Distributed Computing and Networking—CS 461

MIDTERM EXAM — March 30, 2000

1. Briefly describe Nagel's agorithm, and explain what problem it solves. [**5 points**]

2. Consider a PC-based router with a 64-bit wide $\times$ 66MHz PCI bus. This bus has a peak bandwidth of 4.2 Gbps, but it is generally argued that PCI overheads limit you to half the peak bandwidth, so assume that the PCI bus has an effective bandwidth of 2Gbps. Also assume the PC's processor is fast enough to forward 500k minimus-sized packets-per-second. [**10 points**]

   (a) Assuming you want to engineer the router so it is able to keep up with line speeds (i.e., so it does not drop packets on input), and that the average packet size is 1000 **bits**, how many 100Mbps ethernet line cards can the router support?

   (b) Again assuming you want to keep up with line speeds, but this time you are engineering for full-sized packets of 1500 **bytes**, how many 100Mbps ethernet line cards can the router support?

   (c) Suppose the ethernet cards have a microprocessor with enough cycles and memory to determine the correct output port for 80% of the packets that arrive on the interface, and that it is possible for one card to write packets directly to another card without CPU involvement. (The other 20% of the packets have to be sent to the CPU to be classified.) How many ethernet cards can this configuration support if you want to keep up with line-speeds for full-sized packets of 1500 bytes?

3. Consider a network path consisting of four 10 Mbps ethernet links connected by three store-and-forward switches. Assume that each link introduces a propagation latency of 10 microseconds, but the switches introduce no queuing delays. [**10 points**]

   (a) Calculate the path's one-way latency (total packet delay) for sending a 1KB message

   (b) Calculate the effective bandwidth for sending a 1KB message across this path.

   (c) Calculate the effective bandwidth for sending a sequence of 1KB messages if the source must wait for a 20-byte ACK of the previous packet before it can send the next packet.

4. You are hired to design a reliable byte-stream protocol that uses a sliding window (like TCP). This protocol will run over a 622Mbps network. The RTT of the net is 100ms, and the maximum segment lifetime is 50 seconds. How many bits would you include in the AdvertisedWindow and SequenceNum fields of your protocol header? (Do not round to the nearest power of two.) [**10 points**]

5. What two scaling problems is the Internet experiencing? Briefly describe subnets and supernets (CIDR), and explain how each addresses these two scaling problems. [**10 points**]

6. The RED Gateway mechanism maintains two queue thresholds: MIN and MAX. Explain how a RED gateway processes an arriving packet in terms of these two thresholds. Explain why it is possible for the queue in a RED gateway to contain more than MAX packets. [**5 points**]

7. For the two TCP traces shown below [**15 points**]:

   (a) Identify the feature that the version of TCP illustrated in the second trace has that the version of TCP illustrated in the first trace does not have? Explain why this feature results in the different behavior.

   (b) Explain why both traces have heavy losses at approximately 0.5 sec.

   (c) For the first trace, explain why the congestion window is flat between 0.5s and 2.0s.

   (d) In the second trace, explain what the TCP congestion control algorithm is doing between 2.0s and just after 3.5s.

   (e) In the second trace, explain what triggers the change in the slope of the congestion window at just after 4.0s.