

UNIVERSITY COLLEGE LONDON  
DEPARTMENT OF COMPUTER SCIENCE  
3035/GZO1: Networked Systems  
Kyle Jamieson and Brad Karp  
Individual Coursework 3

Distributed: 24th November 2011; Due: 2nd December 2011, 9:05 AM

Answer all of the following problems. Either handwritten or typeset solutions are fine, but if you write by hand, please ensure your answers are legible. Please show all work! We cannot award credit for correct answers if their complete derivation isn't shown. Please state clearly all assumptions you make while solving a problem. This coursework is worth 8% of the total marks for 3035/GZO1.

Please monitor the 3035/GZO1 Moodle forum during the period between now and the due date for the coursework. Any announcements (*e.g.*, helpful tips on how to work around unexpected problems encountered by others) will be sent to the list.

**Hand-in instructions:** Hand in hardcopy for your solutions at the 5th floor reception desk in the CS department (in MPEB) at 9:05 AM on the 2nd of December. There is no provision for electronic submission of this coursework. If you would prefer, you may also submit hardcopy in lecture at the start of lecture (9:05 AM) on Thursday, the 1st of December. Any late submissions should be handed in at the 5th floor reception desk in the CS department (in MPEB).

**Collaboration:** Collaboration is *not permitted* on this problem set; you may not discuss the problems or their solutions with anyone else (whether or not the other person is taking the class), apart from the instructors and teaching assistant. All work you submit must be your own. You may of course refer to all lecture notes and readings, and any other materials you wish (textbooks, papers, or material found on the Internet).

**Late days:** If you wish to use any late days on this assignment, you *must* state how many days you wish to use clearly at the top of the first page. If your written assignment does not include a request to use late days, you will not be permitted to use late days on the assignment later. (Full late-day policies are stated on the 3035/GZO1 web site.)

## 1. TCP and High-Bandwidth Paths

Your physicist friend C. King Higgs is running an experiment on the Large Hadron Collider (LHC) nuclear physics apparatus at CERN. He has enlisted your help as someone knowledgeable in networking in sending the data from Switzerland, where the LHC is located, to his own university laboratory at Caltech in California.

The experiments your friend is conducting will generate measurement data at the impressive rate of 5 Gbits/s. He has arranged two 10 Gbits/s links, one to connect each of CERN and Caltech to the same international Internet Service Provider's (ISP's) backbone, and this backbone has sufficient capacity to carry the 5 Gbits/s data flow between CERN and Caltech.

You decide to use TCP over the Internet to carry the data between CERN and Caltech over this ISP's network. Suppose that the round-trip time (RTT) between CERN and Caltech is 200 ms. Throughout this question, ignore the bandwidth overhead of TCP and IP headers, and assume that the sender sends 1460 user data bytes per TCP segment.

- (a) How large would the receiver's advertised window and the sender's send window both need to be at a minimum in order to achieve a throughput of 5 Gbits/s? In your answer, ignore TCP congestion control—that is, assume that the congestion window is larger in size than the receiver's advertised window and sender's send window at all times. Justify your answer by explaining why TCP requires that window size.

[2 marks]

- (b) Now, let's consider some of the effects of TCP congestion control, and in particular the congestion window (`cwnd`), on this data transfer. Assume that the path between sender and receiver drops no packets, that the receiver's advertised window and sender's send window are not limiting factors, and that the TCP receiver does not use delayed ACKs. At the start of the connection between CERN and Caltech, how long will it take for the sender to reach the window size that will allow the sender to send at 5 Gbits/s (the window size that you found in the previous part of this question)? Justify your answer by explaining any relevant TCP congestion control behaviors that are invoked at the start of a connection.

[2 marks]

- (c) Suppose that the sender is currently using the exact full congestion window (`cwnd`) size needed to sustain 5 Gbits/s throughput. During the sending of this window's worth of data, a router along the path between sender and receiver drops a single packet of the sender's data. Assume that the sender detects this loss upon receiving three duplicate ACKs, and adjusts its congestion window (`cwnd`) accordingly. How long (in seconds) will it take before the sender's window once again reaches the exact full congestion window (`cwnd`) size needed to sustain a throughput of 5 Gbits/s? In your answer, you may assume that the network drops no packets after the single loss described above.

[2 marks]

- (d) Assume that there is no loss between the sender and receiver, and that all window sizes are such that the sender is able to send at 5 Gbits/s. How long would it take for the sender to cycle through the entire possible range of TCP sequence numbers for sent packets, and begin sending packets with sequence numbers it has already used to send previous packets on the same connection?

[2 marks]

- (e) Suppose that the 5 Gbits/s data flow between CERN and Caltech lasts hours. Given the answer to the previous part of this question, what best-effort packet delivery phenomenon present in the Internet would cause the reassembled byte stream produced by the TCP receiver at Caltech to differ from the original byte stream passed to TCP by the sending application at CERN? Justify your answer by explaining exactly how this phenomenon could cause the receiver to receive a different byte stream than what was sent.

[2 marks]

## 2. TCP Retransmissions, RTT Estimation, and Short Flows

A TCP sender carries out estimation of a connection's round-trip time using the improved method advocated by Van Jacobson in his classic paper, Congestion Avoidance and Control, assigned for GZ01. Recall that the sender samples the round-trip time between when each data packet is sent and when the ACK returns for that data packet. It does so by recording the time on its local clock when a data packet was sent, and recording the time on its local clock when the corresponding ACK returns.

- (a) Suppose the sender retransmits a packet after a timeout. In such cases, the sender cannot accurately sample the round-trip time for a packet that it retransmits. Why not? Justify your answer by explaining the problem that could arise, and give an example of a sequence of events that could occur in the best-effort Internet that would cause the TCP sender to arrive at an incorrect single sample of the RTT for a packet that it retransmits.

[2 marks]

- (b) Suggest a different mechanism that TCP could use to sample the round-trip time experienced by a packet that would work correctly, even for packets that are retransmitted. You are allowed to add a fixed-size additional field to the TCP header. In your answer, describe what the sender writes into this additional field, what the receiver does with the contents of this additional field, and why this method yields correct results for retransmitted packets.

[5 marks]

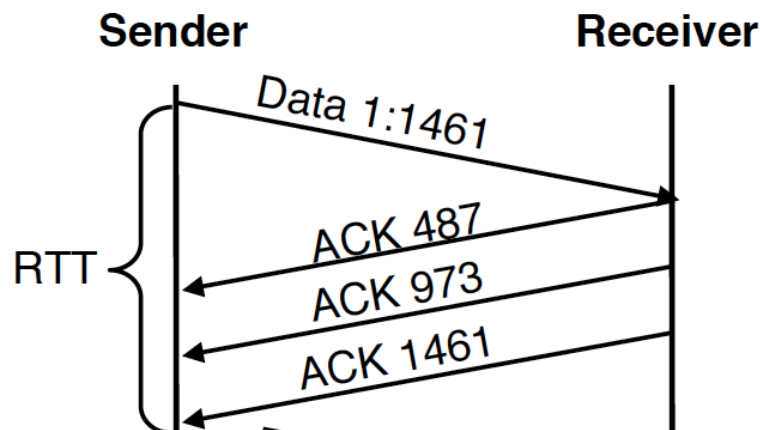
- (c) A busy web server sends predominantly short files of under 8 KB in size to web clients using TCP. A measurement study finds that when this web server retransmits data packets with TCP, 56% of the retransmissions are after a retransmit timeout (RTO), and 44% of the retransmissions are done using fast retransmit. Assuming that this web server fully and correctly implements the congestion control and avoidance algorithms specified in Van Jacobson's paper, Congestion Avoidance and Control, why is it retransmitting after RTOs so often, instead of using fast retransmit?

[3 marks]

### 3. Faster Downloads with TCP

TCP sequence number and ACK number fields in the TCP header are expressed in bytes. However, most implementations of the sending side of TCP congestion control assume that when an ACK returns, it is acknowledging an entire segment of data.

- (a) Your friend Alyssa P. Hacker wants to speed up TCP transfer times for her downloads. In particular, she'd like to receive an entire download of a file hundreds of kilobytes long in a little more than one round-trip time (after the initial three-way handshake), despite TCP's congestion control algorithm. She tells you she has implemented an "enhanced" TCP receiver, and shows you the following protocol time diagram, which begins immediately after her connection's 3-way handshake:



Given the details of TCP's behavior stated at the beginning of this question, what will the protocol time diagram look like after the server receives this traffic from Alyssa's TCP receiver?

[5 marks]

- (b) Based on the protocol time diagram Alyssa has shown you, state a general technique that a TCP receiver can use to trick a TCP sender into sending very quickly (far more quickly than TCP congestion control allows). Assume that the network path between the server and Alyssa's machine is of sufficient capacity to carry a burst of packets of arbitrary length from the server without dropping any packets.


[3 marks]

- (c) How would you change the TCP sender's implementation to defend against the anti-social behavior of Alyssa's TCP sender?

[2 marks]

### 4. Distance-Vector Routing

Suppose you have a network of routers and links as shown below. The large number in the middle of each link is that link's weight; each box is a router whose address is denoted by a letter; and the small number within a router next to each router port is the interface ID of that port on that router.



You run distance-vector routing on all these routers, using only periodic updates (no triggered updates), and without poison-reverse or split horizon.

Assume that the routers are perfectly synchronized, so that every update interval, they all send routing updates at exactly the same time. Assume that all routers and links stay up, and no routing protocol packets are ever lost. In the event of a tie, where the costs of two different next hops to the same destination are equal, a router breaks the tie by choosing the next hop with the lowest interface ID.

- (a) Suppose that all routers are powered on simultaneously.

Show how the all the routers' routing tables evolve over time, until routing fully converges. To do so, show the full routing table in each router immediately upon boot, and then show the full routing table in each router immediately after it has received the next set of updates from all its neighbors. (Recall that all routers have the same update interval, and are perfectly synchronized, so that they all send updates at the same time.)

In your answer, write routing tables in the following format:

Destination	Cost	Interface
<i>router address</i>	<i>cost</i>	<i>interface</i>
...		

[7 marks]

- (b) Suppose that in the previous part of this question, the routers did *not* use the “lowest interface ID” rule to break ties when two or more routes have the same cost to the same destination. Instead, suppose that when an update came in on interface  $i$ , and it resulted in a route with equal cost to one in the routing table, that route would take priority, and interface ID  $i$  would be written into that routing table entry in the local router.

In a realistic deployment, where routing messages are *not* sent by different routers in synchronized fashion, what negative consequence could be observed by users when the routers do not use a consistent tie-breaker rule, and instead use “most recent update wins”? State your answer in terms of end-to-end packet delivery behavior observable by users, and be careful to state any assumptions you make.

[3 marks]

**Problem set total: 40 marks**