

Networked Systems Midterm Exam, COMP3035/GZ01, 2011-2012

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

Write only your name on this exam paper in the box below, and your initials at the top right corner of all other pages of this exam paper.

Name:

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

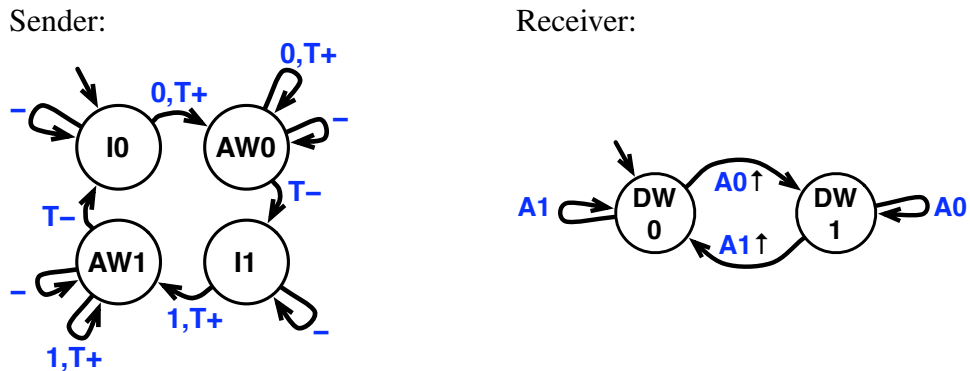
Calculators are permitted

		Earned marks	Possible marks
Part ONE	Question 1		15
	Question 2		20
Part TWO			10
Grand total			50

Part ONE

1. 3G backup link

Alyssa P. Hacker uses the rdt v3.0 protocol described in lecture to transfer important files from her laptop computer to her file server over WiFi and a local Ethernet. The latency between her laptop and file server is about two milliseconds over WiFi, so she runs rdt v3.0 with a 10 millisecond timeout. We have replicated the rdt v3.0 state machine below for your reference.



Key:

-: do nothing

T-: stop timer

0, 1: send data with seq. no. 0, 1

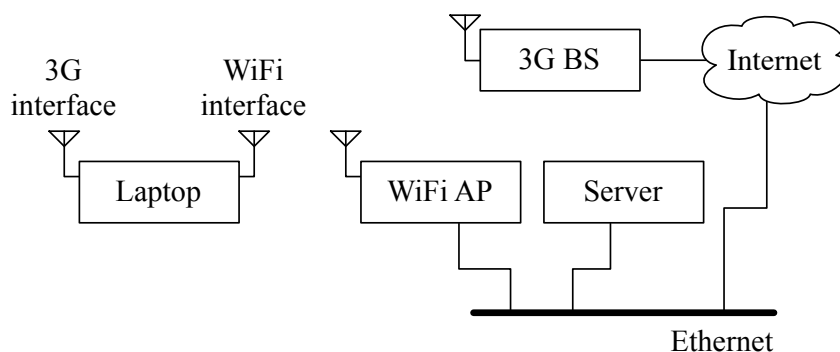
T+: start timer

↑: deliver data to higher layer

A0, A1: send ack with seq. no. 0, 1

One day, fellow 3035/GZ01 student Hayward U. Buzoff moves in next door and sets up a WiFi access point on the same channel as Alyssa's, causing bursts of interference lasting a minute or so, every few minutes, and interrupting Alyssa's file transfers.

Other nearby neighbors use other WiFi frequencies, so switching frequency isn't a good idea for Alyssa. Instead, she adds a 3G (mobile telephone) network interface to her laptop and configures her provider's 3G base station ("3G BS" in the figure below) to route packets onto her Ethernet LAN and to her file server, as shown:




Packets take 100 milliseconds to travel through the 3G base station and over the Internet to or from Alyssa's Server.

Alyssa first configures her rdt v3.0 sender to send the **first, third, fifth, etc. retransmissions on the 3G interface**, while keeping **other retransmissions and all original packet transmissions on the WiFi interface**. She maintains the 10 millisecond timeout on all timers. So packets from Alyssa's laptop to the server now take one of two routes:

1. Laptop to WiFi AP (via WiFi interface), on to Server (via Ethernet).
2. Laptop to 3G BS (via 3G interface), over Internet, and on to Server (via Ethernet).

ACKs from Server to Laptop follow the reverse path of the data they are acknowledging.

- a. When Alyssa begins a file transfer using the above protocol, she observes that files transferred appear to grow in size once received. Complete the timeline below to explain how this might happen, labeling packet and ACK sequence numbers, whether the packet goes over 3G or WiFi, time intervals between each packet and at timer expirations, and labeling sender and receiver states (we've filled in the initial states for you).



[5 marks]

- b. Suggest a change to the way that the rdt v3.0 sender sets retransmission timers that restores exactly-once semantics to this system. You may assume that the identifier of the interface a packet left on is available to the sender. **Explain why** your proposed change fixes Alyssa's protocol.

[5 marks]


- c. Now suppose that the WiFi and 3G bit rates are both 1 Mbit/s, the data packet size is 1000 bits, the frame loss rate on the WiFi link is on average 5%, and the frame loss rate on the 3G link is zero. What is the WiFi link utilization? Ignore processing times at sender and receiver, and show your calculations for full marks. [5 marks]

[Total for Question 1: 15 marks]

2. Careful File Transfer

Kerry Forward has just begun a new job at Cosmic Streams Inc. For his first task, he is asked to implement a careful file transfer protocol between two nodes A and B that are connected together via satellite link. The satellite link suffers from a large propagation delay $t_{prop} = 200$ ms, and a relatively large packet error rate. If sender A wishes to transmit a file, it first breaks the file down into k packets 1KB in size, then transmits each packet at a rate of 1Mbps over the satellite link to B.

- a. Recalling his GZ01/3035 lectures, Kerry devises the following protocol: sender A stores a checksum of the file on its disk. It sends the k file packets to B along with the checksum. B buffers the k packets until the complete file and checksum are received. B writes the file to disk, reads it again, and recomputes the checksum. If its result doesn't match the received checksum, it asks A to retransmit the entire file once again. Since retransmission requests are small, Kerry assumes they consume a negligible amount of transmit time. A timing diagram for Kerry's protocol is shown below.



i. If no retransmissions were necessary, how long would it take for the last data bit of the file to arrive at the receiver under this scheme? Write your answer in terms of k . (Assume the checksum has negligible size). [2 marks]

ii. Suppose the probability of delivering a packet correctly is s . What is the probability of transmitting k packets from A to B correctly? What is the probability that one or more of those k packets is corrupted? [2 marks]

- iii. Let N be the number of times A needs to transmit the entire file. Find the probability that A will need to transmit the entire file $N = n$ times in order for B to receive it correctly. Hint: reason about the cases where A needs to transmit the entire file to B once, twice, and then generalize. [3 marks]

- iv. From 2.a.iii, Kerry deduces that the random variable N has a geometric distribution of the form $P(N = m) = x(1 - x)^{m-1}$. He recalls that the expected value of this variable is $1/x$. Using this knowledge, find the expected time to transmit the entire file in terms of k and s . [2 marks]

- b. To improve the average transmission time required for large files, Kerry's manager, Greta Clue, suggests he implement the checksum functionality at the lower layers of his protocol. Kerry therefore adds code to perform per packet checksums. In particular, Kerry adds a 16 bit checksum to each packet. Using a Stop and Wait protocol, when B receives a packet it checks whether it has been corrupted by recomputing the checksum. If the packet is indeed corrupt, B will ask A to retransmit it. Assume that retransmission requests from B to A are delivered without loss or corruption.
- i. In general, give two possible dangers of implementing a particular functionality at the lower layers of a protocol stack instead of the higher layers. [2 marks]

- ii. Under this scheme, how many times do we need to transmit a single packet on average in order to receive it correctly? Hint: the probability of the number of transmissions is once again geometrically distributed. [3 marks]

iii. Taking into consideration the increase in packet size, what is the average transmission time for the entire file in terms of s and k ?

[3 marks]

c. Greta tells Kerry that after the modifications he's made, the original code that verified the file's overall checksum has become redundant. She asks him to remove it in order to obtain a performance win. Referencing a design principle you've learned in 3035/GZ01, argue whether Greta is correct or not. If she isn't, give an example of a scenario in which Greta's argument would be incorrect.

[3 marks]

[Total for this part: 20 marks]

Part TWO

Multiple Choice

The questions in this section are multiple choice. Zero, one, or more than one choice may be correct, and the number of correct choices varies from question to question. Circle the letters below corresponding to ALL correct choices.

Do not guess. You will be awarded one mark for each choice correctly circled, but you will lose one mark for each choice incorrectly circled. Final results in this section will be normalized across the class.

1. Network Taxonomy

Which of the following statements about classifying networks are true?

- A. Some broadcast networks are circuit-switched.
- B. A datagram network is a kind of packet-switched network.
- C. In some types of packet-switched networks, routers maintain per-connection state.
- D. The Internet is a datagram network.
- E. As originally designed, the telephone network is a datagram network.

2. Layering

Which of the following statements about layered protocol organization are true?

- A. Each layer on a host interacts with its peer host's corresponding layer via the service interface.
- B. In the Internet, bits sent from one host to another pass through the transport layer of routers on each hop.
- C. NAT is an example of a protocol that violates layering on the Internet.
- D. The end-to-end argument asserts that systems should not be structured in a layered manner.
- E. The network layer is the lowest layer that knows how to deliver packets between different types of networks on the Internet.

3. ALOHAnet

Which of the following are true statements about the slotted ALOHA MAC protocol with N nodes (for large N)?

- A. When all nodes constantly have data to send and holding send probability constant, the aggregate throughput of a slotted-ALOHA network decreases with increasing numbers of nodes.
- B. Sending at a rate less than $1/N$ at each node results in throughput less than optimal because of an excessive number of collisions.
- C. Sending at a rate greater than $4/N$ at each node results in throughput near optimal, despite increased collisions.
- D. Slotted ALOHA is an example of an FDMA MAC protocol.
- E. Slotted ALOHA is an example of a TDMA MAC protocol.

4. IP Forwarding

Which of the following are true statements about IP forwarding on the Internet?

- A. IP fragment reassembly can happen at IP routers on the Internet.
- B. Once path MTU discovery completes between two endpoints on the Internet, datagram fragmentation won't occur.
- C. For the most part, fast Internet routers avoid costly IP checksum calculations in the forwarding path.
- D. The IP checksum is implemented by means of a 16-bit cyclic redundancy check.
- E. Routers use the IP checksum to detect routing loops and drop packets stuck in a routing loop, to avoid wasting link capacity.

5. Wireless Link-Sharing Protocols

Which of the following are true statements about the specified wireless link-sharing protocols?

- A. When used to share the same total spectrum bandwidth, Frequency Division Multiple Access (FDMA) suffers more from fading than Carrier Sense Multiple Access (CSMA).
- B. The 802.11 MAC protocol achieves greater efficiency than MACAW by using Collision Detection.
- C. MACAW's Request-to-Send/Clear-to-Send (RTS/CTS) exchange prevents all collisions by hidden terminals.
- D. MACA (the predecessor to MACAW) did not suffer from the exposed terminal problem.
- E. In MACAW, when a sender does not receive an ACK in response to a DATA packet, the sender retransmits the DATA packet again without repeating the RTS/CTS handshake.

6. Sharing in Datagram Networks

Which of the following are true statements about the sharing of datagram networks?

- A. When the queue for a shared link becomes full, an Internet router sends explicit feedback packets to traffic sources demanding that they reduce the rate at which they are sending.
- B. In queueing models used to understand the behavior of shared packet-switched networks, increasing the utilization of a shared link typically incurs a longer average-case queueing delay for that link.
- C. When an Internet router reboots, all end systems whose connections traversed that router before the reboot must send new connection setup messages before the rebooted router will know how to forward packets along these connections once again.
- D. Internet routers typically include enough buffering at their queues to accommodate the worst-case-sized burst of traffic expected to arrive for the same output link.
- E. The total delay required to forward a packet on an asynchronously multiplexed link is independent of the packet's length.

7. Multi-Protocol Label Switching (MPLS)

Which of the following are true statements about MPLS?

- A. Explicit routing in MPLS requires that multiple labels be included in a single packet.
- B. IPv6 eliminates the need for MPLS.
- C. Higher-level protocols switched by MPLS must be modified to be MPLS-aware.
- D. An MPLS label need not be globally unique across all routers or switches; it only has significance at a single switch or router.
- E. Unlike IP forwarding, MPLS forwarding is exact-match.

8. The Domain Name System (DNS)

Which of the following are true statements about DNS?

- A. The root DNS servers store the complete domain name-to-address mapping (*i.e.*, all A records) for all hosts on the Internet.
- B. If domain `foo.com` operates its own nameservers, `ns1.foo.com` and `ns2.foo.com`, for resolution of names in `foo.com` to work, the TLD nameserver for `.com` *must* return glue records for `ns1.foo.com` and `ns2.foo.com` to queriers who look up names in `foo.com`.
- C. Suppose several servers with different IP addresses but the same DNS name all serve the same content, to increase total capacity for responding to requests. A DNS server could support load balancing by choosing which address to return for an A lookup on the basis of current load on the several potential candidate servers.
- D. When made while a user browses to an SSL-enabled URL, a successful DNS cache poisoning attack will generate no warning that the user has been directed to an impostor's web server.
- E. Each root server is a single point of failure, as the root servers are each uniquely responsible for handling queries from clients in a particular geographic area.

[Total for this part: 15 marks]

END OF PAPER