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

UNIVERSITY OF
LIVERPOOL

First Semester Examinations 2016/17 (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 question 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

1. **QUESTION ONE**

A. Alice sends a message to Bob which is 3800 bytes long, and is broken into segments of 900 bytes each. Alice chooses a random start value of 1500 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 259 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) 5 segments

(ii)
1500, 2399
2400, 3299
3300, 4199
4200, 5099
5100, 5299

(iii)
2400, 3300, 4200, 5100, 5300

(iv)
all of them will have ACK number 259

(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 900. Assume that the medium does not reorder messages. What are all possible values of the ACK field in all possible messages currently propagating back to the sender at time t? Justify your answer. **4 marks**

Model Solution:
If the receiver is waiting for packet 900, then it has received (and ACKed) packet 899 and the 3 packets before that. If none of those 4 ACKs have been yet received by the sender, then ACK messages with values

of [896,899] may still be propagating back. Because the sender has sent packets [896, 899], it must be the case that the sender has already received an ACK for 895. Once the receiver has sent an ACK for 895 it will never send an ACK that is less than 895. Thus the range of in-flight ACK values can range from 895 to 899.

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 120.3.80.0/22. 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: 120.3.82.0/23
B: 120.3.80.0/24
C: 120.3.81.0/24

D. Consider an HTTP client that wants to retrieve a Web document at a given URL. The IP address of the HTTP server is initially unknown.

i. Which application layer protocols are needed in this scenario and what are they used for? **3 marks**

ii. Which transport layer protocols do these protocols use? **2 marks**

Model Solution:

(i) DNS to get the IP-address and HTTP for receiving the document.

(ii) DNS uses UDP in default mode, while HTTP uses TCP.

2. **QUESTION TWO**

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
98.13.176.0/22	98.13.176.1	98.13.176.44
98.13.180.0/23	98.13.180.1	98.13.180.22
98.13.184.0/23	98.13.180.2	98.13.184.11
	98.13.184.1	
	98.13.184.2	
	98.13.184.3	

i. Draw a diagram to represent this configuration. **5 marks**

ii. Draw the forwarding table for the host machine with IP address 98.13.180.2 **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.

- 98.13.182.1 **2 marks**
- 98.13.186.1 **2 marks**

Model Solution:

(i) Diagram

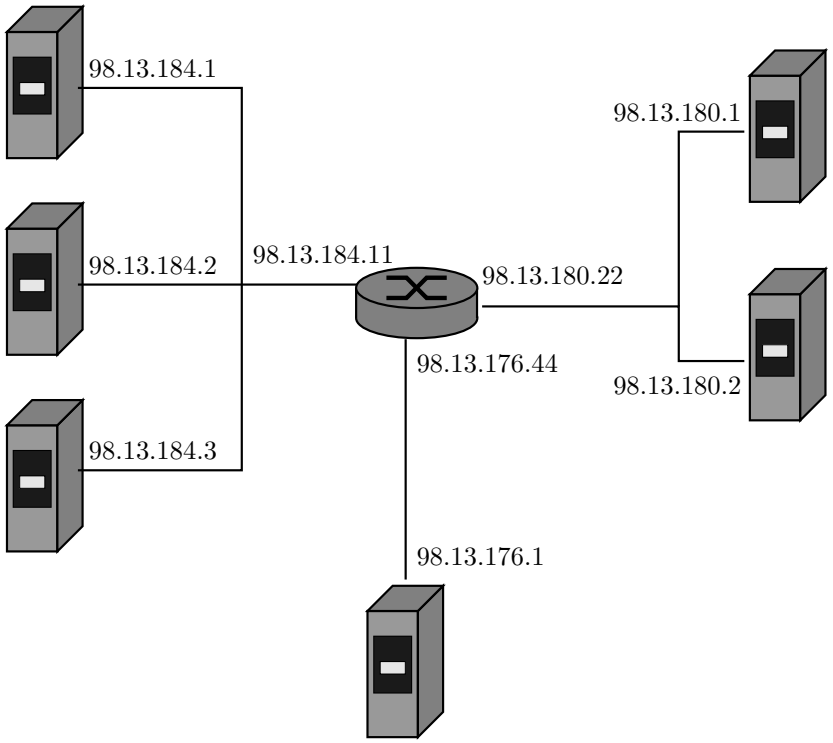

(ii)

Destination Subnet	Next Router	Number of Hops
98.13.176.0/22	98.13.180.22	2
98.13.180.0/23	—	1
98.13.184.0/23	98.13.180.22	2

(iii)

Destination Subnet	Next Router	Number of Hops	Interface
98.13.176.0/22	—	1	98.13.176.44
98.13.180.0/23	—	1	98.13.180.22
98.13.184.0/23	—	1	98.13.184.11

(iv) 98.13.182.1 is not in the IP address range of any of the 3 subnets.
98.13.186.1 is not in the IP address range of any of the 3 subnets.

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 101111. **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:

$$\begin{array}{ll} \overbrace{1001}^{G} & \overbrace{101111}^{D}000 \\ & \underline{1001} \\ & 0101 \\ & \underline{0000} \\ & 1011 \\ & \underline{1001} \\ & 0100 \\ & \underline{0000} \\ & 1000 \\ & \underline{1001} \\ & 0010 \\ & \underline{0000} \\ & \underbrace{010}_{R} \end{array}$$

The CRC bits are 010.

C. When you design a new application-layer protocol you have to define 4 items. What are they? **4 marks**

Model Solution:

- type of messages exchanged (e.g. request, response)
- syntax of message (fields in the message, how delimited)
- semantics of the message (meaning of the information in fields)
- rules for determining when and how a process sends messages and responds to messages

3. **QUESTION THREE**

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

i. Suppose the channel has a signal-to-noise ratio $S/N = 1023$. 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 $36000\,\text{bps}$? How many bits must each signal state encode? **3 marks**

Model Solution:

(i) Shannon:

$$\text{max data rate} = B \cdot \log_2(1 + S/N) = 6000\text{Hz} \cdot 10 = 60000\text{bps}$$

(ii) Nyquist:

$$\text{max data rate} = 36000\text{Hz} = 2 \cdot B \cdot \log_2(M) = 12000\text{Hz} \cdot \log_2(M).$$

Thus $M = 8$ and each state must encode $\log_2(M) = 3$ bits.

B. 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.)

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. Suppose 50 hosts are sharing a broadcast channel. Further suppose at any time each host has a frame to send with probability p. Which of the multiple access protocols from (3E) are desirable if p is low (say 1%)? Why? What about if p is high (say 90%)? **3 marks**

Model Solution:
If p is low, there is little chance to experience collisions. Thus, a random access protocol (CSMA/CD or Slotted ALOHA) is desirable. If p is high, a channel partitioning protocol (TDMA or FDMA) is more efficient.

4. **QUESTION FOUR**

A. Consider RSA with $p = 7$ and $q = 13$.

i. What are n and z? Show all work. **3 marks**

ii. Let e be 5. Why is this an acceptable choice for e? **2 marks**

iii. Find d such that $(e \cdot d \mod z) = 1$. **2 marks**

iv. Encrypt the message $m = 3$ using the key (n, e). Let c denote the corresponding ciphertext. Show all work. **3 marks**

Model Solution:

i. $n = p \cdot q = 91$
$z = (p-1) \cdot (q-1) = 72$

ii. $e = 5$ is less than n and has no common factors with z.

iii. $d = 29$

iv. $m = 3$, so
$c = (m^e \mod n) = (3^5 \mod 91) = (243 \mod 91) = 61.$

B. The *Monoalphabetic cipher*, as discussed in the lectures, is an encryption method that, for every letter in the plaintext, substitutes that letter with a letter determined by a one-to-one mapping M, where M is the key. Let M_1 and M_2 be two such mappings that are given as follows:

M_1 :

$a \mapsto z$	$h \mapsto y$	$o \mapsto x$	$v \mapsto w$
$b \mapsto v$	$i \mapsto u$	$p \mapsto t$	$w \mapsto s$
$c \mapsto r$	$j \mapsto q$	$q \mapsto p$	$x \mapsto o$
$d \mapsto n$	$k \mapsto m$	$r \mapsto l$	$y \mapsto k$
$e \mapsto j$	$l \mapsto i$	$s \mapsto h$	$z \mapsto g$
$f \mapsto e$	$m \mapsto f$	$t \mapsto d$	
$g \mapsto c$	$n \mapsto b$	$u \mapsto a$	

M_2 :

$a \mapsto y$	$h \mapsto x$	$o \mapsto z$	$v \mapsto t$
$b \mapsto w$	$i \mapsto o$	$p \mapsto r$	$w \mapsto s$
$c \mapsto v$	$j \mapsto n$	$q \mapsto q$	$x \mapsto i$
$d \mapsto u$	$k \mapsto m$	$r \mapsto p$	$y \mapsto h$
$e \mapsto l$	$l \mapsto c$	$s \mapsto f$	$z \mapsto g$
$f \mapsto k$	$m \mapsto b$	$t \mapsto e$	
$g \mapsto j$	$n \mapsto a$	$u \mapsto d$	

i. How many different keys does the Monoalphabetic cipher have? **1 mark**

ii. Is the Monoalphabetic cipher a symmetric-key or a public-key cryptography technique? Explain. **2 marks**

iii. Using the Monoalphabetic cipher with the key M_1, **decrypt** the following ciphertext and write the decoded plaintext: `tlxralzdxl` **3 marks**

iv. The *Polyalphabetic cipher* improves upon the Monoalphabetic cipher by applying several Monoalphabetic ciphers according to a cyclic pattern. Let the following be such a cyclic pattern using M_1 and M_2:

$$M_1, M_2, M_1; M_1, M_2, M_1; \ldots$$

Using the Polyalphabetic cipher according to this pattern, **encrypt** the following plaintext and write the encoded ciphertext: `ergo bibamus` **3 marks**

Model Solution:

(i) 26!

(ii) The Monoalphabetic cipher is a symmetric-key algorithm. The same key is used for encryption and decryption.

(iii) Plaintext: `procurator`

(iv) Ciphertext: `jpcx wuvyfaf`

C. 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.

D. 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.