PAPER CODE NO.	EXAMINER : Martin Gairing
COMP211	DEPARTMENT : Computer Science Tel. No. 0151 795 4264

UNIVERSITY OF
LIVERPOOL

First Semester Examinations 2017/18 (Model Solution)

INTERNET PRINCIPLES

TIME ALLOWED : Two Hours

INSTRUCTIONS TO CANDIDATES

This examination consists of two sections. Section A is worth 25 marks and Section B is worth 75 marks. Answer **ALL** questions is Section A and **THREE** questions from Section B. If you attempt to answer more questions than the required number of questions (in any section), the marks awarded for the excess questions answered will be discarded (starting with your lowest mark).

THIS PAPER MUST NOT BE REMOVED FROM THE EXAMINATION ROOM

UNIVERSITY OF LIVERPOOL

Section A

Each of the following questions comprises several statements, for which you should select ALL answers that apply. (2 marks each)

(10 MCQ questions)

The following TRUE/FALSE questions are worth 1 mark each.

(5 True/False questions)

UNIVERSITY OF LIVERPOOL

Section B

Reminder: You have to answer **THREE** questions from Section B.

1. **QUESTION ONE**

A. If you intercepted a message that was encrypted using a Monoalphabetic cipher, but you did not know the key, what type of attack could you use to decrypt the message? Outline how this attack would be performed. **3 marks**

Model Solution:
A brute-force attack could be used to decrypt the intercepted message. For a brute-force attack, one would try each key on the message until the message is understandable.

For a brute-force attack, one would first try to decrypt short words in the message. Once a short word was decrypted to a recognisable word, one would fix the key for the letters of the recognised word. Then one would continue decrypting other words in the ciphertext, until one decrypts the entire ciphertext, or until subsequent words yield gibberish, at which point one would revise the key.

One should keep in mind that the key space is large (26! possible keys) and thus a brute-force attack can be a very time consuming process.

Considering that we do not know any of the plain-text in the message, the only other option is to use statistical methods to estimate the location of common letters. For instance, one would determine the frequency of the letters and letter combinations in the ciphertext, and then compare that with such frequencies of a common text in English. Letters with matching frequencies are likely to be mapped to each other by the cipher.

B. In the lectures we studied an efficient and secure e-mail scheme, which provides secrecy, sender authentication and message integrity.

i. Draw a diagram for the sender side of this scheme. **4 marks**

Model Solution:

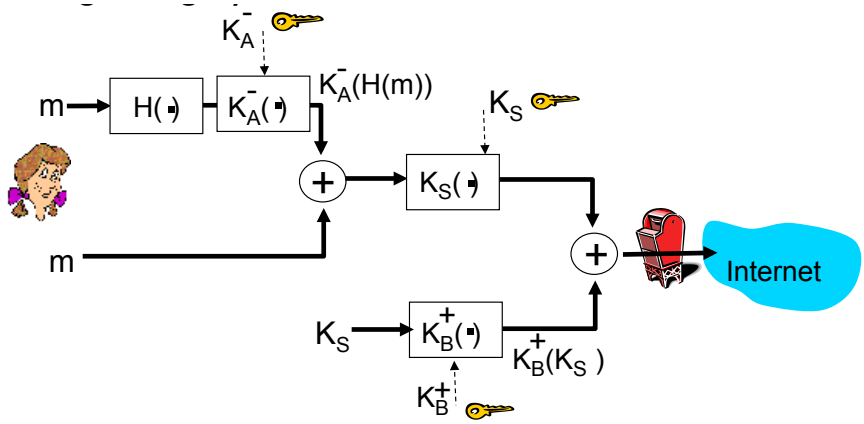

ii. Explain how the scheme ensures secrecy, sender authentication and message integrity and why it is efficient. **6 marks**

Model Solution:
Alice digitally signs a message digest $H(m)$ of the email using her private key K_A^-, verifying that she is the sender of the email, providing sender authentication and message integrity. She then generates a random symmetric key K_S and encrypts message and digest with K_S. Alice also encrypts the session key K_S with Bob's public key K_B^+. Encrypted key and message are send to Bob and only Bob can decrypt the session key $K_B^+(K_S)$ with his private key and then use K_S to retrieve the message m. This ensures secrecy.
The scheme is efficient since the computational intensive public key cryptography is only used for encoding/decoding the fixed size message digest and the session key K_S but not the whole email.

C. Alice wants to communicate with Bob using symmetric-key cryptography (e.g. DES) with a session-key K_S.

i. In the lectures we learned how public-key cryptography (e.g. RSA) can be used to distribute a session key K_S from Alice to Bob. Suppose the private keys of Alice and Bob are K_A^- and K_B^-, while the public keys are K_A^+ and K_B^+. Draw a diagram that shows the message exchange between Alice and Bob which achieves this. **4 marks**

Model Solution:

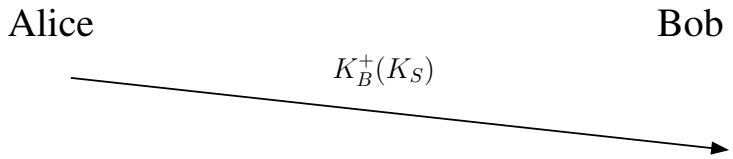

ii. Now we want to distribute the session key without public-key cryptography using a Key Distribution Centre (KDC). The KDC is a server that shares a secret symmetric key with each registered user. For Alice and Bob denote those keys by K_{A-K} and K_{B-K}. Upon request the KDC issues a session key K_S. Design a scheme that uses the KDC to distribute K_S to Alice and Bob. Your scheme should use three messages:

- a message from Alice to the KDC requesting a session key K_S with Bob,
- a message from the KDC to Alice,
- a message from Alice to Bob.

Draw a diagram showing the exchanged messages. **5 marks**

Model Solution:

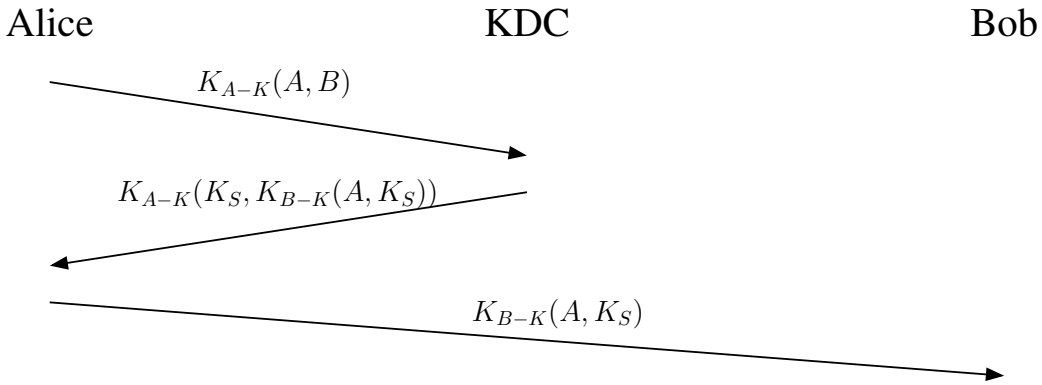

iii. What is the main advantage of first distributing a session key and then using symmetric-key cryptography rather than using public-key cryptography techniques for the whole communication? **3 marks**

Model Solution:
Symmetric-key cryptography is more efficient than public-key cryptography. In particular DES is at least a factor 100 faster that RSA.

2. **QUESTION TWO**

A. Alice sends a message to Bob which is 4500 bytes long, and is broken into segments of 1200 bytes each. Alice chooses a random start value of 2400 for her sequence numbers.

i. How many segments will the message be broken into? **1 mark**

ii. Give the start and end bytes of each segment. **2 marks**

iii. Give the ACK numbers which Bob will use to indicate that each segment was received uncorrupted. **2 marks**

iv. Suppose Bob chooses a random start of 852 for the sequence numbers (of his ACKs), and that he only sends headers (and no data) back to Alice. What will be the ACK numbers used by Alice in response to these ACKs? **2 marks**

v. Draw a brief Message Sequence Chart for the interaction. **3 marks**

Model Solution:

(i) 4 segments

(ii)
2400, 3599
3600, 4799
4800, 5999
6000, 6899

(iii)
3600, 4800, 6000, 6900

(iv)
all of them will have ACK number 852

(v) The MSC should show two vertical lines, one representing Alice and one Bob, with time running down the page. Between the two lines are diagonal arrows representing each message sent by either party, in sequence order, with each arrowhead pointed towards the receiver of the message. Each arrow should be annotated with the sequence numbers or ACK numbers corresponding to that message which the arrow represents.

B. Consider the Go-Back-N protocol with a sender window size of $N = 4$ and a sequence number range of 1024. Suppose that at time t the next in-order packet that the receiver is expecting has sequence number 650. Assume that the medium does not reorder messages. What are the possible sets of sequence numbers inside the sender's window at time t? Justify your answer. **4 marks**

Model Solution:
By the fact that the receiver is expecting a packet with sequence number 650, we know that the receiver has received packet 649, and has ACKed that and all other preceding packets. If all of these ACK's have been received by sender, then sender's window is [650, 653]. Suppose next that none of the ACKs have been received at the sender. In this second case, the sender's window contains 646 and the 4 packets up to and including 649. The sender's window is thus [646,649]. By these arguments, the senders window is of size

4 and begins somewhere in the range [646,650].

C. Consider a router that interconnects three subnets: A, B, and C. Suppose all of the interfaces in each of these subnets are required to have the prefix 71.5.88.0/22. Suppose subnet A is required to support 500 interfaces, and subnets B and C are each required to support 250 interfaces. Provide network addresses for A,B and C (in the form a.b.c.d/x) that satisfy these constraints. **6 marks**

Model Solution:
E.g.:
A: 71.5.90.0/23
B: 71.5.88.0/24
C: 71.5.89.0/24

D. What is the problem of *network congestion*? Which nodes in the Internet first know about this problem? Do they communicate their knowledge to other nodes when using TCP? If so, how? How does the sender of a message using this protocol respond to such communications? **5 marks**

Model Solution:
Network congestion occurs when there is too much traffic for the communication channels and/or for the devices in the network core to process. The consequences are delays in packet delivery and loss of packets.

The only nodes with direct knowledge of the problem are the nodes in the network core; however, these nodes may not have the capability to detect the problem. Under TCP, this knowledge is not communicated directly to the nodes at the network edge.

Under TCP, the Sender (Alice) may infer network congestion on the basis of timing for successful delivery of packets and the extent to which resending of packets is required. When thus inferred, Alice may throttle back (or up) her rate of packet sending.

3. **QUESTION THREE**

A. Suppose 6 host machines and 1 router are connected by a company network consisting of 3 subnets. The configuration is given in the following table:

Subnets	Host IP-addresses	Router IP-Addresses
66.25.48.0/22	66.25.48.1	66.25.48.44
66.25.52.0/23	66.25.52.1	66.25.52.22
66.25.56.0/23	66.25.52.2	66.25.56.11
	66.25.53.1	
	66.25.56.1	
	66.25.56.3	

i. Draw a diagram to represent this configuration. **4 marks**

ii. Draw the forwarding table for the host machine with IP address 66.25.52.2 **3 marks**

iii. Draw the forwarding table for the router. **3 marks**

iv. Suppose an additional host machine is connected to the company network. For each of the following IP addresses, either give the subnet to which this IP address belongs, or state that it is not a valid IP address for any of the subnets.

- 66.25.50.1 **2 marks**
- 66.25.58.1 **2 marks**

Model Solution:

(i) Diagram

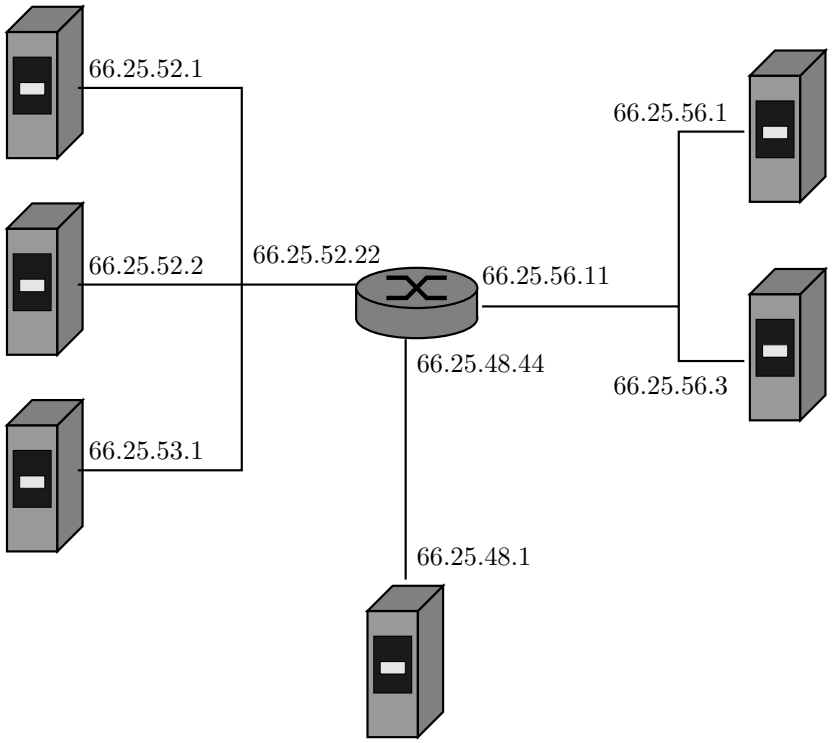

(ii)

Destination Subnet	Next Router	Number of Hops
66.25.48.0/22	66.25.52.22	2
66.25.52.0/23	—	1
66.25.56.0/23	66.25.52.22	2

(iii)

Destination Subnet	Next Router	Number of Hops	Interface
66.25.48.0/22	—	1	66.25.48.44
66.25.52.0/23	—	1	66.25.52.22
66.25.56.0/23	—	1	66.25.56.11

(iv) 66.25.50.1 will be connected to 66.25.48.0/22.
66.25.58.1 is not in the IP address range of any of the 3 subnets.

B. Suppose two network interface cards (NIC) A and B are attached to the same broadcast LAN and that A wants to send a datagram to B but doesn't know B's MAC address. Which protocol can be use to resolve this? Explain step by step how A retrieves B's MAC address using this protocol. **5 marks**

Model Solution:
A can use the ARP protocol to get the MAC address of B:
A broadcasts ARP query packet containing B's IP-address to destination MAC address FF-FF-FF-FF-FF-FF. All machines on LAN (including B) receive ARP query. B replies to A with its (B's) MAC address in a frame sent directly to A's MAC address (unicast). A caches B's MAC address in its ARP table.

C. Suppose Bob joins a BitTorrent torrent, but does not want to upload any data to any other peers (so called free-riding).

i. Bob claims that he can receive a complete copy of the file that is shared by the swarm. Is Bob's claim possible? Why or why not? **3 marks**

ii. Bob further claims that he can further make the free-riding more efficient by using a collection of multiple computers (with distinct IP addresses). How can he do that? **3 marks**

Model Solution:

(i) Yes. His first claim is possible, as long as there are enough peers staying in the swarm for a long enough time. Bob can always receive data through optimistic unchoking by other peers.

(ii) His second claim is also true. He can run a client on each host, let each client free-ride, and combine the collected chunks from the different hosts into a single file. He can even write a small scheduling program to make the different hosts ask for different chunks of the file. (This is actually a kind of Sybil attack in P2P networks.)

4. **QUESTION FOUR**

A. Consider a communication channel with bandwidth $B = 6000\,\text{Hz}$.

i. Suppose the channel has a signal-to-noise ratio $S/N = 511$. What is the *maximum data rate* of this channel? **3 marks**

ii. What is the minimum number of signal states M needed to achieve a data rate of $48000\,\text{bps}$? How many bits must each signal state encode? **3 marks**

Model Solution:

(i) Shannon:

$$\text{max data rate} = B \cdot \log_2(1 + S/N) = 6000\text{Hz} \cdot 9 = 54000\text{bps}$$

(ii) Nyquist:

$$\text{max data rate} = 48000\text{Hz} = 2 \cdot B \cdot \log_2(M) = 12000\text{Hz} \cdot \log_2(M).$$

Thus $M = 16$ and each state must encode $\log_2(M) = 4$ bits.

B. What is a CRC code? What purpose does it serve? Compute the CRC bits defined by the generator 1001 and the data bit string 101110. **6 marks**

Model Solution:

CRC stands for Cyclic Redundancy Check and is a checksum function. A checksum function is a means to assess whether data has been corrupted in transit.

Given the data bitstring $D = 101111$ and the generator $G = 1001$ we have to compute the CRC bits R such that $< D, R >$ is divisible (modulo-2) by G:

$$\begin{array}{ll} \overbrace{1001}^{G} & \overbrace{101110}^{D}000 \\ & \underline{1001} \\ & 0101 \\ & \underline{0000} \\ & 1010 \\ & \underline{1001} \\ & 0110 \\ & \underline{0000} \\ & 1100 \\ & \underline{1001} \\ & 1010 \\ & \underline{1001} \\ & \underbrace{011}_{R} \end{array}$$

The CRC bits are 011.

C. For a communication session between a pair of processes, which process is the client and which is the server? **2 marks**

Model Solution:
Client: process that initiates communication.
Server: process that waits to be contacted.

D. What is the difference between *persistent* and *non-persistent* HTTP? Which version of HTTP supports both? **3 marks**

Model Solution:
Non-persistent HTTP: Only one object is sent over a TCP connection. HTTP/1.0 uses non-persistent HTTP.
Persistent HTTP: Multiple object can be sent over single TCP connection between client and server. HTTP/1.1 uses persistent HTTP in default mode but can also use non-persistent HTTP.

E. Explain the difference between TDMA, FDMA, CSMA/CD and Slotted ALOHA. **5 marks**

Model Solution:

TDMA: channel partitioning protocol; allocation of channel-use is by time-slots.
FDMA: channel partitioning protocol; allocation of channel-use is by frequencies.
CSMA/CD: random access protocol; if channel sensed busy defer transmission; if collision is detected abort transmission and wait random time.
Slotted ALOHA: random access protocol; time divided into slots, when node obtains fresh frame it transmits in next slot; if collision send frame in each subsequent slot with probability p.

F. Why are there different protocols for Inter-AS and Intra-AS routing? **3 marks**

Model Solution:
Intra-AS routing creates routes for packets inside the same autonomous system, while Inter-AS routing is used to create routes between different autonomous systems. In Inter-AS routing the AS admin wants control about how its traffic is routed and who is routing through its network. This policy may dominate over performance. Inside an AS the focus is on performance, since no policy decisions are needed (single admin).
Hierarchical routing also saves table size and reduces update traffic.