Stanford University
Computer Science Department

Fall 2006 Comprehensive Exam in Networking

You are allowed 1 hour to complete this exam.
(i) This exam is closed book, closed lap-top and closed notes.
(ii) Write your solution in your blue-book. Be sure to write your name and student ID clearly on the front of the book.
(iii) Don't panic! Be sure to start by reading the exam all the way through. Then answer the questions in whatever order you choose.
(iv) Show your reasoning clearly. If your reasoning is correct, but your final answer is wrong, you will receive most of the credit. If you just show the answer without reasoning, and your answer is wrong, you may receive no points at all.

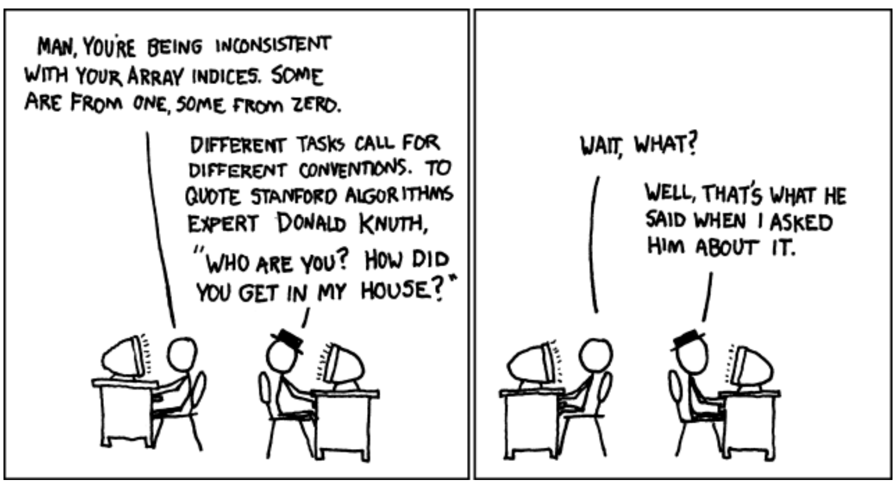

The Stanford Honor Code

In accordance with both the letter and spirit of the Honor Code, I didn't cheat on this exam.

Signature:____

Enter your magic number here:____

Multiple Choice Questions.

Instructions: *in the following questions, check all listed assertions that appear to be correct. There is at least one correct assertion per question, but there may be more. Each correct assertion checked will earn you one point.* ***For each incorrect assertion you check, you will lose one point.*** *If you don't know an answer, checking no assertion will neither earn you nor lose you any points. Yep, you could end up with negative points on this section.*

1. **Layering.** "Layering" is commonly used in computer networks because:
   (a.) It forces all network software to be written in ANSI 'C'.
   (b.) Encapsulation is the lowest overhead method to transmit data.
   (c.) It allows widespread code and implementation re-use.
   (d.) It keeps networks warm enabling them to run faster.

2. **Elasticity Buffer.** An elasticity buffer is used to store bits arriving at a network interface. If the receiving station uses a 200-bit elasticity buffer and the clocks of the transmitter and receiver have a minimum frequency of 99.999MHz and a maximum frequency of 100.001MHz, which of the following statements are true:
   (a.) All packets have to be less than or equal to 12,500 bytes long.
   (b.) All packets have to be less than or equal to 4500 bytes long.
   (c.) The two clocks have a tolerance of +/- 100ppm.
   (d.) The two clocks have a tolerance of +/- 10ppm.
   (e.) The transmitter's clock is always faster than the receiver's clock.

3. **IP.** During normal IP packet forwarding by a router, which of the following packet fields are updated?
   (a.) IP header Source address
   (b.) IP header Destination address
   (c.) IP header TTL
   (d.) IP header checksum
   (e.) Destination UDP address
   (f.) Destination UDP port number

4. **TCP**. In the TCP protocol, the receiver advertises the current size of the receive window to the sender. The sender uses this information to control congestion on the network.
   (a.) True.
   (b.) False.

5. **IP addresses.** You are given the following network prefixes:

171.64.179.0/24	171.64.180.0/24
171.64.181.0/24	171.64.182.0/24
171.64.183.0/24	171.64.184.0/24

   You are to aggregate the prefixes together into the smallest possible number of shorter prefixes. What is the smallest number of prefixes, and what are they?
   (a.) 1 prefix; 171.64.179.0/21
   (b.) 2 prefixes; 171.64.179.0/23, 171.64.182.0/23
   (c.) 2 prefixes; 171.64.179.0/24, 171.64.181.0/22
   (d.) 2 prefixes; 171.64.179.0/22, 171.64.183.0/24
   (e.) 3 prefixes 171.64.179.0/24, 171.64.180.0/22, 171.64.184.0/24

6. **TCP.** The TCP protocol uses a sliding window protocol. The window size varies because:
    (a.) Routers along the route advertise a varying window size to prevent congestion.
    (b.) The destination advertises a reduced window size when its buffers are congested.
    (c.) The destination advertises a reduced window size when packets take a long time to reach it.
    (d.) The source reduces its window size when it detects that congestion is occurring.

7. **Fair Queueing.** Which of the following are true:
    (a.) A fair queueing scheduler used in a router transmits one packet at a time.
    (b.) A fair queueing scheduler used in a router transmits just one bit from each packet at a time before moving onto the next packet.
    (c.) If traffic arriving at each router in a network is leaky-bucket constrained, and if each router uses weighted fair queueing schedulers, then bounds can be placed on the end-to-end delay of each packet.
    (d.) If a fair queueing scheduler calculates the finishing time of two packets, *A* and *B*, such that *A* is scheduled to depart before *B*, then at a later time as new packets arrive, the scheduler may change its mind and schedule *B* ahead of *A*.
    (e.) Weighted fair queueing allows a router to provide each flow with a weighted share of the link capacity.

8. **TCP**. TCP guarantees a reliable, in-order stream of data. Which of the following are true?
    (a.) TCP guarantees that if a byte did not reach the receiver, the sender will find out and be able to retransmit the data.
    (b.) TCP guarantees that any error that corrupts the TCP header is detected by the receiver
    (c.) TCP guarantees that any error that corrupts the TCP payload is detected by the receiver
    (d.) When a sender receives an acknowledgement with ACK sequence number A, it knows that all the data up to and including byte A-1 has been correctly received by the destination
    (e.) When a sender receives an acknowledgement with ACK sequence number A, it knows that all the data with a sequence number greater than A has not been received by the destination

9. Link layer. Which of the following are true?
    (a.) An Ethernet switch can interconnect a 10Mb/s Ethernet network and a 1Gb/s Ethernet.
    (b.) An Ethernet hub can interconnect a 10Mb/s Ethernet network and a 1Gb/s Ethernet.
    (c.) An Ethernet network cannot detect collisions until it has computed a checksum over the frame.
    (d.) 4B/5B is considered more efficient than Manchester encoding because more user data is transmitted in same amount of time.
    (e.) The 802.11b wireless protocol incorporates a link-layer ACK not present in regular Ethernet.

Longer questions

10. **(3 points) IP**. Explain why IPv4 fragments are reassembled at the end-point, rather than at an intermediate router along the path.

11. **(5 points) TCP sequence numbers**.
    (a.) Why does TCP use a 32-bit sequence number, instead of, say, 16 bits?

    (b.) Compute the maximum data rate at which a sender needs to worry about TCP sequence number wrap-around, assuming a maximum segment lifetime of two minutes.

12. **(20 points) TCP Congestion Control.** In this question we'll find approximate equations for the throughput of the TCP AIMD mechanism.
    (a.) Why does TCP use *additive* increase rather than *multiplicative* increase?

    The graph below shows the "sawtooth" evolution of the TCP sender's window size as a function of time. W is the maximum window size (measured in packets). In this question, assume that all packets are P bits long, and that exactly one packet is dropped every time the window size reaches W.

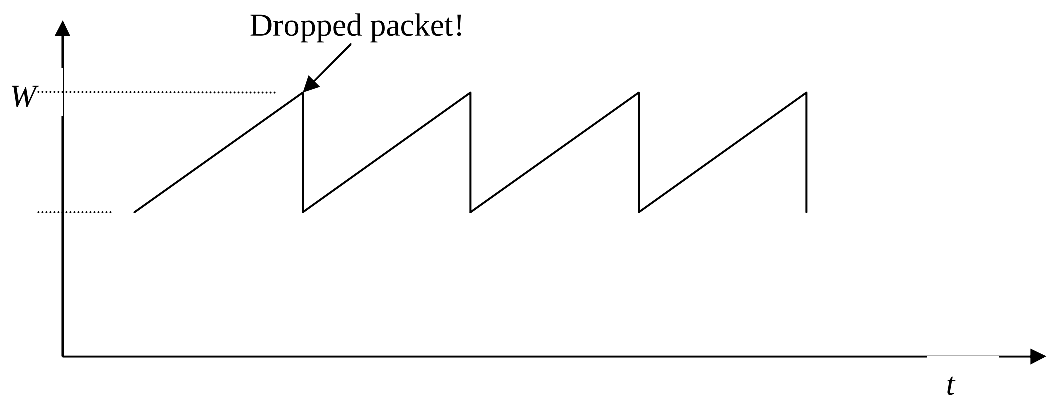

(b.) If we ignore the "slow-start" phase at the beginning of the flow, show that the average rate at which the transmitter sends packets is given by:

$$\overline{R} = \frac{3}{4}\frac{W}{RTT} \text{ packets per second.}$$

RTT is the round-trip time, which we will assume is constant.

(c.) Show that the fraction of packets dropped, L, is given by the following expression:

$$L = \frac{1}{\frac{3}{4}W\left(1+\frac{W}{2}\right)}$$

Hint: Remember that we're assuming that exactly one packet is dropped every time the window size reaches W.

(d.) Using your answers to (b) and (c), and assuming that is very large, show that the average rate at which the transmitter sends is given by:

$$\overline{R} \approx \frac{\frac{3}{4}\sqrt{\frac{8}{3}}}{\sqrt{L \cdot RTT}} P \approx \frac{1.22P}{\sqrt{L \cdot RTT}} \text{ bits per second.}$$

(e.) The *goodput* of a TCP connection is defined to be the rate at which data is sent once. i.e. the rate does not include data that is retransmitted. Is the goodput rate smaller or larger than R?

Throughout this question we assumed that RTT remains constant as the window size changes. We'll now see how this assumption is not only incorrect, but actually leads to a very different conclusion about rate R. Look again at the sawtooth figure. We say that the TCP flow is in equilibrium – the average rate available to the flow is constant; the only reason it follows a sawtooth is that TCP is varying the window size to see if the available rate has changed (which it hasn't). Each time the window size reaches W, the router at the bottleneck link drops a packet, which means the router's buffer is full. Assume that when the window size reaches the "trough" (i.e. bottom of the sawtooth) the router's buffer is empty (this is indeed the case if the buffer is sized accordingly, but you **don't** need to prove it here). Each time the source increases the window size from the trough to the peak the buffer will increase in size (and hence so will RTT) until the buffer overflows again.

(f.) How much does RTT increase each time the window size is increased? In this case, what is the actual average rate, R, at which the source sends?

13. **(20 points) Internet Design.** The Internet was originally designed to: (1) Quickly recover from link and router failure; (2) Efficiently use expensive long-haul links.

(a.) Explain what design choices were made to accomplish these two characteristics
(b.) Comment on whether you think the Internet accomplishes both goals, how it compares to alternative design choices, and whether you think these two design choices are relevant any more.

It was observed that a design choice was made – the so-called "End to End Principle".

(a.) Explain the end-to-end principle.
(b.) Give three examples that suggest the end-to-end principle still characterizes the Internet.
(c.) Give three examples that suggest it doesn't any more.