Department of Computer Science and Engineering, National Sun Yat-Sen University
First Semester of 2003, PhD Qualifying Exam: Computer Networks

(10 points for each problem)

1. Consider a simple protocol for transferring files over a link. After some initial negotiation, A sends data packets of size 1 KB to B; B then replies with an acknowledgment. A always waits for each ACK before sending the next data packet; this is known as *stop-and-wait*. Packets that are overdue are presumed lost and are retransmitted. Suppose that the link can deliver out of order, and that sometimes a packet can be delivered as much as 1 minute after subsequent packets. How many bits for sequence number are enough in the protocol for this situation?

2. When closing a TCP connection, why is the two-segment-lifetime timeout NOT necessary on the transition from LAST_ACK to CLOSE?

3. Suppose a host wants to establish the reliability of a link by sending packets and measuring the percentage that are received; router, for example, do this. Explain the difficulty of doing this over a TCP connection.

4. If a host A receives two SYN packets from the same port from remote host B, the second may be either a retransmission of the original or else, if B has crashed and rebooted, an entirely new connection request. (a) Describe the difference as seen by host A between both cases. (b) Given an algorithmic description of what the TCP layer needs to do upon receiving a SYN packet. Consider the duplicate/new cases above, and the possibility that nothing is listening to the destination port.

5. Suppose TCP operates over a 1-Gbps link. (a) Assuming TCP could utilize the full bandwidth continuously, how long would it take the sequence numbers to wrap around completely? (b) Suppose an added 32-bit timestamp field increments 1000 times during the wraparound time you found above. How long would it take for the timestamp to wrap around?

6. During linear increase, TCP computes an increment to the the congestion window size as $\frac{MSS^2}{cwnd}$. Explain why computing this increment each time an ACK arrives may NOT result in the correct increment. Please revise the increment and show a more precise one.

7. Defeating TCP congestion-control mechanism usually requires the explicit cooperation of the sender. However, consider the receiving end of a large data transfer using a TCP modified to ACK packets that have NOT yet arrived. It may do this either because not all of the data is necessary or because data that is lost can be recovered in a separate transfer later. What effect does this receiver behavior have on the congestion-control properties of the session? Can you devise a way to modify TCP to avoid the possibility

of senders being taken advantage of in this manner?

8. The link-layer sliding window protocol can be used to implement flow control. We can image doing this by having the receiver delay ACKs, that is, not send the ACK until there is free buffer space to hold the next frame. In doing so, each ACK would simultaneously acknowledge the receipt of the last frame and tell the source that there is now free buffer space available to hold the next frame. Explain why implementing flow control in this way is not a good idea.

9. Suppose some repeaters (hubs), rather than bridges, are connected into a loop. (a) What will happen when somebody transmits? (b) Why would the spanning tree mechanism be difficult or impossible to implement for repeaters? (c) Propose a mechanism by which repeaters might detect loops and shut down some ports to break the loop. Your solution is not required to work 100% of the time.

10. You are an Internet service provider; your client hosts connect directly to your routers. You know some hosts are using experimental TCPs and suspect some may be using a ‘‘greedy’’ TCP with no congestion control. What measurements might you make at your router to establish that a client was not using slow start at all? If a client used slow start on startup but not after a timeout, could you detect that?