Networked Systems, COMPGZ01, 2011
Answer TWO questions from Part ONE on the answer booklet containing lined writing paper, and answer ALL questions in Part TWO on the multiple-choice question answer sheet.
Marks for each part of each question are indicated in square brackets
Calculators are permitted

Part ONE

1. **Fast IP Checksums**

   After graduation, you and some other students join a London-based computer networking startup to design and build new Internet protocols for the data center. Your aim is speed, so you and your fellow engineers take a close look at Internet checksum calculation.

   To help get you started, we reproduce part of the 16-bit 1's complement number system:

16-bit 1's complement	Decimal
1000 0000 0000 0000	−32767
1000 0000 0000 0001	−32766
1000 0000 0000 0010	−32765
⋮	⋮
1111 1111 1111 1011	−4
1111 1111 1111 1100	−3
1111 1111 1111 1101	−2
1111 1111 1111 1110	−1
1111 1111 1111 1111	−0

16-bit 1's complement	Decimal
0000 0000 0000 0000	+0
0000 0000 0000 0001	+1
0000 0000 0000 0010	+2
0000 0000 0000 0011	+3
0000 0000 0000 0100	+4
⋮	⋮
0111 1111 1111 1101	+32765
0111 1111 1111 1110	+32766
0111 1111 1111 1111	+32767

   Carefully reading the UDP RFC,[1] your colleague Mo Bitz discovers that a UDP transmitter can specify an all zero (`0x0000`) "flag" value in the Internet checksum field of the UDP header to indicate it did not generate a checksum. Receivers can then save the effort of verifying the checksum.

   The RFC also states that receivers should replace a checksum of `0x0000` (that is to say, the 1's complement sum of the words being checksummed is `0xFFFF`) with `0xFFFF`.

   a. Give a precise yet terse description, in words, of how the sender computes an Internet checksum over data d.

   [3 marks]

   b. Give a precise yet terse description, in words, of how the receiver verifies an Internet checksum over received data d' with computed checksum c appended.

   [3 marks]

[1] RFC 768 for the curious.

c. Why does replacing an Internet checksum of `0x0000` with `0xFFFF` not alter the result of the receiver's checksum verification?

[3 marks]

Thinking about the RFC's substitution rule makes you and Mo wonder how often real packet data sums to `0x0000`.

d. Show that whenever the Internet checksum is `0xFFFF` (that is to say, the 1's complement sum of the words being checksummed is `0x0000`), then all the 16-bit words the Internet checksum is taken over are necessarily `0x0000`.

Hint: One approach is to show that this is true for an Internet checksum made up of one-bit 1's complement numbers, and then assuming it is true for $n-1$ bit 1's complement numbers, show it is true for n bit 1's complement numbers.

[10 marks]

e. Why would a UDP checksum not be computed over all-zero data in practice on the Internet?

[4 marks]

Now you and Mo begin to think about how a router can quickly recompute the IP checksum when it decrements the *one byte* time-to-live field in the IP header in the forwarding operation.

f. Give an algorithm to decrement one byte in a buffer covered by an Internet checksum without inspecting any other byte in the buffer other than the one byte comprising the Internet checksum.

Your algorithm should be slightly different depending on whether the byte in question is a high-order byte or a low-order byte in its word: use the following template for your algorithm:

```
if (high order byte) {

  /* Your expression for a high order byte goes here. */

}
else {

  /* Your expression for a low order byte goes here. */

}
```

Use the variable C for the old value stored in the checksum field, and C' for the new value stored in the checksum field.

[10 marks]

[Question 1: Total 33 marks]

2. **Spanning Tree Protocol**

Having just learned about the Spanning Tree Protocol (STP), UCL student Bill DiBosco decides to build a bridged LAN running STP in his student housing. Keeping in mind that more paths offer more redundancy to link failures, he constructs the following topology:

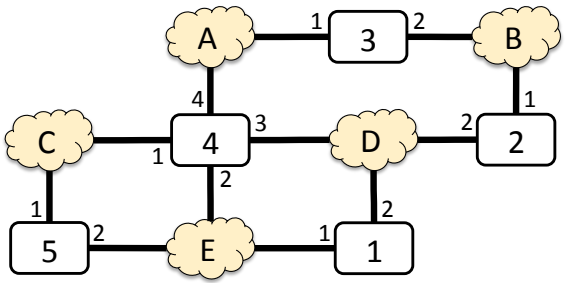

Here clouds labeled with letters represent LANs, and each bridge (also known as a switch) is numbered with its identifier. The numbers just outside each switch in the figure above are port identifiers.

Bill wires up the above topology, and flips a switch connected to the power supply of all bridges so that they power on simultaneously. Each bridge sends out one configuration message of the form (Root-ID, Distance, Switch-ID) as soon as it powers up.

a. Just after every switch has received the first configuration message sent by every other switch, there are five ports in the designated state on the bridges in Bill's network. Which ports are they?

[1 mark for each port, total 5 marks]

b. Does your answer to the previous part depend on the order in which switches send (and receive) configuration messages? Why or why not?

[3 marks]

c. After the STP has converged, Bill examines the state of each switch in the network. **In your answer book**, replicate the following, filling in the blanks and the blank entries in the tables corresponding to the state Bill sees in the following switches. You may abbreviate root ports as "R", blocked ports as "B", and designated ports as "D". Some entries are filled out for you.

i. **Switch 1**
Root identifier: _____

[1 marks]

Configuration message: (1, 0, 1)

Port	**Best configuration message**	**Port state**
1	(1, 0, 1)	D
2	_____	_____

[1 marks]

ii. **Switch 2**

Root identifier: _____

[1 marks]

Own configuration message: _____

[1 marks]

Port	Best configuration message	Port state
1	_____	_____
2	_____	_____

[Each line 1 mark; total 2 marks]

iii. **Switch 3**

Root identifier: _____

[1 marks]

Own configuration message: _____

[1 marks]

Port	Best configuration message	Port state
1	_____	_____
2	_____	_____

[Each line 1 mark; total 2 marks]

iv. **Switch 4**

Root identifier: _____

[1 marks]

Own configuration message: _____

[1 marks]

Port	Best configuration message	Port state
1	_____	_____
2	_____	_____
3	_____	_____
4	_____	_____

[Each line 1 mark; total 4 marks]

v. **Switch 5**
Root identifier: _____

[1 marks]

Own configuration message: _____

[1 marks]

Port	Best configuration message	Port state
1	_____	_____
2	_____	_____

[Each line 1 mark; total 2 marks]

For the following questions, suppose that switch 2 fails.

d. Which LANs (if any) are temporarily disconnected from LAN D as a result of switch 2's failure?

[3 marks]

e. Once the STP has re-stabilized, list the LANs and switches a packet sent from LAN B to LAN D will traverse.

[2 marks]

[Question 2: Total 33 marks]

3. **Wireless MACs and TCP**

When a radio transmits using higher power, its range increases. Consider the below wireless topology, in which all nodes use omnidirectional antennas:

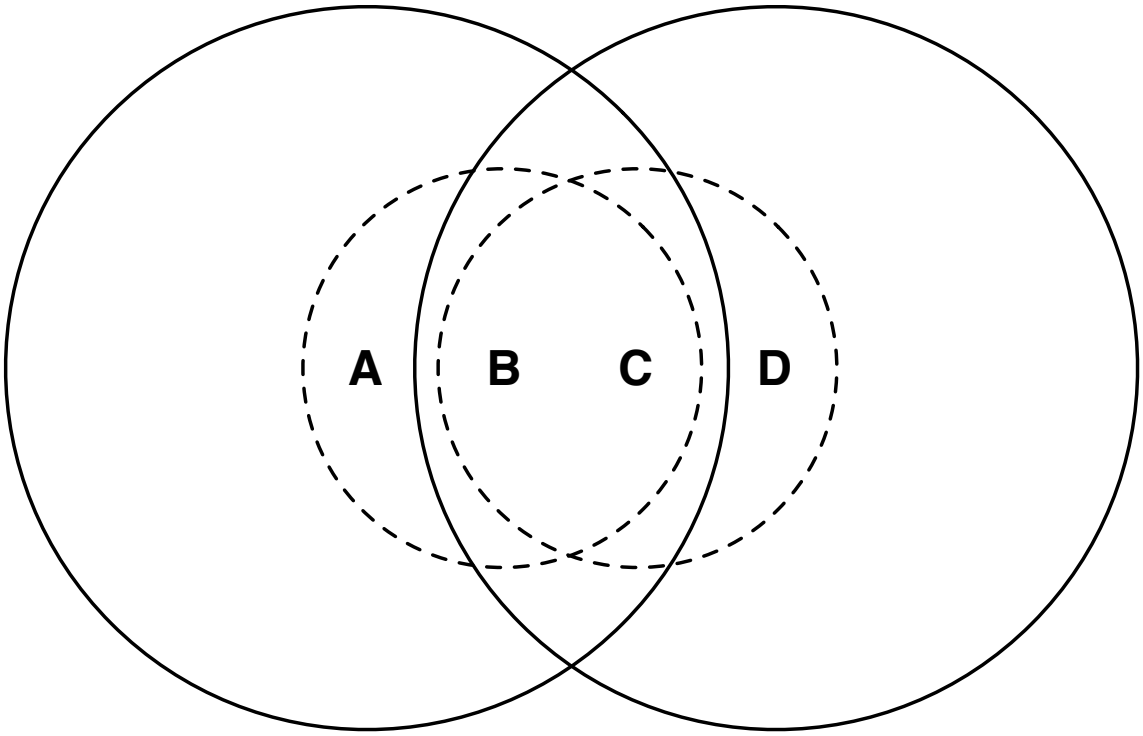

Nodes *A* and *D* in the figure transmit with high power, while nodes *B* and *C* transmit with low power. There are four circles drawn in the figure. Each circle represents the transmit range of the node on which it is centered. The solid circles represent the transmit range of nodes *A* and *D*, while the dashed circles represent the transmit range of nodes *B* and *C*.

Assume that collisions occur when more than one transmission concurrently reaches the receiver, where the transmission range of a sender is determined by the circle centered on it. Assume further that only collisions cause packet losses. Finally, ignore ACKs in all parts of this question.

a. i. Suppose node *A* transmits to node *B*. List all potential hidden terminals with respect to *A in either direction,* (*i.e.,* for cases where *A* sends first *and* cases where the other hidden terminal sends first. Note that a hidden terminal with respect to *A* need not also be sending to *B*—it may be sending to some other node.
In your answer, justify why each node you list is hidden with respect to *A*.

[3 marks]

ii. Again, suppose node *A* transmits to node *B*. List all potential exposed terminals with respect to *A*. Justify why each node you list is exposed with respect to *A*.

[3 marks]

iii. Now suppose node B transmits to node C. As before, list all other nodes hidden with respect to B, in either direction, justifying why each is hidden.

[3 marks]

iv. Again, suppose node B transmits to node C. As before, list all other nodes exposed with respect to B, justifying why each is exposed.

[3 marks]

b. Suppose A sends to B and C sends to D. Both senders send at the same fixed bit-rate. Assume that this bit-rate always decodes successfully when either link is used alone, and never decodes successfully if a collision occurs.

i. Assume that no mechanism is used to detect or avoid collisions. What is the throughput of each transfer as a fraction of its sending rate?

[3 marks]

ii. Now suppose the same two senders send to the same two receivers, but that they use the MACAW MAC protocol discussed in lecture. Assume that any control packets sent by MACAW are short, and thus consume negligible link bandwidth. What is the throughput of each transfer as a fraction of its sending rate?

[4 marks]

The rest of this question concerns TCP.

c. The TCP sender retransmits data segments (packets) for which it does not receive ACKs. It also samples the elapsed round-trip time (RTT) between the sending of a data packet and the return of the ACK for that sequence number, as input to its estimate of the connection's RTT. Suppose the TCP sender retransmits a data packet, and subsequently receives an ACK for that packet. Why should the TCP sender *not* use RTT samples for retransmitted packets in its RTT estimate?

[6 marks]

d. Suppose an attacker sends many packets of the type used in the first step in the usual TCP 3-way handshake to a server, all to the same destination port, but each with a different source port. He sends one such packet every 10 ms, and continues this behavior for hours. He never sends any packets apart from these initial TCP handshake packets (*i.e.,* he does not attempt to pursue any of these handshakes beyond this first step). To hide his own identity, he chooses an unused IP address from elsewhere on the Internet, and sets the source IP address on each of these packets to be this unused IP address.

A typical operating system (OS) implements the server side of TCP connection establishment as follows:

When an application on a server asks the OS to listen for incoming TCP connections to a port, the OS allocates a queue for that TCP port. Each entry in this queue stores state about one incoming connection that has begun, but not yet completed, the full 3-way handshake. This queue is typically fairly short: it's only on the order of ten entries long. While the queue is full, no new incoming connection requests to the port can be processed, because there's nowhere to record state about the progress of the new incoming connection request.

Entries in this queue are freed in one of two ways:

- when a connection completes the 3-way handshake successfully, OR
- when the connection fails to complete the 3-way handshake within a long timeout period (on the order of minutes, to allow any severely delayed messages from the remote host to arrive)

What effect will the attacker's behavior have on legitimate users who wish to connect to the server at the same destination port that is under attack? Refer specifically to the details of the 3-way handshake used by TCP and any relevant aspects of the server's handshake implementation described above in your answer.

Assume that no packets are dropped in the middle of the network. Assume that the total rate at which legitimate users attempt to connect to the server at the destination port is far less than once every 10 ms.

[8 marks]

[Question 3: Total 33 marks]