THE BCS PROFESSIONAL EXAMINATIONS
BCS Level 5 Diploma in IT

April 2007

EXAMINERS' REPORT

COMPUTER NETWORKS

General

The overall performance of students in recent three examinations has improved, which is encouraging. But there is still uneven performance within answers to questions mainly due to not spending sufficient time reading questions, and thus responding with wrong answers. Focusing on reading and understanding what is expected as answers would improve the quality of response and performance as a whole.

Question 1

a) Briefly explain the operation of the Address Resolution Protocol (ARP).

(10 marks)

b) Two LANs, L1 and L2 are interconnected by a Router. A computer, C1, located on LAN L1 wishes to communicate with a server, S2, located on LAN L2. Initially C1 knows the IP address of S2 but not its MAC address. Hence, C1 uses the ARP protocol to obtain the appropriate MAC address to use. Determine the MAC addresses that would be used in the LAN frames that carry the IP datagram from C1 to S2.

(6 marks)

c) Referring again to the network described in part (b), if the server is able to support more than one TCP connection with C1, explain how TCP port numbers can be used to differentiate between these two connections.

(9 marks)

Answer Pointers

a) The ARP protocol allows nodes to determine both the MAC and IP address of another node.

Consider two nodes, A and B. A wishes to determine the MAC address of B.

Node A issues an ARP request. This will be carried in a LAN frame with the destination address set to broadcast. The data field of this LAN frame will contain the ARP request PDU. This PDU contains the IP address of B, the IP address of A, the MAC address of A and a blank field to indicate that the MAC address of B is required.

All devices on the LAN will receive this ARP request (MAC broadcast). They will each examine the ARP PDU and check the destination IP address. Only the device that has that address will respond, i.e. B.

Node B will therefore issue an ARP reply. This will be transported within a LAN frame that has a destination address equal to the MAC address of A and an ARP reply PDU in the data field. This PDU will contain all of the addresses that were in the original ARP request, however, B will add its own MAC address into the blank field.

On receipt of the ARP reply, A has discovered the MAC address of B.

b) Note that C1 is on LAN L1 and S2 is on LAN L2 and the two are connected by a Router. Two LAN frames are needed to carry a datagram from C1 to S2:

C1 to the Router and the Router to S2.

C1 to the Router: MAC destination = the MAC address of the Router; MAC source = MAC address of C1.

Router to S2: MAC destination = the MAC address S2; MAC source = MAC address of the Router.

c) The server has only one IP address. Hence it is important to be able to differentiate which IP datagram is intended for which connection.

TCP port numbers allow for multiplexing at the Transport layer. TCP port numbers are 16 bit numbers and each TCP PDU contains both a source port and destination port number. In this way, port numbers are the equivalent of a source and destination address at the Transport layer.

Hence, each connection is assigned a unique port number within the server – call these TP1 and TP2. Therefore one connection from C1, will be distinguished by using the TP1 port number and the other by using the TP2 port number. Both of these will be carried over the same IP datagram flow but when they arrive at the server, the server's TCP protocol will separate the PDUs based on their destination port numbers.

Examiner's Guidance Notes

A fair number of students answered a) very well. On the other hand a significant proportion of students did not seem to know exactly what the Address Resolution Protocol was for. So the marks tended to be high or very low. Most students did not know the answer to b), even if they did understand the basic operation of ARP. The reason was that they did not understand clearly that the router segments the network and that the ARP operation is carried out twice – once (by the client) to find the router's MAC address and once (by the router) to find the server's address. Some students realised that routers do not forward broadcasts and that therefore ARP could not work across the router anyway. In c), a reasonable proportion of the students understood the point about the way that TCP port numbers can be used to differentiate between two or more connections – not all at quite the same level.

Question 2

a) Explain the operation of the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Local Area Network access protocol.

(10 marks)

b) Show by means of a diagram the frame format used within the IEEE 802.3 CSMA/CD LAN.

(9 marks)

c) Why does the CSMA/CD LAN impose both a minimum and a maximum size frame limit?

Answer Pointers

a) The access protocol follows the following set of procedures.

A node wishing to transmit will:

- check that the network is free from transmission (carrier sense)
- it will wait if a transmission is detected
- if the network is clear then the node will start to transmit
- it will continue to monitor the network whilst transmitting its own frame
- if another transmission is detected whilst the node is transmitting then a collision is said to occur
- the node will continue to transmit to ensure that the collision is spread across the whole network and will then stop
- the node will wait a random period of time before attempting a re-transmission
- if no collision occurs then the node completes the transmission of its own frame.

b)

Destination Address (6 bytes)	Source Address (6 bytes)	Length (2 bytes)	Data (46 to 1500 bytes)	PAD	CRC (4 bytes)

c) Minimum:

A key part of the operation of CSMA/CD is the ability to detect the presence of other transmissions and collisions. The worst case is that a node begins transmission and that this transmission has to propagate along the full length of the network. If another node begins to transmit just before this transmission reaches it then a collision will occur at the far end of the network. It is therefore crucial for this collision to be detected and that won't happen until it has propagated back along the full length of the network. This 'round trip delay' is a key network parameter and it is therefore important for all nodes to transmit for at least this time. Hence, this defines a minimum frame size.

Maximum:

Once a node has been transmitting its frame for a time greater than the round trip delay then no collisions will occur and this node has full control over the LAN. All other nodes will detect the transmission and simply be waiting for the network to go quiet. In order to ensure that no one node is able to take control of the network and lock other nodes out then it is important, it does not transmit for longer than a defined maximum time. This therefore sets a maximum limit on the frame size.

Examiner's Guidance Notes

a) & b) most students who attempted these questions answered well.

c) This part was a little harder, but still a fair proportion of students did know why there is a maximum and minimum frame size on a CSMA/CD LAN. So, overall, the students scored well on this question and it was the most popular question. Nearly 90% of the students attempted this question. That was higher than the proportion of students attempting any other question. The average score on this question was 15.71. That was considerably higher than the average score for any other question.

Question 3

a) Show by means of a diagram, the cell format used within the ATM network.

(5 marks)

b) What is the difference between a Virtual Path and a Virtual Channel?

(6 marks)

c) Explain the purpose and function of the ATM Adaptation Layer (AAL) protocol.

(9 marks)

d) What is meant by the Available Bit Rate (ABR) service?

(5 marks)

Answer pointers

a)

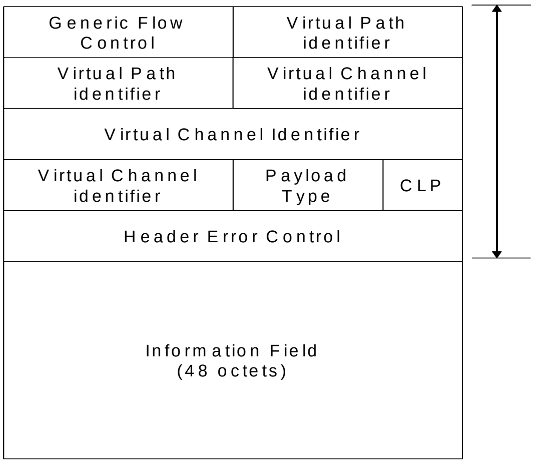
Generic Flow Control
Virtual Path identifier
Virtual Path identifier
Virtual Channel identifier
Virtual Channel Identifier
Virtual Channel identifier
Payload Type
CLP
Header Error Control
Information Field (48 octets)

b) A virtual channel defines a single point to point connection, identified by its virtual channel identifier (VCI). A virtual path however, is a bundle of virtual channels that share the same end-point. Hence, a virtual path can be considered as a container that contains several virtual channels. Each virtual path is identified by its unique virtual path identifier (VPI).

c) The lower layers of the ATM protocol suite are responsible for the transmission of the 53 octet ATM cells. At the higher layer we have the applications between transported over the network. Hence, there is a requirement in the middle to convert the application to and from an ATM cell stream. This is the function of the AAL layer. The AAL protocols are end to end protocols and hence, only present in the end-stations. The basic function of AAL is to segment data from the higher layer protocols into cells and to reassemble a received cell stream into data structures acceptable by the higher layer protocols. Where an application requires a strict timing relationship to be maintained between communicating end-stations then it is the responsibility of the AAL protocol to achieve and maintain this. Equally, the AAL must overcome the problem of lost cells and provide flow and timing control. In this way, it is the AAL that provides the required Quality of Service demanded by the application.

d) The Available Bit Rate service is one class of service offered by the AAL protocol. It is intended for applications that require a variable bandwidth from the network. Nodes are allowed to transmit cells and providing that capacity exists within the network, these will be carried. However, if the network is unable to carry the traffic being submitted to it, then it will provide feedback to force ABR traffic sources to modify their cell generation rates. In this way the feedback mechanism is able to ensure that ABR traffic sources make best use of the capacity that is currently available within the network. It can be considered as an enhanced form of a best effort delivery protocol.

Examiner's Guidance Notes

A lot of the responses were quite accurate in describing the cell format within ATM, but a great many were very vague. The majority of students understood that a virtual path is a collection of virtual channels, which was an important component of the answer to this part of the question. Their efforts at giving a precise explanation of what a virtual channel actually is were more variable. On the whole there was an understanding shown of the fact that the AAL is responsible for mapping the higher level protocols onto ATM cells and of the fact that the AAL sits between the ATM lower levels and the applications. Less obvious to many of the students was the responsibility of the AAL for the problem of lost cells and the provision of timing and flow control. Many of those who did not know exactly what ABR is suggested the Available Bit Rate was the number of bits available on the network, failing to notice that the question was about the ABR service. So the central concept that this service adapts to variations in the available bandwidth was lost.

Question 4

a) The Data Encryption Standard (DES) is an example of a block cipher.

Taking DES as example, explain:
i) What is a block cipher?
ii) What are the limitations to a block cipher?
iii) What are the advantages of chaining?

(9 marks)

b) Explain the difference between public key and private key encryption.

(8 marks)

c) What are the advantages and disadvantages of the use of a public key in encryption?

(8 marks)

Answer Pointers

a)

i) A block cipher works on fixed-sized blocks of data. With DES, a message is split into blocks of plaintext, each 64 bits long. A 56 bit key is used to encrypt each block of text into a 64 bit block of ciphertext. The blocks of ciphertext are transmitted and decoded by the receiver.

ii) The stream of enciphered blocks could be intercepted and altered by a person knowing the key. They could then add some extra blocks without the receiver being able to tell. In addition, repetitive blocks of text generate identical blocks of ciphertext and this help someone trying to break the code.

iii) With chaining the problems mentioned in (b) can be addressed, since this approach links together the enciphering of each block of text with the preceding block.

b) In a private key arrangement, both the receiver and sender can use the same key. In a public key arrangement the receiver and sender use different keys. The two keys are pairs, in the sense that if a sender uses the public key to encode the message, only the private key can be used to decipher the encrypted message. Also the private key cannot be obtained from the public key. Therefore the public key can be given to anyone who needs to send a message to the holder of the private key.

c) The advantage is much greater flexibility since the public key can be given to anyone who needs it, and the encryption is still secure. On the other hand, if anyone at all can get hold of the public key, then it is not possible to tell if the sender is a trusted source (though this problem can be dealt with by separate procedures of message authentication).

Examiner's Guidance Notes

Most students did not know clearly what is a block cipher and lost most of the marks on this part – including all the sub-parts. A minority of students did know what a block cipher is and explained it well. But few, even of those, understood the principle behind chaining. b) was better done – and done very well by many. But many were misled into thinking of the fact that there are a dual pair of keys in public key encryption and thought that this is what was being asked In(c), there was a simple answer to this question, in terms of the advantages of using a public key approach, provided the students got onto the right track in b). The public key approach is more simple to administrate.

Question 5

a) Explain, in general terms, how the link state approach to routing differs from the distance vector approach.

(10 marks)

b) In the context of TCP/IP give one example of a protocol that embodies the distance-vector approach, and one that embodies the link state approach.

(2 marks)

c) Explain why the link state approach to routing scales better than the distance vector approach.

(13 marks)

Answer Pointers

a) In the link state approach to routing, unlike the distance vector approach, each router had complete information on the topology of the network – i.e. each router knows where all the other routers are in the network and all the interconnections between them. In the link state approach each tests the connections it has to other (directly connected) routers by sending messages at intervals to check these neighbour routers are responding. Other routers will then use that information to update their view of the topology of the network. In the case of the distance vector approach the overall information on the topology of the network is effectively distributed between the routers, and they are not capable of forming an overall view of the topology. Each router sends updates to other routers, but only to routers that are directly connected to it. The information gradually propagates through the network, since each router must be connected to at least one other router.

b) In the context of IP RIP embodies the distance vector approach and OSPF embodies the link state approach.

c) In the link state approach to routing, each router propagates information at intervals – as with distance vector routing. However, the amount of information included in each message is very small in the case of link state routing, whereas with distance vector routing the information generated in the update messages can become very great. Each router periodically tests the connection to each directly connected router, and this is the link state information that is transmitted to the other routers.

Examiner's Guidance Notes

a) was about the difference between link state routing and the distance vector approach to routing. As usual, some students answered this exactly. But the big problem overall was that students did not really understand the distance vector approach. They used the right terminology – pointing out that in the link state approach there was a database of the links covering the whole network (a very important point), whereas in the distance vector approach there was a table and each router had only a partial view of the network. In b), surprisingly, most students could not give the names of protocols that embodied these two main approaches to routing. In c) as a result of confusion, the comparisons made between the two systems ('why does link state routing scale better than the distance vector approach') in c) were also, very often, off the main point.

Question 6

a) Briefly explain the overall approach to security embodied in using firewalls.

(8 marks)

b) Why must all firewalls on the same system be configured the same way?

(5 marks)

c) Explain the idea behind packet filtering and give two examples of criteria that might be chosen as the basis of filtering datagrams.

(6 marks)

d) To use a packet filter effectively, a firewall grants access rather than restricting access. Explain why that approach is more effective.

(6 marks)

Answer Pointers

a) A firewall is the standard mechanism for controlling internet access, in order to protect an organisation's network from outside interference – i.e. to stop outsiders gaining access to restricted information, changing information or interfering with an organisation's intranet. A firewall is placed at the connection(s) to an external network like the Internet. The configuration of a firewall depends on the policies of the organisation, and these differ from organisation to organisation.

b) All the firewalls taken together establish a security perimeter (there needs to bone firewall at each external connection). Unless the same policy is pursued consistently with all the firewalls, it may be possible to get round a restriction placed on one firewall by finding a different point of entry.

c) A filter is a blocking mechanism – often available in a commercial router. Datagrams can be blocked using different criteria. Each datagram is considered separately and no record is kept of the blocked datagrams.

d) The number of well known ports is growing, so it would be necessary to continually update the information in the server if the approach is to block each service separately. Also many applications assign ports dynamically, so they cannot be put in a list. Therefore the configuration should be made to allow access to networks, hosts and ports that are approved and to block everything else.

Examiner's Guidance Notes

There was a tendency with quite a lot of students to amalgamate all possible aspects of security with the role of a firewall – even to include, in some cases, user authentication with a password and user name combination. So, in many cases, the central focus was not entirely clear and this problem came out even more in some of the subsequent parts of the question. Many students could see that the problem with having several firewalls configured differently was that the system was only as strong as the weakest firewall. c) required an explanation of the principle of filtering packets. Those students who only had a vague notion of the precise function of a firewall could not give a precise answer to this question. The key is in looking at the origin and destination of packets (including the destination ports). Many students thought that the actual data in the packets was scrutinised 'for errors'. Only a few students understood what was meant by an approach of granting access rather than restricting access. What it means is that everything is denied that is not explicitly allowed. Unless the students grasped this then the answer was unlikely to be convincing.