CS164 Final Exam

Problem 1. Consider the *go-back-n* sliding window algorithm.

(a) For a data-link layer protocol, the sender and receiver are connected by a single physical-layer link. Briefly explain why, in this type of network, the receiver does *not* need a Receive Window Size greater than one packet.

(b) Does the same situation also apply to a higher-layer protocol, like IP? Explain your answer.

Problem 2. Alice wants to transfer some data to Bob over the Internet using TCP/IP. Assume that: (i) the Round-Trip Time, from Alice to Bob and then back to Alice, is 1 second; (ii) the Maximum Segment Size on the medium is 10,000 bits including all overhead; (iii) the links are running at a data rate of 50,000 bits/sec., and (iii) Bob's advertised receive window is very large. Alice uses the TCP *slow start* congestion avoidance algorithm, and Bob transmits a "zero length" acknowledgement packet back to Alice as soon as he sees the end of each of her arriving packets. There are no errors. Draw and label a timing diagram that shows the exact times at which Alice transmits her first 10 packets.

Problem 3. A department wants to interconnect its three Ethernets X, Y, and Z.

(a) Suppose they buy three two-port bridges, and connect $B1$ to networks X and Y, $B2$ to networks X and Z, and $B3$ to networks Y and Z. What are the paths that packets follow to get from one network to another, and how are these paths determined?

(b) Now suppose they buy three two-port routers, and connect $R1$ to networks X and Y, $R2$ to networks X and Z, and $R3$ to networks Y and Z. Does using routers instead of bridges make a difference? Compare the two approaches in terms of what paths are used, how the paths are determined, the path lengths and maximum total throughput they can support.

Problem 4. *Two streams* of traffic (starting from Alice and Bob, respectively) meet at a certain router on their way to the *same destination*. Assume that the initial part of the traffic looks like the following: (i) Alice sends four 1K byte packets *A1, A2, A3, A4*, which arrive at the router at respective times 0.1, 0.3, 0.8, and 1.1 seconds; (ii) Bob sends three 2K byte packets *B1, B2, B3*, which arrive at the router at respective times 0.05, 0.06, and 0.08 seconds; and (iii) the outgoing link heading towards their common destination runs at 5K bytes per second. Find the *order in which these packets are transmitted* on the outgoing link, when the router uses:

(a) A single First-In First-Out queue for the outgoing link;

(b) Separate queues for each stream that are served in Round-Robin order; or

(c) Fair Queueing.

Problem 5. Briefly explain *bit stuffing* and *character stuffing*. Which, if any, of these techniques is used in the following and why?

(a) ATM (cell layer).

(b) ATM (AAL3/4 and AAL5).

(c) Ethernet.

(d) Token Ring.

Problem 6. In the "Cheating Husbands" story, Queen Margaret found a way to eliminate all cheating husbands within 3 days using a modulo 3 protocol. Now suppose that the wives have not recorded the *exact number* of cheating husbands they know. Instead, each wife simply remembers (i) whether or not she knows about *any* cheating husbands, and (ii) if so then the number of cheaters modulo 4. Thus, if the total number K of cheating husbands that a particular wife knew was 0, 3, or 13 for example, then we assume that she remembers that $K = 0$, $K \equiv 3(mod\ 4)$ or $K \equiv 1(mod\ 4)$, respectively. Modify Queen Margaret's protocol to handle this new situation. Note that her modulo 3 protocol *does not work* for this problem, so no credit is granted for just writing it down.

Problem 7. Consider the problem of designing a broadcast or multicast protocol from the final project. Assume that none of the application programs using the broadcast or multicast protocol knows all other host names in the network. However, each router does know all the host names in its domain. Also assume that all the hosts are connected to a single broadcast network (like an Ethernet). Packets can be lost or damaged by network errors, but we assume no failures of a host or application program occur.

(a) Suppose one of the hosts sends a broadcast message. Can we guarantee that all hosts in the network eventually receive the broadcast message? Justify your answer.

(b) Now suppose that one of the hosts sends a multicast message to a particular group of target hosts. Can we guarantee that each target host receives a copy, even if only domain names are specified? If not, explain why. If so, show how it could be done.