Networked Systems Midterm Exam, COMP3035/GZ01, 2009-2010

Answer **ALL** questions. Write your answers **on this exam paper**.

Write your name at the top right of the front page of this exam paper.

Write your initials at the top right corner of all other pages of this exam paper.

Marks for each part of each question are indicated in square brackets

Calculators are permitted

1. **Always/Sometimes/Never**

   For each of the following statements, indicate when that statement is true. Each answer should consist only of the single word ”Always,” ”Sometimes,” or ”Never.” Each answer is worth two marks.

   a. According to the end-to-end argument, designers of layered networking protocols should place functionality only in the end systems, and avoid placing functionality in routers in the core of the network.

   [2 marks]

   b. When a TCP sender does not receive an acknowledgement for a data packet, that data packet was lost in the network between sender and receiver.

   [2 marks]

   c. In an isochronously multiplexed network, end-to-end delay for the delivery of a user’s frame may vary because of queueing in the middle of the network, depending on the demand from other users.

   [2 marks]

   d. A single parity bit computed over 128 data bits can detect an error when bit-flips occur in exactly 93 of the 128 data bits, and no bit-flip occurs in the parity bit.

   [2 marks]

e. In Huffman coding, a more frequently occurring symbol in the input is encoded with a shorter-length string of bits than a less frequently occurring symbol in the input.

[2 marks]

f. During the Lulea IP lookup algorithm, looking up an IP address in the forwarding tree data structure requires traversing all three levels of the tree.

[2 marks]

g. Virtual output queueing suffers from head-of-line blocking.

[2 marks]

h. Suppose that packet p arrives on input port i of a switch, and packet q arrives later on the *same* input port i of the same switch. If the switch uses virtual output queueing (as presented in lecture) and schedules its crossbar with the iSLIP algorithm (also as presented in lecture), then packet p will reach its output port before packet q reaches its output port.

[2 marks]

i. A standard NAT (without any special configuration) prevents a host outside the NAT from originating a new TCP connection to a host inside the NAT.

[2 marks]

j. A TCP sender may safely discard data from its send buffer with any sequence number less than the greatest sequence number seen on an ACK received from the TCP receiver. Important: assume in your answer that sequence numbers do not wrap around from $2^{32} - 1$ to zero during the lifetime of a single connection.

[2 marks]

[Total 20 marks]

2. **DNS Lookups**

You are the network administrator for a private 10.0.0.0/8 network. You configure a local nameserver at IP address 10.0.0.1 with the name and IP address of a root nameserver, and a client workstation at IP address 10.0.0.2 with the IP address of the local nameserver 10.0.0.1.

For this question, assume the following DNS nameservers have the following records in their respective databases:

10.0.0.1:

a.root-servers.net.	IN	A	198.41.0.4
.	IN	NS	a.root-servers.net.

a.root-servers.net:

a.gtld-servers.net.	IN	A	192.5.6.30
a0.org.afilias-nst.info.	IN	A	199.19.56.1
net.	IN	NS	a.gtld-servers.net.
org.	IN	NS	a0.org.afilias-nst.info.

a.gtld-servers.net:

ns1.isc-sns.net.	IN	A	72.52.71.1

a0.org.afilias-nst.info:

freebsd.org.	IN	NS	ns1.isc-sns.net.

ns1.isc-sns.net:

www.freebsd.org.	IN	A	69.147.83.33

a. Fill in the following table to reflect the sequence of DNS messages sent over the network when the client workstation at IP address 10.0.0.2 makes a DNS query for domain name www.freebsd.org. For queries, specify the domain name in the "Contents" column. For answers, for each record included in the answer, specify the domain name, type of record ("A" or "NS"), and value of the record in the "Contents" column.

Note: the first row is partially filled in for you, and there are more blank rows in the table than you will need to use in your answer.

Query/ Reply?	From IP:	To IP:	Recursive (Y/N):	Contents:
Query	10.0.0.2			www.freebsd.org.

[20 marks]

b. Now assuming that your local nameserver is a caching DNS nameserver, fill in the sequence of messages if client 10.0.0.3 initiates the same query immediately after the first query, so that no cache entries expire.

Note: the first row is partially filled in for you, and there are more blank rows in the table than you will need to use in your answer.

Query/ Reply?	From IP:	To IP:	Recursive (Y/N):	Contents:
Query	10.0.0.3			www.freebsd.org.

[5 marks]

[Total 25 marks]

3. **Spanning Trees**

Consider the following switched Ethernet topology, where the switches (numbered in boldface) interconnect the LANs shown labeled with letters in the picture. The ports on each switch are numbered as shown in the figure.

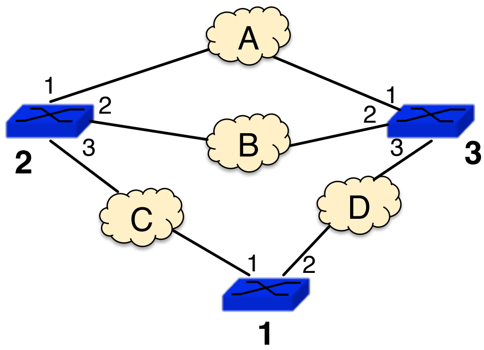

Switches run the spanning tree and learning algorithms as described in lecture. Recall that each switch X periodically sends configuration messages of the form (Root-ID, Distance-to-root, X).

a. Supposing that switch 1 is the first to send its configuration message, write down one possible sequence of events. Include in your list configuration messages sent (and their contents), Root-ID variables set (specifying switch IDs), and ports blocked (specifying switch ID and port IDs).

[10 marks]

Host 10.0.0.1 on LAN A broadcasts an ARP request messsage for host 10.1.0.1 on LAN D. Show the state of each forwarding table after A sends the ARP request.

b. Forwarding table at switch 1:

Ethernet address	Port

[5 marks]

c. Forwarding table at switch 2:

Ethernet address	Port

[5 marks]

d. Forwarding table at switch 3:

Ethernet address	Port

[5 marks]

[Total 25 marks]

4. **Reliable Data Transfer**

Tom Freilink and Sue Ackerley decide to deploy a point-to-point wireless link between their London flats, which are separated by 750 m, with a clear line of sight between them. They build directional antennas from empty cans of Pringles potato chips, and carefully mount their antennas outside their flats' windows, pointing them directly at one another.

Tom and Sue find that their wireless link encounters unpredictable bursts of bit errors, each burst lasting a few consecutive frame durations, where bursts are typically tens of seconds apart. During a burst of bit errors, the receiver typically cannot successfully receive any frames—one or more bits in every received frame is corrupted, such that the CRC-16 checksum over the entire frame (included in the frame's link-layer header) does not match the received data, and the receiver drops the frame.

Assume it takes 1 millisecond to send a fixed-size frame on the link, that all frames sent are of this fixed size, and that the link round-trip time is 3 microseconds.

a. Tom wants to increase the amount of forward error correction (FEC) on all transmitted frames on the link, so that frame transmissions will succeed during these bursts of bit errors. Is FEC a good reliability solution for the particular characteristics of this link? Explain why or why not.

[5 marks]

b. Sue suggests a different approach: implementing stop-and-wait reliability on the link, with a maximum of 5 retries per frame. Assume that Sue tunes the retransmit timer for stop-and-wait to match the link's round-trip time. Is stop-and-wait appropriate for this wireless link's characteristics?

[5 marks]

c. Suppose that Tom and Sue implement stop-and-wait on their link, and send data with TCP on top of this reliable link layer. Moreover, suppose that the version of stop-and-wait they implement uses binary exponential backoff (BEB, as described in lecture for wired Ethernet) between successive retransmission attempts on the same frame.

   i. What will be the effect of this stop-and-wait link layer on the distribution of round-trip times measured by TCP on the link, as compared with the distribution of round-trip times TCP would observe on the same link *without* any stop-and-wait reliability? Justify your answer with relevant details of the behavior of this stop-and-wait scheme.

[10 marks]

ii. Which part of TCP's retransmission algorithm might interact badly with a link-layer retransmission scheme that employs binary exponential backoff (BEB)? (Hint: think about an extreme case, where a long network path consists of many consecutive links, each using this sort of link-layer retransmission scheme.)

[10 marks]

[Total 30 marks]

END OF PAPER