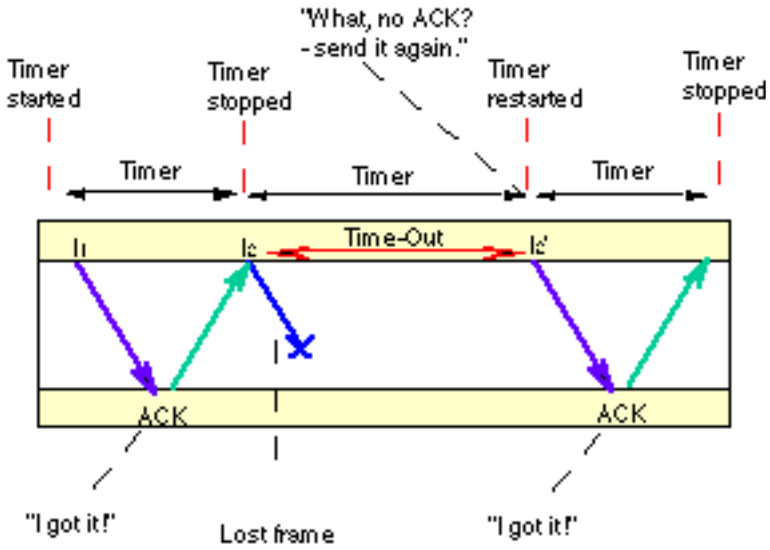

stop and wait protocol transmits a Protocol Data Unit (PDU) of information and then waits for a response. The receiver receives each PDU and sends an Acknowledgement (ACK) PDU if a data PDU is received correctly, and a Negative Acknowledgement (NACK) PDU if the data was not received. In practice, the receiver may not be able to reliably identify whether a PDU has been received, and the transmitter will usually also need to implement a timer to recover from the condition where the receiver does not respond.
Under normal transmission the sender will receive an ACK for the data and then commence transmission of the next data block. For a long delay link, the sender may have to wait an appreciable time for this response. While it is waiting the sender i said to be in the "idle" state and is unable to send further data.

The blue arrows show the sequence of data PDUs being sent across the link from the sender (top to the re ceiver (bottom). A Stop and Wait protocol relies on two way transmission (full duplex or half duplex) to a low the receiver at the remote node to return PDUs acknowledging the successful transmission. The acknowl edgements are shown in green in the diagram, and flow back to the original sender. A small processing dela may be introduced between reception of the last byte of a Data PDU and generation of the correspondin ACK.

When PDUs are lost, the receiver will not normally be able to identify the loss (most receivers will not receiv anything, not even an indication that something has been corrupted). The transmitter must then rely upon timer to detect the lack of a response.

In the diagram, the second PDU of Data is corrupted during transmission. The receiver discards the corrupte data (by noting that it is followed by an invalid data checksum). The sender is unaware of this loss, but start a timer after sending each PDU. Normally an ACK PDU is received before this the timer expires. In this cas no ACK is received, and the timer counts down to zero and triggers retransmission of the same PDU by th sender. The sender always starts a timer following transmission, but in the second transmission receives a ACK PDU before the timer expires, finally indicating that the data has now been received by the remote node

5

Mark

4. (a) The High Level Data Link Control (HDLC) protocol uses a technique known as 0-Bit Insertion to provide transparency. Calculate the number of bits which will be transmitted when an HDLC link serialises the following bytes:

0xFE 0xF1 0xF0 0xFF [6 marks]

First note that this is heaxdecimal value for bytes.

Byte values in binary are:

1111 1110 * 1111 0001 * 1111 0000 * 1111 1111

Since transmission always lsb first, the data needs to be re-written in binary transmission order:

0111 1111 * 1000 1111 * 0000 1111 * 1111 1111

The bit-stream now under-goes 0-bit insertion to provide transparency to HDLC:

0111 11 (0) 11 * 1000 1111 * 0000 1111 * 1 (0) 111 11 (0) 11

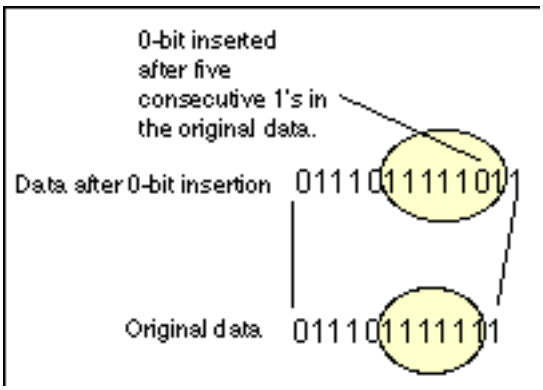

There are therefore 3 0-bits inserted.

The total length is therefore (4 x 8) + 3 bits

35 bits.

6

(b) Describe the differences between a Local Area Network (LAN) and a Metropolitan Area Network (MAN). [4 marks]

The Local Area Network (LAN) is by far the most common type of data network. As the name suggests, a LAN serves a local area (typically the area of a floor of a building, but in some cases spanning a distance of several kilometers). Typical installations are in industrial plants, office buildings, college or university campuses, or similar locations. In these locations, it is feasible for the owning Organisation to install high quality high-speed communication links interconnecting nodes. Typical data transmission speeds are one to 100 megabits per second.

A wide variety of LANs have been built and installed, but a few types have more recently become dominant. The most widely used LAN system is the Ethernet system developed by the Xerox Corporation.

In summary, a LAN is a communications network which is:
- local (i.e. one building or group of buildings)
- controlled by one administrative authority
- assumes other users of the LAN are trusted

A Metropolitan Area Network (MAN) is one of a number of types of networks (see also LAN and WAN). A MAN is a relatively new class of network. There are three important features which discriminate MANs from LANs :

Mark

The network covers a larger distance than a LAN. A MAN typically covers an area of between 5 and 50 km diameter. Many MANs cover an area the size of a city, although in some cases MANs may be as small as a group of buildings or as large as the North of Scotland.

A MAN (like a WAN) is not generally owned by a single organisation. The MAN, its communications links and equipment are generally owned by either a consortium of users or by a single network provider who sells the service to the users. This level of service provided to each user must therefore be negotiated with the MAN operator, and some performance guarantees are normally specified.

A MAN often acts as a high speed network to allow sharing of regional resources (similar to a large LAN). It is also frequently used to provide a shared connection to other networks using a link to a WAN.

4

(c) Why is HDLC preferable to sharde Ethernet in a MAN environment? [2 marks]

It is full duplex (but so also is Ethernet). The principal advanatge is that there is no length constraint, as thereis in shared Ethernet. HDLC links may be of any arbitary length. HDLC also does not presuppose a particular encoding scheme such as manchester encoding, and may be used over any type of synchronous physical link.

2

(d) With the help of diagrams explain how a Router may connect the two types of network. [8 marks]

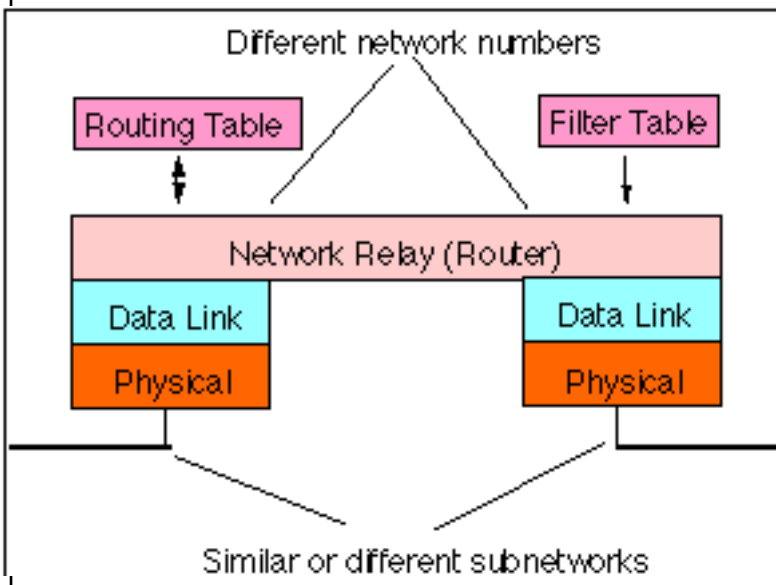

A router is an Intermediate System (IS) whic operates at the network layer of the OSI refer ence model. Routers may be used to connect tw or more IP networks, or an IP network to an in ternet.

A router is most suited for the connection of LAN to a MAN. The router allows two separ ately administered networks to communicat without forming one homogenous network. Th two networks may have different media, and be long to different IP networks (in the case of IP) The router also provides routing of packets t destinations reachable via the MAN and can cor trol access to/from the MAN.

A router consists of a computer with at least tw network interface cards supporting the IP proto col. The router receives packets from each interface and forwards the received packets to an appropriate out put interface. The router uses the IP address, along with routing information held within the router and store in a routing table, to determine the destination for each packet. A filter table may also be used to ensure tha unwanted packets are discarded. The filter may be used to deny access to particular protocols or to prever unauthorised access from remote computers.

Routers are often used to connect together networks which use different types of links (for instance an HDL(link connecting a WAN to a local Ethernet LAN). The optimum (and maximum) packet lengths (i.e. the Maxi mum Transfer Unit (MTU)) is different for different types of network. A router may therefore uses IP to pro vide segmentation of packets into a suitable size for transmission on a network.

Routers :
- Are more expensive than Bridges or Switches
- Work at Network Layer (e.g. IP) and support one or more protocols
- Connect separate networks into an internet
- May protect networks from unauthorised access

8

Mark

5. (a) Computers in a network are identified by either a name or a address. Explain th following terms relating to addresses:

(i) An address cache [2 marks]

The cache is an area of temp. storage that keeps recently resolved addresses. The cache is consulted prior t performing an address resolution across the network. If the required entry is i the cache this value is usec thus saving a network exchange. The contents of the cache may become stale after a period of time, and tehei fore cache entries are normally aged ti.e. deleted aftre a fixed time period) to ensure that old (possibly wrong

2 addresses are not used.

(i) A network address [2 marks]

An address is a unique identifier used by the computer protocols to identify an entity in a network. A typica address could be 8036565901, or a binary expansion of such a number, for example, 1000 0000 0011 011 0101 0110 0101 1001 0000 0001.

The fields may help to determine where an entity is located, but this is not necessarily so; for example, th MAC hardware address used in Ethernet has the form <manufacturer><serial number>. This says nothin

2 about the location of the host computer on network.

(ii) A network name [2 marks]

The major distinction between names and addresses is whether they are intended to be human-readable or ma chine-readable. Names range from simple names of only local applicability, such as mail used to access ma service after gaining access to a computer providing this facility, to universal names. An example of a trul universal name is

<galaxy><star><planet><country><network><host><port>

The DNS provides names for computers which have internet addresses.

2 **(c) Explain the operation of the Domain Name Service (DNS)) [7 marks]**

Once there were only a few computers connected to the first internet, called the ARPANET, at that time every one knew each others IP address, so communication was easy, one simply typed the appropriate sequence c digits representing the IP number for each destination. After a while, the number of computers started t grow, and people began to forget the strange numeric IP numbers. So IP names came into being, and eac computer held a table of names and their associated addresses, which had to be updated as new computer were connected to the network.

Soon new computers were being connected to the network too quickly for everyone to keep up. Someone ha the bright idea of keeping just one central list, and such a list was created and stored at Stanford University Too add a new computer, one simply told the people at Stanford, and they added your name and IP address t the list. Every week, or so, you had to transfer the list to your own computer (using ftp).

A little while passed, and the network grew. Eventually, there were just too many computers, the people a Stanford became overloaded with requests to add and change the network information: the file of all address ees was getting too big, and there was a constant demand for the users to download new copies of the file The solution was to create a distributed database - called the Domain Name System (DNS).

In the DNS, there are a set of root domain servers (rather like the old Stanford computer), but they don't actu ally store much information. Instead they contain the IP addresses of other servers which have informatio about specific groups of addresses known as "domains". The root server is said to delegate responsibility fc each domain to a lower domain server. In turn, each of these servers may delegate other domains to othe

Mark

servers. Before long, there were many many domain servers each responsible for the groups of users in a lo cal area. Each server maintained pointers allowing them to find out information about other domains by send ing query messages to the other domain servers. In this way, any DNS server can resolve the name of an computer to an IP address of any user irrespective of whether that user is in the same local domain or is regis tered with some remote domain.13The term address resolution refers to the process of finding an address of computer in a network. The address is "resolved" using a protocol in which a piece of information is sent by client process executing on the local computer to a server process executing on a remote computer. The infor mation received by the server allows the server to uniquely identify the network system for which the addres was required and therefore to provide the required address. The address resolution procedure is complete when the client receives a response from the server containing the required address.

Example of the use of the DNS

This example considers a login from a computer X to a remote computer Y using a DNS server Z. The pro cess is shown in the figure below:

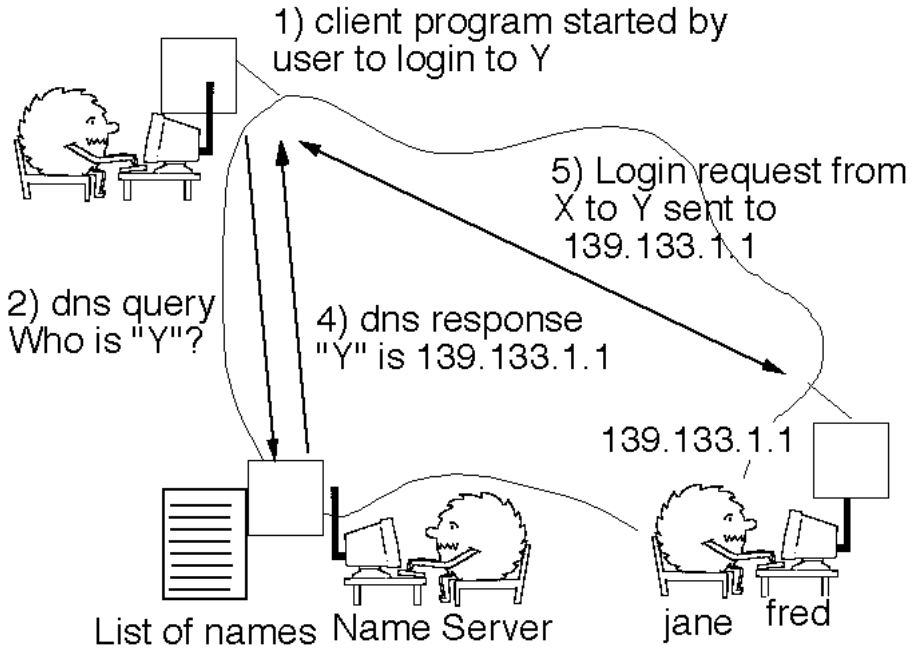

The process may be described in six steps:

A client program starts on the local computer (X) and attempts to resolve the network layer address of the re mote computer from a known name using a known dns server (Z).

A dns query is sent to the server in an IP packet from X to Z.

The server (Z) processes the query and consults local dns entries and (possibly) the entries of other remot dns servers.

The dns server returns a response with the requested information (assuming success) in an IP packet from Z to X.

The local computer (X) then makes a direct connection to the remote computer (Y).

The remote computer starts a process (server) to handle the requested login. All further packets between X and Y are directed to the respective client and server processes.

7

Mark

(d) (i) Sketch the protocol encapsulation used to construct this message [3 marks]

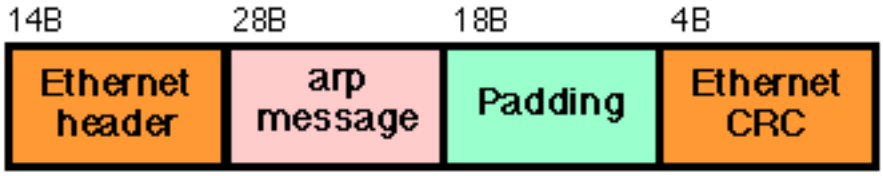

An arp message encapsulated i an Ethernet MAC frame (note the need for Ethernet Pad ding to ensure the minimur Ethernet PDU size)

3

(ii) By observing the Medium Access Control (MAC) header, determine if this is an ar request (query) or an arp response [2 marks]

An Ethernet network uses two hardware addresses which identify the source and destination of each fram sent by the Ethernet. The destination address (all 1's) may also identify a broadcast packet (to be sent to a connected computers) or a multicast packet (msb=1) (to be sent only to a selected group of computers).

An arp request is sent with thebroadcast destination address (since it is not known which computer will reply a specific destination addres can be used).

2

An arp reply is directed to the source of the request (i.e. has a destination address corresponding to the sourc address in the request).

(iii) What is the target IP address which is to be resolved? [2 marks]

2

The required address is the last 4B of the arp message:

0x8b85 cc50 or, more commonly expressed as 139.133.204.80

The arp exchange is:

dent -> (broadcast) ARP C Who is 139.133.204.80, gordon ?
gordon -> dent ARP R 139.133.204.80, gordon is 8:0:20:96:10:1a

Complete arp decode follows:

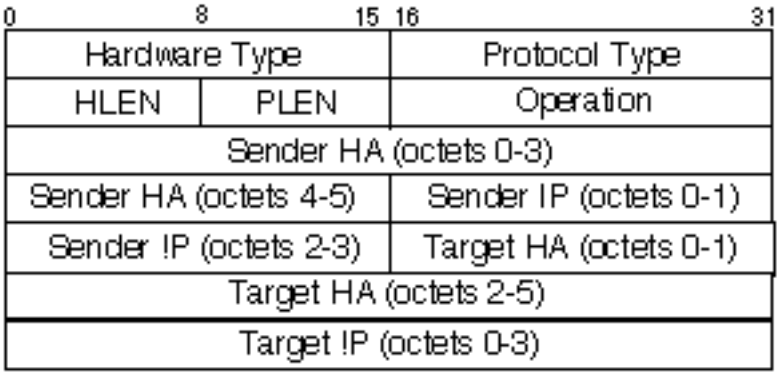

ARP: ----- ARP/RARP Frame -----
ARP:
ARP: Hardware type = 1
0001
ARP: Protocol type = 0800 (IP)
0800
ARP: Length of hardware address = 6 bytes
06
ARP: Length of protocol address = 4 bytes
04
ARP: Opcode 1 (ARP Request)
0001
ARP: Sender's hardware address = 8:0:20:b:b0:83
0800 200b b083
ARP: Sender's protocol address = 139.133.204.17, dent
8b85 cc11
ARP: Target hardware address = ?
ffff ffff ffff
ARP: Target protocol address = 139.133.204.80, gordon
8b85 cc50