THE BCS PROFESSIONAL EXAMINATION
Diploma

April 2002

EXAMINERS' REPORT

Computer Networks

The majority of candidates were poorly prepared for this examination and produced answers that ranged from weak to totally irrelevant. Other candidates could have scored higher marks if they had read the question properly. However, most questions were answered well by at least one candidate, indicating a lack of breadth in the preparation undertaken by the majority.

There was an even split between candidates answering too few and too many questions, perhaps out of desperation. Lots of candidates were unable to attempt all parts of their chosen questions.

Question 1

a) **Explain in detail, using diagrams as appropriate, why the minimum packet size for IEEE 802.3 is set at 64 bytes. (8 marks)**

b) **"Ethernet is not suited to being used in a real time environment". Do you agree with this statement? Give your reasons. (5 marks)**

c) **What advantages are there in IEEE 802.3 adopting Manchester encoding rather than straightforward binary encoding? (5 marks)**

d) **As IEEE 802.3 based network speeds increase, either the minimum frame size must be increased, or the maximum length of cable must decrease. Why is this the case? If we want to keep a minimum frame size of 64 bytes, what would be the maximum permissible length of cable if the transmission speed is 100Mbps? (7 marks)**

Part a) answers demonstrated a wide variation in knowledge and understanding. Some answers were totally correct, indicating an in-depth knowledge of the subject, others were plagued by guesswork and a lack of arithmetic ability.
Part b) produced answers ranging from a single word (and of a yes or no choice, a number of candidates chose the wrong one) to a well constructed argument
Part c) where attempted, produced a range of correct answers but with varying detail.
Part d) was poorly answered, with the vast majority of candidates being unable to perform the simple arithmetic required.

Answer Pointers

a) Minimum frame size is required so that a sending station can detect a collision before sending the last bit of a frame:
The maximum length of an Ethernet network is 2.5km with 4 repeaters.
In this case, the end-end propagation delay is 25.6 □s (τ).
Therefore, in a worst case, a collision would be detected in 2τ = 51.2□s.
At 10Mbps, 51.2 □s is equivalent to 512 bit lengths = 64 bytes.
Thus 64 bytes is the minimum size of frame that lets us detect a collision in the time taken to transmit that frame.

b) Ethernet's behaviour (wrt the binary exponential backoff used in the event of a failed attempt to transmit) is non-deterministic. Therefore it is possible

that particular stations may have to wait a long time before they get a chance to transmit their data.
In the worst case (a heavily loaded, busy segment) stations may in theory never be able to transmit. This slight possibility makes Ethernet unsuitable for deployment in “hard” real-time environments where guarantees wrt maximum time before transmission are required.

c) Using a straight binary encoding scheme of 0v = 0 5v = 1 leads to ambiguities (eg if a station sends the bit string 0001000 others might interpret it as 10000000 or 01000000 because they cannot tell the difference between an idle sender and a 0 bit).

Bit synchronisation is difficult - there is not a separate clock line and so some clock drift between machines is inevitable. Because every bit period contains a transition when Manchester encoding is used (high-to-low or low-to-high representing 0 and 1), it is easier for the receiver to synchronise with the sender.

d) Limiting the packet size, transmission speed or length of cable is necessary because a transmitting station must be able to detect if a collision has occurred within the time taken to transmit that particular frame. We must detect a collision within the time taken to transmit 64 bytes, or 512 bits. At 100Mbps, 512 bits takes 5.12□secs to transmit. Assuming a propagation speed of 200m/□sec, a bit can travel 1024m in this time.
In the worst case, this 1024m corresponds to a round trip, i.e. a 512m separation between stations. Therefore the maximum permitted cable length would be 512m.

Question 2

a) **Explain the role of ARP and, using examples as appropriate, give a summary of its operation. (10 marks)**

b) **Explain, using the IP address “192.33.45.104”, how a subnet mask is used to extract the network number and the host number. (5 marks)**

c) **Outline the essential features of the algorithm used by IP based networks to route packets to their destination. (10 marks)**

Part a) answers demonstrated that most people had heard of ARP, even if they could not remember the workings of the protocol.
Part b) was answered correctly in most cases, although a greater level of detail would have been preferred. Some candidates discussed variable length subnet masking which demonstrated a better knowledge of the subject than would normally have been expected.
Part c) was very poorly answered with the majority of candidates attempting this section demonstrating nothing more that total ignorance of the subject. Again, there were a couple of good answers that demonstrated sound knowledge of IP.

Answer Pointers

a) ARP (Address Resolution protocol) is used to provide a translation between internet (IP) addresses and Ethernet addresses. When a message is sent on a network running TCP/IP, the destination address specified is an IP address which only has relevance in software. For the message to be delivered over a

CSMA/CD network, it is necessary to translate the IP address into a hardware address that is recognisable by the CSMA/CD hardware inside the host machines This is achieved by maintaining tables that map IP addresses to hardware address If an IP address is used that is not found in a stations ARP table, it sends an ARP broadcast message around the network. When a station receives an ARP broadcast that contains its IP address, it is required to respond with its 48-bit Ethernet address. This address is then stored in the ARP table so that future packets can be sent directly.

b) If the bits within a netmask are set to 1 then the corresponding part of an IP address is a network number, if the bits are set to 0 then the corresponding part of the IP address is a host number. Therefore, by performing bitwise operations, the address mask can be used to retrieve either the network number or the host number from a 32-bit IP address

c) The IP routing algorithm is as follows:

```
RouteDatagram( Datagram, RoutingTable)

Extract destination IP address, D, from the datagram and compute the
network prefix, N
    IF N matches any directly connected network address THEN
        deliver datagram to destination D over that network
    ELSE
        IF the table contains a host-specific route for D THEN
                send datagram to next-hop specified in table
    ELSE
        IF the table contains a route for network N THEN
                send datagram to next hop specified in table
    ELSE
        IF the table contains a default route THEN
                send datagram to the default router specified in table
    ELSE declare a routing error
```

Algorithm first checks to see if the destination is on the current network. If so, then can use direct delivery to host.
If not, then check routing table for a specific route to the destination machine, sending the datagram along that route if one is found
If no host-specific route exists, then check routing table for a route to the *network* to which the destination is attached and send if appropriate
If not found, then use a default route, if one exists.

Question 3

a) **Describe the sub-layers found in the OSI Network layer and explain the need for them. (6 marks)**

b) **What is the role of the LLC and MAC layers in the IEEE architecture? How do these layers compare with those in the ISO OSI 7 Layer Reference Model? (6 marks)**

c) **Suggest a way in which a connectionless transport layer service could be provided over a virtual circuit network layer.**

Describe how your approach affects the volume of traffic over the network. Suggest a way in which the volume of control data may be reduced. (6 marks)

d) ***Routers, bridges*** **and** ***repeaters*** **are used to connect differing networks. Under what circumstances would each of these technologies be used? (5 marks)**

e) **Explain the difference between flow control and congestion control. (2 marks)**

Part a) – nobody actually read the question and rather than dealing with just the network layer, candidates launched into an irrelevant discussion of the entire OSI stack.
Part b) was better answered with most candidates recognising the role of LLC and MAC protocols, even if they could not map those protocols onto the correct layers of the OSI model.
Part c) was very poorly answered with hardly any candidates being able to describe a method to perform the task described.
Part d) was on the whole well answered, although again there were some candidates who demonstrated total confusion.
Part e) was answered appropriately by about 40%. However, answers such as "Flow control relates to controlling flow" are not useful and indicate nothing more that the mental flailing of the desperate.

Answer Pointers

a) The sublayers are:

SUBNET ACCESS SUBLAYER
Includes all of the functions necessary to access a particular subnet and is specific to that type of subnet. This layer merely access the subnet and does not attempt to modify the service

SUBNET ENHANCEMENT SUBLAYER
Performs a mapping of the facilities provided by a subnet to a common level of functionality. This may involve enhancement of the service or a de-enhancement of the service (reducing its functionality)

INTERNET SUBLAYER
Provides the global network service to the transport layer.
Includes functions such as routing and switching between subnets.

b) EITHER:

If a connection request can contain user data, then the datagram message can be sent within a connection request. The destination transport entity immediately returns a disconnect request. The sender then has to reply with a disconnect confirm. This results in three messages for each datagram, .i.e. for short messages, efficiency is as low as 33%

OR:

If the connection request cannot contain user data, then we need the following message sequence:

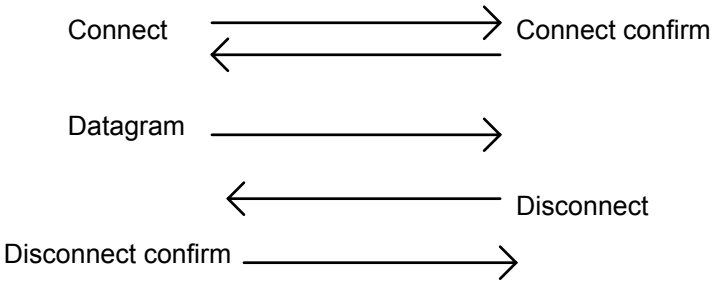

This results in at least five messages for each user message, i.e. efficiency is 25%
Efficiency may be improved if the destination holds the connection open for a time, assuming that more datagrams will be generated by the sender. This avoids having to wait for a connection to be opened for each datagram, and reduces the volume of control traffic across the network. The connection can then be closed after a period of inactivity.

c) Routers connect dissimilar networks at the network layer, and are used to link networks that operate different network protocols. Bridges are used to link networks that use either different physical media, or different media access control (MAC) protocols. Repeaters link networks that use the same MAC protocols over identical, or very similar, physical media.

d) Flow control relates to the control of the flow of information between a single sender and receiver pair, i.e. ensuring that the sender does not transmit more messages than the receiver can handle. Congestion control is a global mechanism designed to ensure that no more messages enter a *network* than can be handled.

Question 4

a) **Explain how the quality of a transmission is affected by physical considerations. (4 marks)**

b) **Describe the differences between asynchronous and synchronous transmission. (4 marks)**

c) **Describe THREE different techniques by which frame boundaries may be encoded within a transmitted bit stream.**
Explain *character stuffing* and state which technique it is associated with and why it is needed. (9 marks)

d) **Under what circumstances would you expect pipelining to improve data link layer protocol throughput? (4 marks)**

e) **Briefly describe how:**
i) **Negative Acknowledgement**
ii) **Piggy-back Acknowledgement**
may be used to enhance the performance of data link protocols. (4 marks)

Part a) answers were variable. Most candidates were able to accurately describe attenuation, resistance and electromagnetic interference. Lesser numbers considered that bit rate would have an impact and some answers, such as “if the cable is up a muddy mountain” or “good news travels faster than bad” whilst good for comic relief, do not demonstrate even a passing knowledge of the subject.

Part b) was answered correctly by most of the candidates that attempted this section but, as usual, a proportion confused the two transmission types.
Part c) was well answered, up to a point, with most people describing the BSC encoding method or a derivative. Bit encoding was also a popular choice. Character stuffing was well defined by a number of candidates.
Part d) indicated that most people had not heard of pipelining, although again some candidates had and were able to produce a well-reasoned answer.
Part e) was poorly answered, on the whole with most candidates demonstrating a poor understanding of both acknowledgement methods.

Answer Pointers

a) Transmission quality may be affected by :

- type of transmission medium used
  - some media are less prone to corrupting data than others.
  - choice of medium is dependent upon desired bit rate and distance involved.
- the bit rate
  - the higher the bit rate, the less time each bit is on the line. Therefore, the more likely we are to miss a bit and corrupt the message.
- distance involved
  - the longer the distance, the more the signal will be degraded by electrical resistance or attenuation.

b) Asynchronous Transmission

- Individual bytes are transmitted with variable time intervals between them
- A character is framed by a start and stop element.
- When the receiver detects a start element, it begins to sample the incoming signal, to allow it to assemble the individual bits
- The local clock that controls the sampling rate is resynchronsied for each character, to minimise the chance of error due to “clock drift”

Synchronous Transmission

- A constant rate clock determines the exact time interval at which bits are transmitted
- The receive clock is synchronised by means of either a separate clock line, or by transitions in the line signal.
- Special SYNC characters are transmitted at the start of the block to allow the receiver to determine character boundaries
- More efficient than asynchronous transmission as start and stop bits are not required
- If synchronisation is lost, the entire block must be retransmitted

c) <u>Character Count</u>

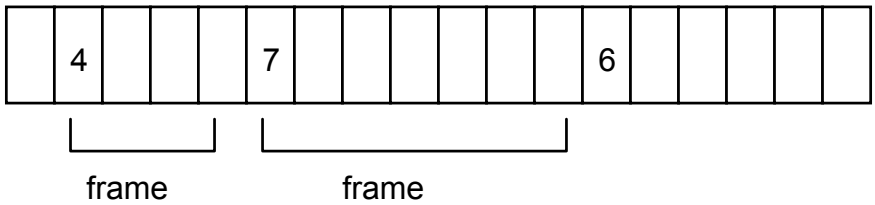

- Character count holds total size of message (including itself)
- Problems if character count is corrupted
- Even if an error is detected, it is impossible to re-synchronise to the start of the next frame

Character oriented Protocols

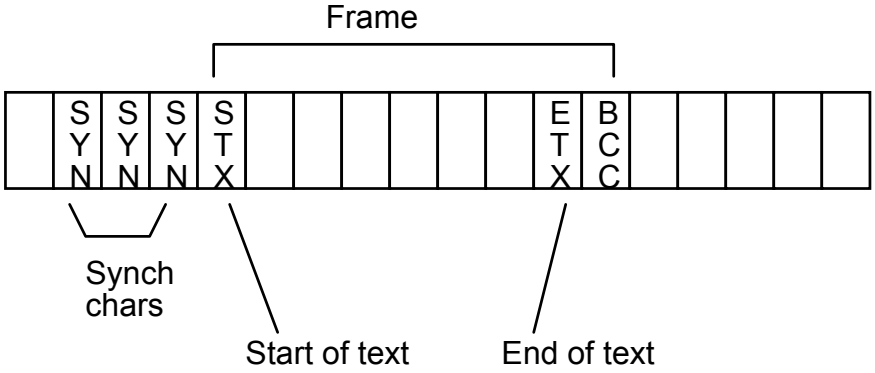

- Use delimiters to mark the start and end of frame
- Example is BSC protocol

Bit oriented protocols

- Like the character oriented approach, but a special bit pattern is used as the STX/ETX delimiter

Character Stuffing

Problems occur when framing binary data using a character-oriented protocol (for example 0x02 appears the same as an STX).
So, when we want to show a control character, we insert a DLE (data link escape) character before it. Therefore, an <STX> appears as <DLE> <STX>. (0x10 0x02 in BSC protocol)

(e) Pipelining would improve the throughput of the data link layer when there is normally more than one message waiting to be sent to a particular destination.
It is particularly useful when the round-trip delay is long compared to the transmission time of a message.

(f) Negative Acknowledgement

- Instead of sending an acknowledgement when a frame has been correctly received, send a Negative Acknowledgement if a frame is either corrupted, or if it hasn't arrived after a suitable timeout period has expired.

Piggyback Acknowledgement

- Instead of sending an acknowledgement message when a frame has been correctly received, send an acknowledgement embedded within a normal data packet. This reduces the bandwidth used by acknowledgements and allows an increase in data throughput

Question 5

5. Two minicomputers, running multi-tasking operating systems, are connected via an Ethernet local area network, with a 10Mbps capacity. A task in one computer performs a request-reply message transaction on a server task in the other computer.

a) **Discuss the factors which affect the transaction response time. (10 marks)**

b) **Identify and explain which overheads affect data throughput between remote tasks. (5 marks)**

This was the chance for candidates to demonstrate an understanding of a wide range of network topics and relate them to real life operations. Most candidates were able to identify at least some of the issues but very few were able to demonstrate the broad knowledge required.

Part a) was the better answered, with most candidates being able to identify at least some of the moe relevant issues.
Part b) was poorly answered, with most candidates either not reading or not understanding the question and choosing merely to repeat their answer to Part a)

Answer Pointers

a) Transaction response time = message delay in transferring request to server task + processing time + delay in transferring the reply
Factors that contribute to message transfer delays include:
Queuing time in source processor
Local operating system time to get message into communications system Processing time within various layers of the communications software in both source and destination processor
Access delay to ring whilst waiting for token to be received
Transmission time = length of message / 16Mbps
Possible retransmission times due to errors (including the time to send a NAK from the receiver or the expiry of a timeout at the source)
Operating system overheads at the destination

b) Application data is limited by the following:
Protocol headers and trailer containing sequence numbers, addresses, CRC checks, etc. (often > 50 bytes per message)
Protocol supervisory frames such as connection initialisation and termination, or acknowledgements
Errors resulting in retransmission of messages
Re-initialising the destination's network interface after receiving one message involves providing a new receive buffer and resetting the hardware. This results in a "dead time" within which no message can be received
Processing overheads in the OS and communications system of both sender and receiver reduce the message transfer rate

Question 6

a) **How many bits can occupy a 1000-metre Token Ring, operating at 4Mbps, containing:**
i) **2 equally-spaced stations?**
ii) **100 equally-spaced stations?**
You should assume a propagation speed of 200m/µsec. (10 marks)

b) **Explain the role of each field within the AC byte of an IEEE 802.5 control token. (8 marks)**

c) **Explain how transmitting and receiving stations on a Token Ring network make use of the FS byte within an information frame. (7 marks)**

This was not a popular question, perhaps due to the declining use of Token Ring technology in the marketplace.

Part a) demonstrated a range of answers, ranging from the totally correct, to those that were plagued by arithmetic errors. Hardly anyone appreciated the delays introduced by NICs in the attached computers.
Part b) was poorly answered, indicating a general lack of knowledge about the way that Token Ring works
Part c) was correctly answered by a number of candidates although again, answers were very general with a disappointing grasp of detail.

Answer Pointers

a)

(i) 1000m ring with 2 stations -> 500m separation of stations
4 Mbps transmission speed -> 4 bits every □sec.
With a propagation speed of 200m/□sec, each bit will travel the ring in 5□sec. Without delay, the maximum number of bits on the ring will be 20.
However, as each station has a 1 bit delay, each bit will take 5 + (2 x 0.25) = 5.5□sec to travel around the ring and a maximum of 22 bits can occupy the ring at once.

(ii) 2000m ring with 100 stations -> 20m separation of stations
4 Mbps transmission speed -> 4 bits every □sec.
With a propagation speed of 200m/□sec, each bit will travel the ring in 10□sec. Without delay, the maximum number of bits on the ring will be 40
However, as each station has a 1 bit delay, each bit will take 10 + (100 x 0.25) = 35□sec to travel around the ring and a maximum of 140 bits can occupy the ring at once.

b) AC byte contains 4 fields to store information about priority levels and the status of the frame:

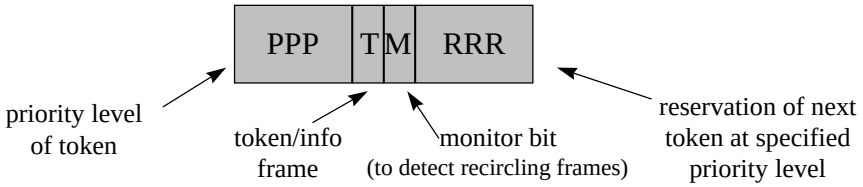

c)

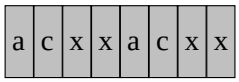

a: address recognised bit
c: frame copied bit
x: undefined bit

A and C bits are duplicated as a check against possible corruption during transmission. A bit is set when a receiving station detects its own MAC address is the same as the destination address in the information frame. C bit is set iff the receiving station successfully copied the entire frame into its frame buffer. Thus, when the frame arrives back at its source, if the A bit is set, the destination station is known to exist on the network, if the C bit is set, the sender knows that the frame was received correctly.