

# CS536 Homework 2

due Tue, Mar 10th, 1:30pm

March 4, 2009

Submit your homework electronically in a text or PDF file. Email your solution to the professor by 1:30pm on Tuesday, Mar 10th. Late submission will not be accepted, nor will collaboration.

Note that Problem 3 requires a diagram. If you prefer, you may hand-draw that diagram and submit it at the beginning of class on Tuesday, Mar 10th. Make sure to indicate that in your writeup, and on your diagram make sure to include your personal information.

To receive credit, all submissions should contain this statement at the beginning: "By turning in this homework submission, I certify that this work was done solely by me." Questions about this homework should be directed to the TAs or the Professor, not to fellow classmates.

## Problem 1 (10 pts)

Suppose Client  $X$  initiates an SSH session with Server  $S$ . At about the same time, Client  $Y$  initiates an SCP session with Server  $S$ . Provide possible source and destination port numbers for

- a. The segments sent from  $X$  to  $S$ .
- b. The segments sent from  $Y$  to  $S$ .
- c. The segments sent from  $S$  to  $X$ .
- d. The segments sent from  $S$  to  $Y$ .
- e. If  $X$  and  $Y$  are different hosts, is it possible that the source port number in the segments from  $X$  to  $S$  is the same as that from  $Y$  to  $S$ ?
- f. How about if they are the same host?

## Problem 2 (10 pts)

UDP and TCP use 1s complement for their checksums. Suppose you have the following four 8-bit bytes: 01101101, 11011011, 00010001, 01000010, What is the 1s complement of the sum of these 8-bit bytes? (Note that although TCP

and UDP use 16-bit words in computing the checksum, for this problem you are being asked to consider 8-bit sums.) Show all work. Why is it that UDP takes the 1s complement of the sum; that is, why not just use the sum? Is it possible that a 1-bit error will go undetected? How about a 2-bit error?

### Problem 3 (10 pts)

Consider the rdt3.0 protocol. Draw a diagram showing that if the network connection between the sender and receiver can reorder messages (that is, that two messages propagating in the medium between the sender and receiver can be reordered), then the alternating-bit protocol will not work correctly (make sure you clearly identify the sense in which it will not work correctly). Your diagram should have the sender on the left and the receiver on the right, with the time axis running down the page, showing data (D) and acknowledgement (A) message exchange. Make sure you indicate the sequence number associated with any data or acknowledgement segment.

### Problem 4 (10 pts)

Consider a reliable data transfer protocol that uses only negative acknowledgements. Suppose the sender sends data only infrequently. Would a NAK-only protocol be preferable to a protocol that uses ACKs? Why? Now suppose the sender has a lot of data to send and the end-to-end connection experiences few losses. In this second case, would a NAK-only protocol be preferable to a protocol that uses ACKs? Why?

### Problem 5 (10 pts)

Consider transferring an enormous file of  $L$  bytes from Host  $A$  to Host  $B$ . Assume an MSS of 1460 bytes.

- a. What is the maximum value of  $L$  such that TCP sequence numbers are not exhausted? Recall that the TCP sequence number field has 4 bytes.
- b. For the  $L$  you obtain in (a), find how long it takes to transmit the file. Assume that a total of 66 bytes of transport, network, and data-link header are added to each segment before the resulting packet is sent out over a 10 Mbps link. Ignore flow control and congestion control so  $A$  can pump out the segments back to back and continuously.
- c. Suppose now the link were a 100 Gbps link. Retaining other parameters from (b), how long will it take to transmit the file now?
- d. If the file were instead were longer than  $L$ , some sequence numbers would be re-used. Assume now that the Host  $B$  acknowledges segments from

Host  $A$ , and that there exist two paths between  $B$  and  $A$ . If a handful of acknowledgements take an alternate path, which experiences severe queueing delay, they might be confused with new acknowledgements if they arrive after the data with the same sequence number is sent. If the propagation delay is 300ms on each path (in both direction), how much queueing delay would be necessary to cause this confusion to occur? What implications might this have for the design of high-bandwidth network transport protocols?

## Problem 6 (10 pts)

In this problem we consider the delay introduced by the TCP slow-start phase. Consider a client and a Web server directly connected by one link of rate  $R$ . Suppose the client wants to retrieve an object whose size is exactly equal to  $15S$ , where  $S$  is the maximum segment size (MSS). Denote the round-trip time between client and server at  $RTT$  (assumed to be constant). Ignoring protocol headers, determine the time to retrieve the object (including TCP connection establishment) when

- $4S/R > S/R + RTT > 2S/R$
- $S/R + RTT > 4S/R$
- $S/R > RTT$

## Problem 7 (20 pts)

**TCP Vegas.** The TCP Vegas techniques are primarily designed with the idea that segments are only lost due to buffers which are full. But suppose there is a physical medium problem (a so-called “dirty link”), and that every third segment is dropped due to bit errors. Compare the impact this will have on TCP Reno with the impact it will have with TCP Vegas. How (if at all) will the Vegas retransmit, congestion avoidance, and slow start modifications react differently from TCP Reno. Next, if you knew with certainty that every third segment would be lost due to bit errors, how would you modify Vegas to work around this problem.

## Problem 8 (20 pts)

**XCP.** When you critically evaluate a paper, one good strategy is to identify assumptions (implicit or explicit) made by the authors in their design, implementation, or evaluation, and then consider how their system might behave if these assumptions turn out to be invalid. Sometimes the authors explicitly state they are only targeting the cases in which these assumptions are true, but other times they represent details which are overlooked limitations of their system.

Identify two such assumptions about XCP. Pick one, and analyze how XCP would behave if the assumption were not true.