CSE 123: Computer Networks
Fall Quarter, 2010
FINAL EXAM

Instructor: Alex C. Snoeren

Name **SOLUTIONS**

Student ID ____________________

Question	Score	Points
1	20	20
2	15	15
3	15	15
4	30	30
5	30	30
6	30	30
7	10	10
Total	150	150

This exam is **closed book**. You are allowed one 8.5x11-inch, double-sided sheet of paper containing whatever you would like (a "crib sheet"). The exam contains seven questions of differing point values. Each question is clearly labeled with its value. Please answer all questions in the space provided. You have three hours to complete this exam. As with any exam, I suggest you read through all the questions first before answering any of them. Note that the final question is a freebie; you will receive full credit regardless of your answer. I would, however, appreciate your feedback if you have time remaining after completing the remainder of the exam.

GOOD LUCK!

1. (20 pts) True/False. Determine whether each of the following statements is true or false. No explanation is necessary; partial credit will not be awarded.

   a) All nodes connected to the Internet must implement UDP.

   False

   b) Channel noise leads to signal attenuation.

   False; just a drop in the SNR ratio.

   c) Media Access Control is a function of the data-link layer.

   True

   d) Switches decrement the TTL field in the IP header.

   False; only routers do.

   e) RTS/CTS combats the Hidden Terminal problem.

   True

   f) Token Ring is an example of a contention-based MAC protocol.

   False

   g) A one-bit parity scheme has Hamming distance 1.

   False; it has a Hamming distance of 2.

   h) FEC can be more efficient than ARQ in a broadcast environment with many receivers.

   True

   i) IS-IS is a distance vector routing protocol.

   False; it's link state.

   j) BGP exchanges link weights.

   False; BGP exchanges path vectors.

2. (15 pts) Short answer. Concisely answer each of the following questions.

a) Why are digital encoding schemes with more frequent transitions preferred?

Assists in clock recovery and avoids baseline drift.

b) What is the minimum value for the timeout of a reliable transmission protocol? Why?

One RTT; anything shorter woudld timeout before the ACK had a chance to arrive.

c) How does 802.11 implement virtual carrier sense?

It uses both the NAV and, optionally, RTS/CTS.

d) What is the difference between congestion control and flow control?

Congestion control prevents overrunning buffers in the network, while flow control prevents overflowing the receiver's.

e) What are the drawbacks of sentinel-based framing?

Sentinel-based framing requires the sender to do some form of stuffing, and the receiver to reverse the process. This adds both complexity and overhead to the transmission process, as well as decreasing the efficiency of the link.

3. (15 pts) More short answer. Provide brief answers for the questions below.

a) What is the difference between the bit rate and baud rate of a signal?

Baud rate is the speed at which symbols can be communicated across a channel; bit rate depends on the information rate of each symbol. I.e., if each symbol corresponds to more than one bit, bit rate will exceed baud rate.

b) Why does distance-vector routing scale better than link-state routing? Be specific.

Routers in distance vector protocols only receive updates from their immediate neighbors, while routers must manage LSPs from every router in the network in link-state protocols.

c) Explain the difference between traffic shaping and traffic policing.

Traffic shaping controls the rate of sending by buffering excess incoming traffic, while traffic policing drops any excess.

d) Give one advantage and one disadvantage of window-based flow control vs. rate-based flow control.

Window-based flow control is simple to implement and does not require fine-grained clocks, but leads to more bursty behavior than rate-based flow control.

e) List two reasons why intra-domain routing protocols are not suitable for inter-domain routing.

1) ASes do not wish to expose their internal topology to other ASes, and 2) as discussed in part b), the number of messages each router has to process scales with the number of routers in the network, not the number of neighbors, so link-state will not work well in very large networks.

4. (30 pts) Fair Queuing.

Answer the following questions about the network of four routers below. Each link has capacity of 1 Mbps. You can assume there is no contention on the access links, or for router backplane resources; i.e., the only constraints are the link capacities between routers. There are four flows in the network, labeled $F1 \ldots F4$, that traverse the routers indicated. $F4$ shares links with every other flow, traversing $R3 \rightarrow R1 \rightarrow R2 \rightarrow R4$.

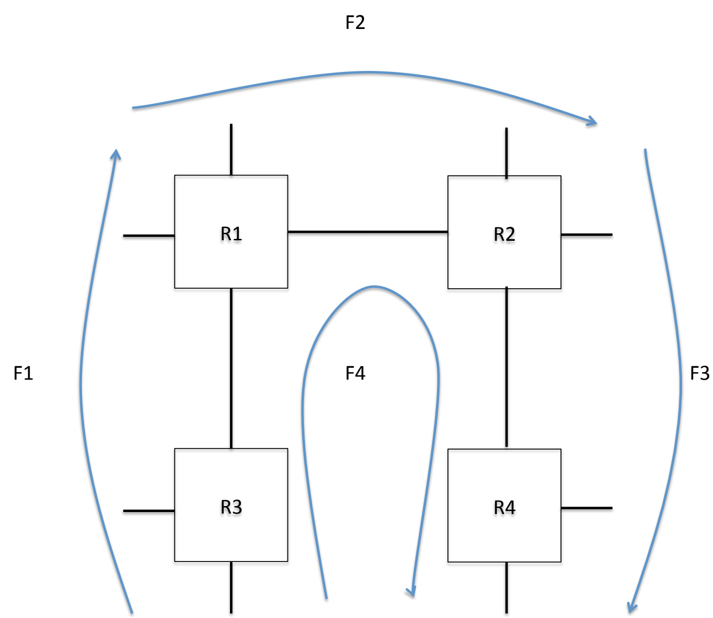

a) Assume each router implements FIFO queuing. If each flow consists of an identical, 1-Mbps constant bit rate UDP flow with equal packet sizes, what will the resulting rate be for each flow? You can assume that FIFO drops packets with uniform probability.

$F1$ and $F4$ compete at $R3$ 1:1, giving each 1/2 Mbps. $F4$ then competes at $R1$ with $F2$ at 1/2:1, giving $F2$ 2/3 Mbps and $F4$ 1/3 Mbps. Hence, $F4$ competes at $R2$ with $F3$ at 1/3:1, resulting in a final allocation for $F3$ of 3/4 Mbps and 1/4 Mbps for $F4$.

b) Now consider the case where all routers implement fair queuing. What would be the throughput of each flow now?

Every flow's fair share at each link is 1/2 Mbps. $F4$ will always be at least 1/2 Mbps when it competes with each other flow, so all flows get precisely their fair share: 1/2 Mbps each.

c) Finally, assume all routers implement WFQ, and each flow is assigned a weight equal to its number (i.e., flow $F4$ gets weight 4). What are the resulting throughputs?

$F1$ and $F4$ compete 1:4 at $R3$, giving $F1$ 1/5 Mbps and $F4$ 4/5 Mbps. $F4$ then competes with $F2$ at 4:2, so $F2$'s fair share is 1/3 Mbps, giving $F4$ 2/3 Mbps. Finally, $F4$ competes with $F3$ at 4:3, so $F3$'s fair share is 3/7 Mbps, resulting in a final allocation for $F4$ of 4/7 Mbps.

5. (30 pts) Routing & Token buckets.

a) Consider the network of five routers below, with edge weights labeled.

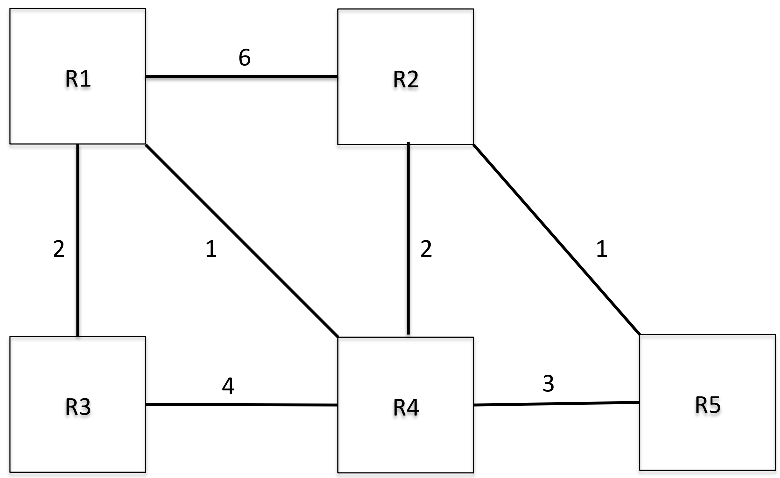

Show the step-by-step operation of Dijkstra's algorithm for $R3$ using the table below.

Step	Confirmed	Tentative	Comments
1	($R3$, 0 , -)		
2	($R3$, 0 , -)	($R1$, 2, $R1$) ($R4$, 4, $R4$)	
3	($R3$, 0 , -) ($R1$, 2, $R1$)	($R4$, 4, $R4$)	
4	($R3$, 0 , -) ($R1$, 2, $R1$)	($R4$, 3, $R1$) ($R2$, 8, $R1$)	It is cheaper to reach $R4$ through $R1$
5	($R3$, 0 , -) ($R1$, 2, $R1$) ($R4$, 3, $R1$)	($R2$, 8, $R1$)	
6	($R3$, 0 , -) ($R1$, 2, $R1$) ($R4$, 3, $R1$)	($R2$, 5, $R1$) ($R5$, 6, $R1$)	It is cheaper to reach $R2$ via $R4$
7	($R3$, 0 , -) ($R1$, 2, $R1$) ($R4$, 3, $R1$) ($R2$, 5, $R1$)	($R5$, 6, $R1$)	
8	($R3$, 0 , -) ($R1$, 2, $R1$) ($R4$, 3, $R1$) ($R2$, 5, $R1$)	($R5$, 6, $R1$)	$R5$ is the same cost through $R2$ or $R4$
9	($R3$, 0 , -) ($R1$, 2, $R1$) ($R4$, 3, $R1$) ($R2$, 5, $R1$) ($R5$, 6, $R1$)		Done.

b) Now consider a Bellman-Ford-based distance-vector protocol that uses split horizon and poison reverse. Assume that the protocol has converged, and then node $R5$ fails. Will the protocol reconverge? If so, show the sequence of updates that result. (After each update, show the nexthop and cost for $R5$ as seen by each of the remaining routers.) If not, explain why not.

This will not converge, because the topology contains a cycle of three routers (e.g., $R1, R3, R4$) where only one is directly connected to the failed node, $R5$. (In fact, when analyzing this problem, you can safely remove $R2$ from the network as no router will ever use it as a nexthop for any other.) The remaining subset of the network looks just like the case on Slide 17 in lecture 13. Hence, the routers will count to infinity because they each only poison one of the others at a time.

c) Consider a token bucket with maximum rate $R = 20$ Mbps. Suppose we want to make sure the maximum rate can only be sent for at most 5 seconds at a time, and at most 150 Mb can be sent over any 10-second window. Compute the required values for the token refresh rate, r, and the bucket depth, b.

If we send the maximum rate for 5 seconds, we send $20 * 5 = 100$ Mb. We can only send 150 Mb in *any* 10-second window, including those we send at max rate, so $150 - 100 = 50$ Mb, which we can send in the remaining 5 seconds, so $r = 50/5 = 10$ Mbps. Now, in order to ensure the bucket has enough tokens to sustain a 5-second burst at 20 Mbps, we require $b = (R - r) * 5 = (20 - 10) * 5 = 50$ Mb.

6. (30 pts) TCP congestion control.

a) (15 pts) Consider a TCP Reno (i.e., one that implements fast retransmit and recovery) flow that has exactly 50 segments to send. Assume that during the transmission, exactly four packets are lost: the 4th, 5th, 22nd, and 48th; no other losses occur. Using the graph below, plot the evolution of the congestion window as each segment is sent. You may measure cwnd and time in whatever units you find convenient. (Do not inflate the window due to duplicate ACKs.)

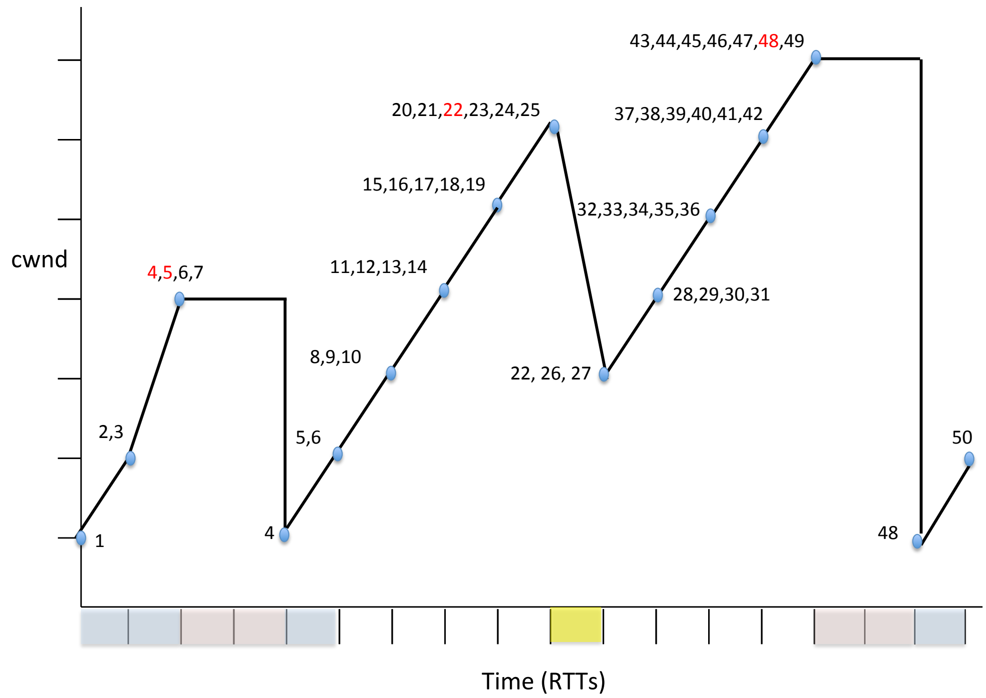

b) (10 pts) Label your plot above indicating the regions where slowstart, timeout, congestion avoidance, and fast retransmit occur.

Timeouts, shown in red, occur after the loss of segments 4, 5, and 48. Slowstart periods, depicted in blue above, are until the first loss, and one RTT after each timeout. A fast retransmit of segment 22 is highlighted in yellow. Everything else is congestion avoidance. (Note I picked 2 RTTs for timeout value, you could use anything 1 RTT or greater.)

c) (5 pts) Give one reason why TCP performs poorly when traversing overflowing FIFO queues, and explain how RED attempts to address the issue.

TCP flows are bursty, so when they encounter a full FIFO buffer, they will experience large numbers of losses in a window. RED spreads losses out. Also, TCP doesn't realize there's congestion until the queue is full, but RED will signal congestion earlier by dropping before the queue fills.

7. (10 pts) You will receive full credit for this question regardless of how you respond, so DO NOT ANSWER IT UNTIL YOU ARE FINISHED with the remainder of the exam. There is no penalty if you don't get to it.

   a) What topic covered in this course did you find the most interesting?

   b) What topic did you find least interesting, and why?

   c) How long did you spend on each project? Which one did you prefer?

   d) Was there a networking topic you wish we had covered?

   e) Is there anything you'd suggest the professor to do differently next time he teaches this course?