THE BCS PROFESSIONAL EXAMINATIONS
BCS Level 5 Diploma in IT

April 2008

EXAMINERS' REPORT

Computer Networks

General Comments

The responses to questions were of varying quality, some were excellent answers to questions and some quite poor answers. There was some evidence that students had not read the questions properly and thus had given a general answer to the topic rather than addressing the precise points in the questions have raised. It is a practice if the students write a brief set of relevant points before expanding them to answers.

Question 1

a) Produce a sketch of the ISO Reference Model with each of the seven layers clearly labelled. Clearly show which layers are mainly associated with the operation of the network infrastructure and which are mainly associated with the application.

(9 marks)

b) What is meant by the term '*peer to peer protocol* '?

(6 marks)

c) A computer is connected to an Ethernet LAN. Assuming that it is running a web browser to access the Internet using the TCP/IP protocol produce a protocol layer diagram to represent this communication. Clearly show how the protocols being used map onto the ISO Reference Model.

(10 marks)

Answer Pointers

a)

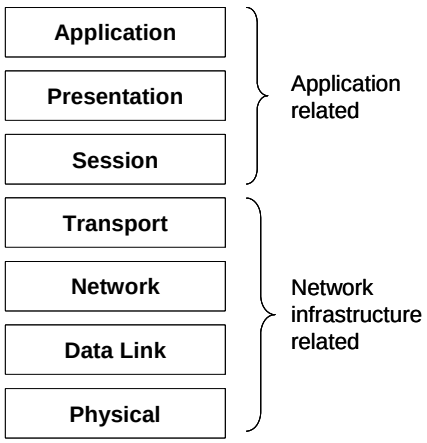

b)

In a layered protocol architecture data flows vertically through the layers in each end station and horizontally over the network medium between end-stations. However, a peer to peer protocol is defined as the communication that takes place between two layers at the same level in different end-stations. Conceptually a peer to peer protocol operates horizontally and directly between the two layers as shown below.

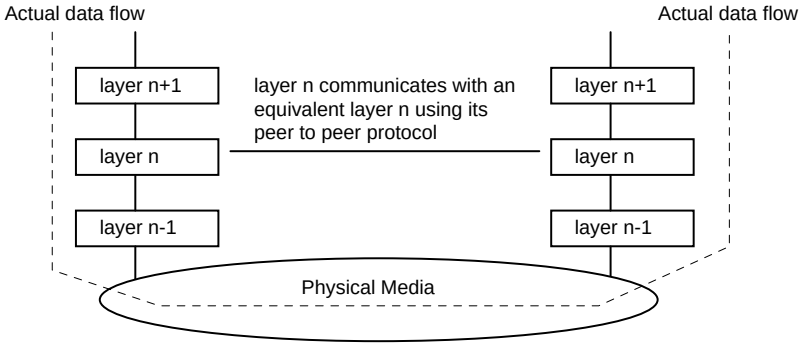

A peer to peer protocol does not consider data flow vertically within the model and simply assumes that the layer n protocol in one end-station communicates directly with the layer n protocol in a different protocol as if they were directly connected. This protocol is defined in terms of the commands and responses it issues together with a description of the protocol data unit format to be used.

c)

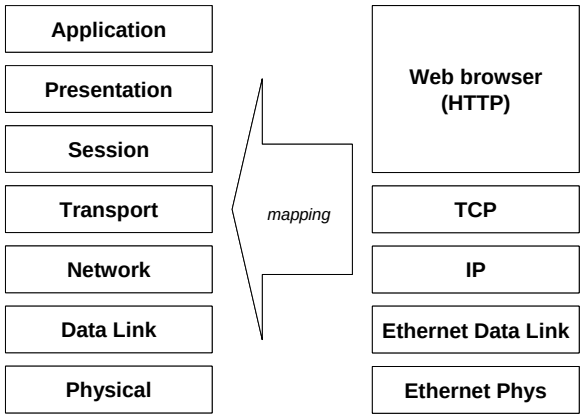

Examiner's Comments

Part (a): Most students scored the 7 marks for correctly labelling each layer. A fair proportion got full marks by also identifying correctly which layers were related to the application and which were related to the network infrastructure. Many more scored some of these final two marks. A substantial number were inclined to split the transport layer between the two alternatives. However, as the model answer shows the transport layer does belong to the network infrastructure – it is the first level that is responsible for end-to-end delivery of the data. On the other hand, the functions of the session and presentation layers are clearly related to the application.

Part (b): This part was also done well, on the whole. One problem is that some students misinterpret the term 'peer to peer protocol' to refer to a situation outside the direct context of the OSI model. Microsoft, for instance, use the term peer to peer to distinguish a set-up where all the computers have the same status from one in which there is at least one machine which is the server. Some students deal with this sense of 'peer to peer' in their answers. However, the context of the question is pretty clear, so if the students understood what the term 'peer to peer protocol' meant in terms of the OSI model they would presumably answer in the correct sense.

Part (c): Part (c) was more challenging, since it involved mapping the OSI model to a real world situation. But many students managed to get at least some marks for mapping the given scenario to the OSI model and a few got high marks.

Question 2

a) What are the basic principles of operation of a Frame Relay network?
(10 marks)

b) When a Frame Relay virtual circuit connection is established between two end-stations, a number of traffic characteristics, such as the *committed burst size*, are specified within the Call SETUP message. Why are parameters of this type important to allow a Frame Relay network to provide dynamic bandwidth allocation?
(10 marks)

c) What is the function provided by the Data Link Connection Identifier (DLCI) ?
(5 marks)

Answer Pointers

a) A Frame Relay network provides a connection orientated service. An end-station connects to the network through a single point of attachment but is able to establish multiple virtual circuits over this physical connection. The total bandwidth assigned to support these connections is limited by the physical capacity of the point of attachment but, within that limit, bandwidth assignment to each virtual circuit can vary with time to dynamically adjust to the needs to each connection.

An end-station must therefore begin by requesting a connection with a remote end-station. This connection will define the traffic characteristics required within what is termed a traffic contract. Once the connection is established, data is passed as a sequence of unacknowledged packets. Each frame relay simply passes the packet onto the next relay. Errors detected along the route will simply result in the packet being discarded. The only feedback that that the network provides to the end-stations is that relating to the existence of traffic congestion. Here notification bits are set to indicate that the packet has passed through an area of traffic congestion; it is then up to the end station to decide whether or not to moderate their rate of traffic generation.

On conclusion of the data transfer phase, the connection must be formally terminated..

b) One of the key objectives of the Frame Relay network is to provide the end-station with a dynamic level of bandwidth that matches the demands of the application or service being delivered. To achieve this it is important to know in advance the bandwidth demands of each connection. These are described within what is known as a traffic contract. This contract will define the average traffic level to be generated, the burst size which defines a level above the average that should be transmitted and an excess burst size that should be accommodated only if spare capacity exists within the network at that time.

The traffic contract is therefore defined by a series of parameters such as the committed burst size and these are conveyed within the initial connection establishment. On receipt of the Call SETUP message, each frame relay will examine the traffic contract parameters and will only allow the SETUP message to progress if that relay is able to support the traffic demands of the connection. Should a relay be unable to satisfy this traffic demand, then the call request will be terminated at that point.

During the course of the data transfer phase, the bandwidth utilisation per virtual circuit is monitored and compared against the traffic contract. If at any point the data being transmitted exceeds the limits imposed by the traffic contract then packets will be dropped to bring the traffic back within limits.

c) When a virtual circuit is established there is a need to differentiate it from others. This is achieved by assigning a unique Data Link Connection Identifier (DLCI) to a virtual circuit. The DLCI can be 10, 16 or 23 bits in length. However, whilst each virtual circuit needs to be uniquely identified, the DLCI need only be unique on each section of a connection, i.e. between an end-station and the Frame Relay; between Frame Relays. Along the full length of a virtual circuit several DLCIs can be used to define that circuit.

Examiner's Comments

Overall: There is a lot of fairly difficult detail for the students to master to answer questions on frame relay, so it is not surprising that the marks scored were generally lower than for Question 1. Sometimes the students give quite long answers that have correct points in them (for instance explaining the evolution from X.25 to frame relay) that are nonetheless unrelated to the actual questions. Sometimes they get mixed up between the principles of frame relay and some other network technologies – as, for example, when they stress that frame relay guarantees error free delivery (which is not true at all).

Part (a): In general students were able to score some marks here, but were not able to furnish sufficient accurate detail to score very highly.
Two students (from different centres) provided responses that followed the indicative answer word for word for a significant portion of the answer, and then followed each other word for word for most of the rest of the answer. The most likely explanation is that a standard source was used to provide the indicative answer to part (a) – as given above. However, there is a question as to whether it is reasonable to suppose that students could give such word-perfect answers without using a crib. There is, obviously, no objection to students memorising answers.

Part (b): This part of the question required a good deal of precision in order to score substantial marks. A few students achieved this and most managed to say something relevant and to score a few marks.

Part (c): This question was a little easier in the sense that 2 marks were allocated for indicating that the DLCI was used to identify each virtual circuit in frame relay. On the other hand the final 3 marks were more difficult because it was necessary for the students to know that the DLCI did not have to be unique along the whole length of a virtual circuit but only for an individual section of the connection.

Question 3

a) What is the difference in the quality of service offered to applications by the TCP and UDP protocols?

(8 marks)

b) What are TCP and UDP port numbers used for?

(5 marks)

c) Explain the basic principles of how TCP is able to overcome transmission errors and ensure a reliable end to end data transfer.

(12 marks)

Answer Pointers

a)
UDP
UDP operates above IP and provides a connectionless, unacknowledged service to the upper layers. There is no error detection or recovery and no flow control provided by UDP. It is a best effort service.

TCP
TCP operates above IP and provides a connection orientated service in which end to end data transfer is guaranteed with flow control and congestion avoidance capabilities. It is a reliable service.

Both TCP and UDP provide services to higher layer protocols however multiple higher layer protocols can be multiplexed onto a single UDP or TCP layer. Each of these higher layer protocols are then differentiated by means of UDP/TCP port numbers. Port numbers are 16 bits in length. Therefore the port number identifies the particular higher layer protocol to which a given data stream is destined. Some of these port numbers are pre-defined, for example port 80 refers to the higher layer http protocol.

b)
Every octet transmitted through TCP is uniquely identified by a 32 bit sequence number which increases by one for each new octet.

c)
Data is transmitted and acknowledged by the receiving end station. Acknowledgements are identified by means of the ACK bit and acknowledgement number within the TCP header.

A positive acknowledgement is indicated by virtue of the fact that the ACK bit is set and then the acknowledgement number will indicate the number of the first non-acknowledged octet. In other words all octets up to an including acknowledgement number -1 have been successfully received.

If data is corrupted or lost in transit then this must be detected by the transmitter. If an acknowledgement has not been received within a given time – determined by a timer – then the transmitter simply sends the data again. It is the responsibility of the receiver to ignore any duplicates it receives. Hence, the transmitter will continue re-sending data until a positive acknowledgement is received.

Examiner's Comments

Part (a): For this first part, where it was necessary to identify the general characteristics of the TCP and UDP protocols students scored quite well and many scored full marks or nearly full marks.

Part (b): A lot of students were not able to explain the concept of ports effectively. If they did not know or say that port numbers are 16 bits in length then they would lose two marks. But it should have not been difficult for them to explain that ports identify the higher layer protocol to which the data stream is destined (or something that would show a basic understanding of that idea in other terms). A number of the students referred to the idea of 'data types' in connection with ports. From a certain point of view one could see the data going to different applications as different 'data types' but this was not really explained clearly by the students who used that concept.

Part (c): A lot of the responses to this question were weak. There were a few very good ones.
It is surprising that there should be so many weak responses since this question deals with a very basic aspect of contemporary networks. One the basis of previous experience in marking these papers, I expected this to be a question of which most of the students would score quite well.
One specific issue that came up repeatedly was that students had the view that TCP, in setting up a connection, creates a definite path through the network that all datagrams have to follow, and this explains why datagrams 'do not arrive out of order'. In fact TCP sets up a <u>virtual</u> connection and it uses IP to actually carry the datagrams over the network, so that individual datagrams are routed individually and may arrive out of order. The receiver is expected to put the TCP segments that travel inside the datagrams back into order if necessary.

Question 4

a) In the context of IPv4, explain the differences between Class A, Class B and Class C addresses.

(6 marks)

b) Explain what is meant by a *netmask*. A network manager is instructed by his supervisor to install four small IP networks, each with 32 computers attached, connected by a router. The manager is only allocated one Class C network. How can the network manager make use of netmasks when configuring this network? Identify one possible suitable value that could be chosen for the netmask.

(10 marks)

c) Routing within small networks that use IPv4 is often handled using distance vector protocols. Explain the general behaviour of distance vector protocols. Your answer should draw attention to problems which can occur within such networks.

(9 marks)

Answer Pointers

a)
The candidates should explain that IP addresses consist of a network number part and a host number part and that addresses are 32 bit numbers . They should then explain that rather than having a fixed split of bits between network/host, this varies with class A having 8 bits in network part and 24 in host (thus 2^24 hosts per network), class B having 16/16 (thus 2^16 per network) and class C having 24 bits in network part and 8 bits in host (thus 2^8 per network). If they miss some detail of the above, but give us extra information such as the values of leading bits in network part, of the fact the two of the host numbers are reserved for broadcast and network number etc we would still award all six marks.

b)
The candidates should explain that an early improvement to the fixed bit allocation of class based addresses was the introduction of netmasks. The netmask is thus being used to indicate which of the address bits are considered to belong to the network part, the rest being the host part. The candidates will tell us that the netmask is a 32 bit binary number with 1s in the positions related to the network part. They candidates will then explain that the choice of a suitable netmask can allow the manager to split up his single Class C network into four subnetworks . The candidate will note that he has to configure his router and hosts with this knowledge but that routers external to his site will not need this information . The netmask of 255.255.192.0 should be proposed.

c)
Distance vector protocols operate by routers making announcements of networks which they can reach together with the distance (metric) of their route to that destination. The announcements are normally made at regular fixed intervals, typically about once every 30 seconds. The major problem that occurs with such protocols is the formation of routing loops .

Examiner's Comments

This question was attempted by about 68% of the candidates. Around 10% of those who attempted the question presented very good answers. Unfortunately, around 15% produced quite poor answers. Part a) was in general answered reasonably well with many candidates identifying the key aspects of the answer. In part b), many candidates failed to notice that the number of computers on the sub networks was 32 and thus, allowing for other IP addresses reserved as the network number and the broadcast address, there is a need for 34 in each block. Thus, the class C network has to be split into four blocks, each with 64 potential addresses. For students who missed this key aspect, marks were still awarded to match the understanding and knowledge that was presented. Part c) of the question was answered poorly by many candidates. In several cases, there was clear confusion in the candidate's answer mixing aspects of distance-vector routing with link-state routing. Some candidates correctly identified slow convergence times as another potential problem and marks were awarded appropriately.

Question 5

a) In a traditional Ethernet, built using co-axial cable, there are rigid requirements that govern the maximum size of the network. Produce a description, accompanied by a sketch diagram, to illustrate a maximally sized network and the components involved.

(10 marks)

b) Describe the general behaviour of the CSMA/CD algorithm used to control media access within an Ethernet. Your answer should also cover the behaviour of the *truncated binary exponential backoff* algorithm used when collisions occur.

(15 marks)

Answer Pointers

a)
Four marks will be awarded for a reasonable diagram and six marks for the accompanying description. Considering the ten marks together they will be allocated as follows. The candidates should note that maximum co-axial segment length is 500m. The candidates should show that the maximum path can be at most 2.5Km (2500m) and that this can be made up of three general 500m segments with machines attached but that the final two segments should be link segments without hosts attached . The segments are joined by repeaters and thus a maximum of four on the longest path. The candidates should identify the presence of terminators , transceivers (MAUs), host interfaces and transceiver drop cables (AUI, Attachment Unit Interface cable) .

b)
The candidates should comment that CSMA/CD is a completed distributed approach with no central control station involved (2 marks). Individual hosts with data to transmit follow the algorithm:
1/. Check if media is busy
2/. If busy, go back to 1/.
3/. If not busy, commence transmission.
4/. Monitor during transmission to check if media condition indicates a collision
5/. If no collision, work finished

6/. If collision, then issue JAM
7/. After sending JAM, delay a random amount of time then go to 1/.
The students should then expand on the behaviour of the *truncated binary exponential backoff* algorithm. They should explain that the delay is calculated based on a random number of "slot times", where a slot time is the essentially the round trip time of a maximally sized Ethernet, namely 51.2 microsecs. The random number of slot times to delay (r) is calculated as $0 <= r < 2^k$ where k is selected as the min (10,n), where n is the re-transmission attempt number. (6 marks)

Examiner's Comments

This question was attempted by about 62% of the candidates. None of the answers presented were of high quality and a large number were very poor. Part a) was the best answered part of the question; but even so, many of the answers were incomplete. The answers to part b) lacked a lot of detail. Very few candidates demonstrated any detailed knowledge or understanding of the behaviour of the *truncated binary exponential backoff* algorithm.

Question 6

a) Simple serial communication systems often use a method known as parity checking as a means of error control. Explain what is meant by a parity bit and how it is calculated.

(5 marks)

b) Assume a small block of eight, 8-bit characters are to be transmitted. Describe how a combination of both transverse and longitudinal parity can be used to locate any single bit error that might occur within the block of characters.

(10 marks)

c) It is becoming increasingly important that we can exchange information in a secure manner.

i) Explain the main problem that affects the use of symmetric key cryptography systems.

ii) How does a public key cryptography system overcome the problem you have identified in 6c(i)?

(10 marks)

Answer Pointers

a)
In the simplest form, a parity bit is an extra bit attached to the data bits of a character so that overall some rule governing the number of ones and zeros transmitted is obeyed . There are various approaches to parity. Even parity is when we make sure that an even number of 1s is transmitted and Odd parity when an odd number of 1s is transmitted . The students should also of course comment that a receiver needs to check that the parity of received information matches the agreed rule .

In this question we introduce the idea that parity can be used in more sophisticated ways. Transverse parity (row parity) is the more complete term for single parity bits being attached to each character as in the previous part of this question. Longitudinal parity (column parity), is the term used when an extra "parity character" is added at the end of a small block of characters. Each longitudinal parity bit is set so that some agreed rule is met by all bits in that position, one from each character of the block. The candidate should then describe, probably by means of a simple example, how this can be used to spot a single bit error and indeed its location.
In the example I give here, I have 7 bit characters and even row parity. I have four characters followed by the longitudinal parity character using odd parity.

This represents a "correct" transmission

pbit	**d0**	**d1**	**d2**	**d3**	**d4**	**d5**	**d6**	
1	**0**	**1**	**0**	**1**	**0**	**1**	**0**	**character one**
1	**0**	**0**	**0**	**0**	**1**	**1**	**1**	**character two**
1	**1**	**1**	**1**	**0**	**0**	**0**	**0**	**character three**
0	**1**	**0**	**1**	**0**	**1**	**0**	**1**	**character four**
0	**1**	**1**	**1**	**0**	**1**	**1**	**1**	**long. parity**

This represents a transmission with d4 of character three corrupt

pbit	**d0**	**d1**	**d2**	**d3**	**d4**	**d5**	**d6**	
1	**0**	**1**	**0**	**1**	**0**	**1**	**0**	**character one**
1	**0**	**0**	**0**	**0**	**1**	**1**	**1**	**character two**
1	**1**	**1**	**1**	**0**	**1**	**0**	**0**	**character three**
0	**1**	**0**	**1**	**0**	**1**	**0**	**1**	**character four**
0	**1**	**1**	**1**	**0**	**1**	**1**	**1**	**long. parity**

Checking our reception we see that the transverse (row) parity bit for character three does not meet the even parity rule we chose. Similarly, we note that the longitudinal (column) parity bit for d4 does not meet the odd parity rule. We have thus located that bit d4 of character three is corrupt. If we assume only single bit errors can occur, we can of course thus correct this.

c)
i)
Candidates should first observe that, by definition, a symmetric key cryptography system uses the same key for both the encryption process and the decryption process. Thus, to be able to decrypt a message that has been received, the receiver must be in possession of the key . The main problem therefore is that of key distribution. The key needs to be distributed in a secure manner and the key needs to be replaced if there is hint it may have been compromised .

ii)
Candidates should first observe that, by definition, public key cryptography systems use a different key for the encryption process (the public key) to that used by the decryption process (the private key). You can therefore publish your public key to anyone who can then use it to encrypt a confidential message before sending it to you . As the private key is needed for decryption, only you can decrypt the message and there is no problem if the message itself is intercepted by other people.

Examiner's Comments

This question was attempted by about 45% of the candidates. A small number of the answers were exceedingly good, gaining full or almost full marks. However, a large number of the answers offered were very poor. Part a) was answered reasonably well. A small number of students offered good answers or part b). Many of the answers to part b) were somewhat confused or incomplete. Many candidates did not offer an answer to part c). Of the answers that were presented, some identified most of the points expected whereas other students that were somewhat confused or missed important aspects.