THE BCS PROFESSIONAL EXAMINATIONS
Diploma in IT

October 2006

EXAMINERS' REPORT

Computer Networks

General

The overall performance in this examination has significantly increased with the total pass rate in October 2006 presentation just above 70%. Together with the April pass rate of over 60%, this is very encouraging indeed. It appears that the performance in this paper has been improving steadily. Questions 2, 4 and 5 were favourites and question 2 was attempted by over 80% of the students. Quite few discussions were well presented and the answers were of high quality. However, there are still some 30% of students who have been under-performing, and it is hoped that as the performance steadily improves, this percentage will come down.

Question 1

1. *a)* A transmission system uses a coding scheme that defines a 'symbol' as a voltage that can have one of four possible values. If the system operates at a transmission rate of 200 symbols per second, determine the data rate measured in:

i) Baud
ii) Bits per second **(6 marks)**

Show, with the aid of a diagram, how the bit pattern '10101110' would be encoded using Manchester encoding. **(9 marks)**

b) The Link Access Protocol (LAP-D) used within ISDN has a frame structure that begins and ends with a flag sequence of '01111110'. Show how zero-bit insertion ("bit stiffing") is used to ensure that this flag sequence cannot occur elsewhere within the frame. **(10 marks)**

Part a) Answer Pointers

A symbol is a voltage level that can have one of four possible values. Four unique values represents the combinations generated by two bits, i.e. $2^2=4$. Therefore one symbol equals two bits. One baud is one symbol per second.

If the transmission system operates at 200 symbols per second then the data rate is:

200 Baud

400 bits per second

Marking Scheme: 2 marks for determining that one symbol represents two bits; 2 marks for the speed in baud and 2 marks for the speed in bits per second.

Part b) Answer Pointers

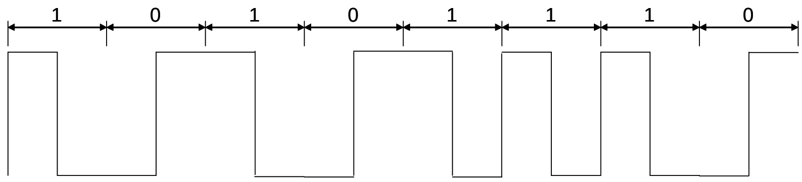

Marking Scheme: 2 marks for the logic 1 encoding, 2 marks for the logic 0 encoding, 5 marks for the overall bit pattern produced.

Part c) Answer Pointers

In order to avoid the flag sequence appearing elsewhere within a data stream zero bit insertion is used. When five consecutive 1s are detected then an additional zero bit is inserted in the data stream as shown below. At the receiver, when five consecutives 1[st] are detected and the next bit is a zero then that zero is removed; again as shown below. In the example shown below five consecutive 1s are detected twice and so two additional 0s are added to the transmitted data stream.

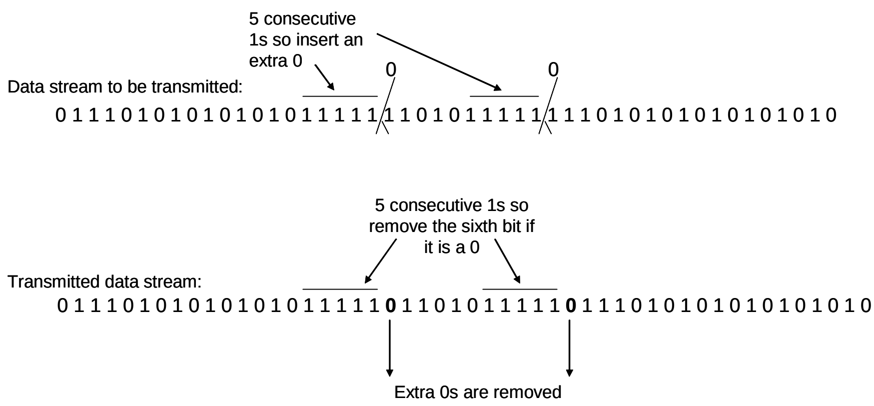

Marking Scheme: 5 marks for the insertion process – detect five consecutive 1s and add an additional 0; 5 marks for the removal process – detect five consecutive 1s and remove the next bit if it is a 0.

Examiner's Comments

The students found the first part of part (a) difficult. Most did not know that the baud rate was simply the same thing as the number of symbols per second and could not work out the bit rate from the information supplied. However, a few students did understand the question and answered it correctly.

The second part of part (a) was answered correctly by most of the students who attempted it. There is some confusion over which way the logic works – whether a zero is a high-low transition or low-high transition – and the same, but opposite, problem for a one. A surprising number of students were well aware of this problem and one went so far as to do the question both ways. More usually, they quoted which book they were using to determine which way the encoding works. Some pointed out that it differs between different books (e.g. Stallings and Tanenbaum). Halsall says that either way is possible in the case of differential Manchester encoding, but not, admittedly, in the case of (non-differential) Manchester encoding.

Under the circumstances I gave the marks either way. The students who attempted this question were, in general, well schooled in the issues involved.

In respect of part (b) of question one, about bit stuffing, a good proportion of students understood what was involved. In that case they scored the marks easily. On the other hand a substantial proportion of the students did not understand the issue.

Overall, therefore, the marks were not so high on this question (since many students could only answer parts of it). But those who attempted it nearly all scored some marks from it.

Question 2

2. By considering the ISO Reference Model, explain what functions are performed by the Data-Link and Network Layers. **(8 marks)**

i) What is meant by the term '*peer to peer protocol* '? **(5 marks)**

ii) Consider two wide area networks; one using IP and the other using X.25 as its Network Layer protocol. Explain how the quality of service offered by these two networks differs. **(12 marks)**

First part) Answer Pointers

Data-Link layer

- provides data transfer over a point to point link, e.g. LLC, LAP-D
- defines the protocol for transmitting frames over the attached network, i.e. IEEE 802 MACs
- offers both a connectionless (unreliable) or connection orientated (reliable) data transfer service
- provides network dependent local addressing, i.e. IEEE 802 MAC address

Network Layer

- responsible for inter-networking
- includes provision of a network wide addressing scheme
- routing traffic through an inter-network
- providing a connectionless or connection-orientated service to the Transport layer.

Marking scheme: 4 marks for the data link layer and 4 marks for the network layer.

Second part) Answer Pointers

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.

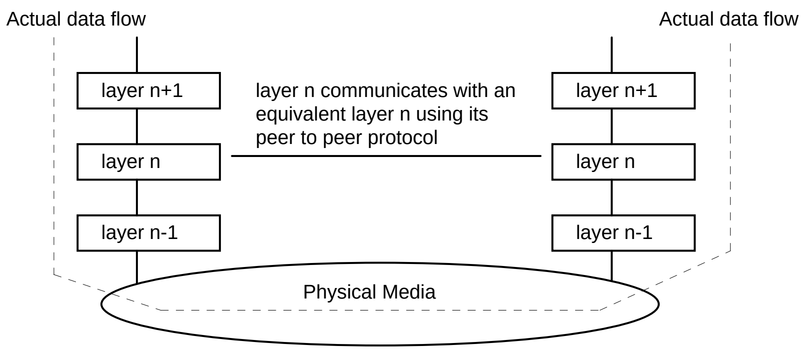

A peer to peer protocol does not consider data flow vertically within the model and simply assumes that the layer n protocol in one end-station communicates directly with the layer n protocol in a different protocol as if they were directly connected. This protocol is defined in terms of the commands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

Marking Scheme: 3 marks for recognising that a peer to peer protocol operates horizontally between layers at an equal position within the model, i.e. layer n to layer n. 2 marks for describing how such a protocol is defined in general terms.

Third part) Answer Pointers

QoS offered by IP

- connectionless
- datagram service with a fragmentation capability
- no error correction – detection and discard only
- no flow control
- best-effort service/no guarantees

QoS offered by X.25

- connection-orientated
- full error detection and correction by packet re-transmission
- flow control provided with both local and end to end significance
- virtual circuits established with logical channel number assignment to differentiate them
- guaranteed delivery

Marking Scheme: IP: 2 marks for noting that it is a best effort service; 2 marks for it being a datagram service with no error correction and 2 marks for identifying that it does not provide flow control. X.25: 2 marks for noting that it offers a reliable data transfer service using virtual circuits; 2 marks for error recovery through packet re-transmission and 2 marks for identifying that it provides flow control.

Examiner's Comments

First part: Very many students understood the role of the data-link and network layers in the OSI model quite adequately and answered this well.

Second part: Quite a lot of students understood what is meant by a peer-to-peer protocol in terms of the OSI model – but probably a minority. Some of the students answer by saying that a peer-to-peer protocol is 'joining two computers directly' (sometimes meaning in a physical sense). Others think it is to do with having a network that does not have a server, so that all the computers are peers. This does have some justification, but, even in that case, there is not a specific protocol that is involved. So no marks were given unless the students addressed the issue as intended.

Third part: A very substantial number of students were able to address the issue of the difference in quality of service between an X.25 network and an IP network is a satisfactory way. Some dealt with the issue in more developed way and some in a less deep way. But, since there were 12 marks at issue, it was possible to award marks in an appropriate way.

Generally, the marks were good for this question. Even if all the marks were lost for the second part, that was only 5 marks, so most of the marks could still be gained.

Question 3

3. *a)* What functions are performed by a router? **(8 marks)**

b) Why is it important for routers to know about all of the possible routes through a network topology? **(5 marks)**

c) Explain the main differences between the operation of routing protocols based on the following principles:

i) Distance vector routing (Bellman Ford)
ii) Link state database routing **(2 x 6 marks)**

Part a) Answer Pointers

A router is responsible for three key functions. These are traffic routing where the addressing information contained within a received packet/datagram is examined to determine how it should be forwarded. This requires a router to maintain a routing table that allows it to determine the next hop (outgoing port) on which to send a received packet/datagram. This table is constructed through the use of a routing protocol that operates between routers and from which routes through the network can be determined and classified.

The second function is that of quality of service control through packet classification/prioritisation. In today's networks it is increasingly important for traffic to be differentiated and a different quality of service offered to higher priority traffic. Routers use fields within received packets/datagrams – such as addresses/port numbers/protocol fields – to classify each packet or traffic flow. These are then assigned to different output queues from which traffic is transmitted based upon a scheduling algorithm. High priority queues would expect to use more of the available bandwidth on the outgoing port.

The third and final function is providing connectivity of different network technologies. A router operates at layer 3 of the ISO model and as such, is independent of the underlying hardware of the network. In this way a router can connect, for example, a Token Ring LAN to a CSMA/CD LAN, or a CSMA/CD LAN to a wide area network, such as ISDN.

Marking Scheme: 3 marks for traffic routing; 3 marks for QoS; 2 marks for connecting different network technologies.

Part b) Answer Pointers

Within a network, the cost of a particular link can change due to increased traffic levels. Hardware failure can also result in increased error rates on a link and ultimately for complete link failure. The key point of an efficient routing protocol is to ensure that a route can always be found, even if certain links have failed. Hence, it is essential for a router to know all possible routes through a network such that, if the cost of a link increases or a link fails, an alternative can be used without having to re-compute a new router by communicating with the other network routers. Recovery from link failure is therefore not only transparent to the end-stations but also occurs quickly.

Marking Scheme: 3 marks for the reason for why routes cease to remain available; 2 marks the need to efficient re-routing when links fail.

Part c) Answer Pointers
Distance-vector routing:

- Each router measures the cost of the links to each of its adjacent routers.
- This information is then compiled into a table.
- Each router then sends a copy of its link costs table to each of its adjacent routers. Note that a router will always record a cost of zero to itself.
- A router analyses the tables it receives from its adjacent routers and where multiple routes are available between the same two points, the lowest cost route is chosen and this one only is recorded.
- The router's table therefore contains one route to each destination point and this route is the lowest cost of those that exist. In other words, routers do not maintain a complete record of all possible routes through a topology.
- Routers continue to exchange their tables on a regular basis.

Link-state database routing:

- Each router determines the costs to each adjacent router and the identity of each attached network. This information is built into a 'link state database'.
- Each router then broadcasts its link state database to each adjacent router.
- These routers then re-broadcast each received link state database on all links except for the one on which it was received. Hence, this 'flooding' mechanism ensures that every router will receive a copy of the link state database of every other router.
- In this way, each router will receive a copy of the link state database of every other router. From these each router can determine a complete map of the network topology with all routes identified.
- The shortest routes through the topology are then calculated from the complete network map using an algorithm such as Dijkstra's algorithm.

Marking Scheme: Distance vector – 6 marks for a general description as above, the key point being that each router maintains a partial network map, showing only the shortest routes. Link state – 6 marks for a general description as above, the key point being that each router maintains a complete map of the network showing all possible routes but that the shortest routes are determined by a routing algorithm.

Examiner's Comments
Part (a): This question was done quite well. Nearly every student who answered had some basic idea of what a router does. On the whole the basic functions of a router were explained quite well. In addition a fair number of students understood that modern routers perform extra functions like security (through access lists) and marks were allocated to that in the marking scheme.

Part (b): This was also fairly straightforward. Some of the students missed the most basic idea that the advantage for a router of knowing all the routes through a topology is very clear if a link goes down, since it has to seek a different route in that case. Many students thought that dealing with congestion was the main issue, and that is an issue, so marks were given for that.

Part (c): The difference between distance vector routing and link state routing is more technical and answers to this were more variable in quality. Not very many of the students really understood that in the latter case each router can form an idea of the whole topology of the network or what difference there is in the two cases in terms of the messages sent between the routers.

Question 4

4. *a)* Explain, giving reasons, what network device/s you would use in each of the following cases:

i) Linking a LAN in a building to another LAN in the next building so that data frames can be selectively be forwarded from one LAN to another.

ii) Linking a LAN in a building to another LAN in a building situated at the other side of a field so that a number of data frames can be exchanged between the LANs. **(12 marks)**

b) Referring to data frames in *a)* and assuming that the data frequently involves characters, suggest giving reasons a framing format that also incorporates burst error detection information. **(13 marks)**

Answer Pointers

Explanation should be as follows:

a) i) Use of a bridge as a bridge filters input and output traffic so that only data frames intended for another network is are routed to that network. Hence the two LANs behave like a single LAN. Bridges are governed by two factors –the filtering rate and the forward rate. Bridges isolate problems in a LAN and do not allow them to propagate to the other LAN.

6 marks

ii) A bridge attached to a local LAN in one building with a fibre modem and optical fibre connecting the other LAN through another fibre modem. Use of fibre minimises external signal interference as data frames pass between the LANs and provides high bandwidth.

6 marks

b) Burst errors are errors that involve changes to a small set of bits near a single location. CRC codes using a mutually agreed polynomial G(x) are particularly useful in detecting burst errors, better than checksums. Explanation of CRC codes here.

6 marks

In the case of character transmission, distinction has to be made between data and control information and byte stuffing is used. Explanation of byte stuffing here.

4 marks

Hence the frame format:

SOH	Block of data with byte stuffing	EOT	CRC

3 marks

Examiner’s Comments

Part (a) gave two slightly physical different scenarios and asked the students to devise an architecture for a LAN. This was a very open question and the students came up with many divergent answers – most of them workable, even if some solutions were a little exotic. It was a good opportunity for the students to show that they had a grasp of networking architecture and, generally speaking, they responded to that well.

Part (b): some students could explain quite adequately what a burst error was and got marks for that. A reasonable number knew that CRC codes were appropriate for identifying burst errors. On the other hand hardly any saw that the issue of byte stuffing was alluded to in the way the question was posed (although a very small minority did see that). Virtually none of the students posited the frame structure in the model answer. On the other hand many chose Ethernet framing as their answer. Since that is not absurd, a few marks were given for that.

Question 5

5. *a)* Discuss the main characteristics of an ATM network with particular reference to routing of data in the network. You are encouraged to use appropriate diagrammatic illustrations. **(17 marks)**

b) What is congestion and how is it controlled in an ATM network? **(8 marks)**

Answer Pointers

5. a) Discussion points: The major objective of ATM is to integrate real-time data such as voice and video signals and non-real-time data such as computer data. ATM supports several service models, variable frequency (video), burst transfers (computer data).

The ATM layer. ATM is a packet switched, virtual circuit network architecture. ATM packets or cells has 53 bytes with 5 bytes as header (VCI label) and 48 bytes of data load which includes payload type (PT), cell-loss priority (CLP) bit and header error control (HEC) bytes. VCI indicates the virtual channel to which the cell belongs, PT indicates type of payload, CLP priority setting bit and HEC error detection byte.

An ATM protocol stack consists normally of two layers, the ATM layer and the physical layer (PHY). At the end points of the network a further top layer ATM adaptation layer (AAL) is included which is analogous to the Internet transport layer.

A diagram of showing the three ATM layer stacks at the end with a group of two ATM layer stacks including data flow together ATM cell format here.

9 marks

Routing: Before a source can begin sending cells to a destination, the ATM Network should first establish a virtual circuit from source to destination. A VCI is associated with each link on the virtual circuit. A VCI label can change between intermediate switches in the route.

A brief discussion on virtual channels and virtual paths.

A diagrammatic illustration of an ACM network with switches, VCI labels, path taken by ATM cells here.

8 marks

b) ATM ABR (available bit rate) a protocol that takes a network-assisted approach towards congestion control which is markedly different from that of the Internet's TCP protocol. The ATM virtual circuit approach means that each switch on the source-to-destination path maintains state information about the virtual circuit and these states allows switches to track the behaviour of individual senders- tacking the average transmission rate, and take appropriate congestion control actions such as signalling the sender to reduce the packet transmission rate.

8 marks

Examiner's Comments

Again the first part of the question was posed in an open way. The students responded well to that and showed that they did have a fair grasp of ATM networking technologies and its advantages. On the other hand only a few students could really explain the question of routing in an ATM network in a developed way.

Part b): Very few students could explain clearly what congestion was – many confused it with flow control (this is a common mistake). This weakness naturally had a big impact on how they dealt

with the question of how ATM deals with congestion. Many students explained that ATM was very careful in allocating resources before communication takes place, which is true. But they did not realise that congestion can still occur in an ATM network. Obviously, in that case they could hardly explain how ATM approaches congestion control.
On the other hand a fair number of students did have some grasp of how ATM seeks to deal with congestion, and that is impressive since it is a very technical subject.

Question 6

6. *a)* Explain the two main approaches used to control errors in transmitted data streams. **(5 marks)**

b) In a data communication network, a sender encodes all 7-bit ASCII characters using Hamming code before transmission to a receiver. Assuming that a receiver in the network receives the bit pattern 10111001001,

i) Use the Hamming code scheme to check the bit pattern, indicate any error and show how the error can be corrected:
ii) Indicate the ASCII character transmitted
iii) Determine the code efficiency of the encoder
iv) Outline the limitation of the Hamming code scheme and suggest how this can be overcome.

(10 marks)

c) What is a digital signature and how is it generated? Briefly explain in what way a digital signature ensures message authentication. **(10 marks)**

Answer Pointers

6. a) The two main approaches: - Error detection
- Error correction

Error detection approach uses only enough redundant information to allow the receiver that to deduce that an error has occurred.

2 marks

Error correction approach uses enough redundant information along with each block of data sent to enable the receiver to deduce what the transmitted data must have been.

3 marks

b) Data received: 10111001001
Check bits-powers of 2. I.e. 1,2,4,8..= 1011

i) Modulo-2 sum of data = 1000
Hence bit 8 error
ii) The 7-bit ASCII code= 10110000
iii) Code efficiency= 7/11 = 63.6%
iv) Hamming code limitation: can only correct single bit error and code modification essential to correct burst errors.

10 marks

c) Points for discussion:

Digital signature authenticates that a message received from a sender is indeed come from that sender.

1 mark

Generating digital signature: the sender uses a hash function to convert the message into a message digest which is then encrypted with sender’s private key. This is digital signature.

3 marks

The message along with the digital signature is sent to the receiver. The receiver decrypts the digital signature using the public key to obtain the message digest. A message digest is then created by the receiver using the same hash function and the two message digests are compared. A match proves the authentication of the message.

6 marks

Examiner's Comments

Part a): In the main, students understood that error detection and error correction were the two main approaches to dealing with errors and know the difference between them in a fairly clear way.

Part b): A good number of students knew how to do this question, which involved checking a group of digits received to see if there was an error, using the check bits that were sent. However, quite a lot of students wanted to encode the bits first, even though they were already encoded. In other words they were trying to do the question by rote, without having read in properly. No marks were given in the latter case.

Part c): On the whole this was well done. Although the use of private and public keys for encoding a message and giving the receiver a digital signature at the same time (to authenticate the sender) is quite involved, a fair proportion of students had understood the matter properly, although some efforts to answer this, understandably, got quite confused.