15-441: Computer Networks
Homework 2

Assigned: September 25, 2002.
Due: October 7, 2002 in class.

In this homework you will test your understanding of the TCP concepts taught in class including flow control, congestion control, and reliability.

You must solve the homework individually. Make sure you provide all your answers in the space provided in the end. For full credit, show your work in clearly readable writing. Points will be deducted for unreadable answers. The following notation should be used for all the calculations.

Parameter	Description	Units
MSS	Maximum Segment Size	bytes/packet
W	Window Size	packets
$W * MSS$	Window Size	bytes
RTT	Round trip time	seconds
X or B	Throughput	bits/sec
ρ	Utilization	dimensionless percentage
L	Loss rate	dimensionless proportion/probability

Table 1: Notation of parameters to be used throughout the assignment

1 TCP Basics

1. True or False?
   (a) Host A is sending a large file to host B over a TCP connection. Assume host B has no data to send to host A. Host B will not send acknowledgments to host A because host B cannot piggyback the acknowledgments on data.
   (b) The size of the TCP `RcvWindow` never changes throughout the duration of the connection.
   (c) Suppose host A is sending a large file to host B over a TCP connection. The number of unacknowledged bytes that A sends cannot exceed the size of the advertised receiver buffer.
   (d) Suppose that the last `SampleRTT` in a TCP connection is equal to 1sec. Then the current value of `Time-outInterval` for the connection will necessarily be $\geq$ 1 sec.
2. Suppose host A is sending a large file to host B over a TCP connection. The two end hosts are 10msec apart (20msec RTT) connected by a 1Gbps link. Assume that they are using a packet size of 1000 bytes to transmit the file. Also assume for simplicity that ACK packets are extremely small and can be ignored.
   (a) At least how big would the window size (in packets) have to be for the channel utilization to be greater than 80%.
   (b) Assuming infinite initial threshold, no losses and competing traffic, approximately how long (in seconds) would it take for the normal slow start mechanism to achieve 80% utilization?

3. TCP waits until it has received three duplicate ACKs before performing a fast retransmit. Why do you think the TCP designers chose not to perform a fast retransmit after the first duplicate ACK for a segment is received?

2 TCP Security

Consider a misbehaving TCP receiver. The receiver modifies its TCP such that upon receiving a data segment containing N bytes, the receiver divides the resulting acknowledgment into M, where $M \leq N$, separate acknowledgments each covering one of M distinct pieces of the received data segment. For e.g. if it receives data acknowledging bytes 1 to 1000, then the receiver, for $M = 2$, will send 2 ACKs for 501 and 1001.

Consider a normal TCP sender sending data to this misbehaving TCP receiver. The sender sends a 1500 byte packet with sequence number 1. The receiver sends back $M = 3$ ACKs.

1. What packets will the sender send next in response to the 3 ACKs?

2. Assuming no losses and negligible packet transmission time, derive an expression for the sender window size during the slow start phase, in terms of n (number of RTTs) and M.

3. Can you think of a simple enhancement to the sender which can prevent the receiver from launching this attack?

3 RED/AQM

Consider a bursty greedy source. Would RED drop more packets for this source or Droptail? Explain briefly your answer.

4 Fast Recovery and TCP Modelling

Note: Remember to use Table 1 for all symbols and units. Points will be lost for incorrect units.

1. Harry Bovik says that when during steady state, achieved TCP throughput should be 75% of the bottleneck bandwidth. His reasoning is that when the throughput reaches 100% of the bottleneck, there is a packet drop causing the window to be halved to 50% and then it increases linearly back to 100% repeating the same cycle. Is he right or wrong? Explain briefly your answer.

2. Let w represent the congestion window size in TCP during steady state. The idealized congestion window size varies from $\frac{W}{2}$ to W during steady state and then a loss occurs. Suppose an additional constraint is added such that the maximum receiver window size for a connection is set to $\frac{3}{4}W$. Assuming a loss occurs when the congestion window has a size of W,

   (a) Find the loss rate L for this TCP connection in terms of W. **Hint : Calculate the number of packets sent per loss cycle**

   (b) Derive an expression for the achieved throughput B in terms of RTT, L and MSS.

5 Congestion Control

Harry Bovik thinks that if the sources in the Internet use additive increase additive decrease (AIAD) the system should still converge to fairness and efficiency.

Harry decides to use phase-plots (see pg. 269 of the book) to check his intuition for the two-user case. Figure 1 shows a simple phase plot. In the graph, the state of the system is represented by a point (x1, x2), where x1 is stream 1's current rate and x2 is stream 2's current rate. The total capacity of the link is C bits/s. Assume that after the packet loss the window sizes have been reduced according to AIAD and now the system state is in either one of the labelled regions.

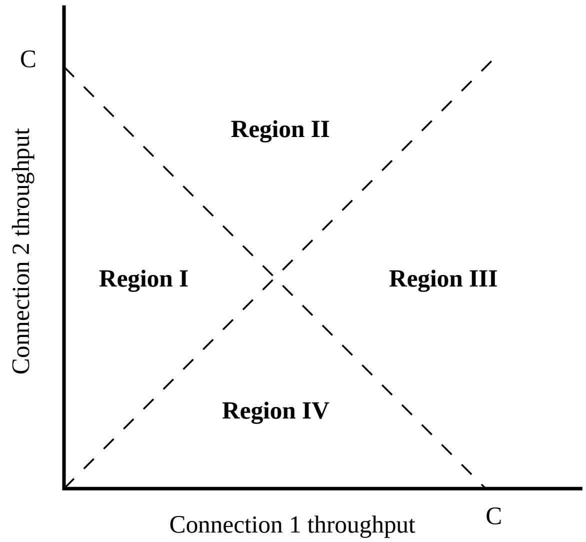

Figure 1: TCP Phase-plot

1. On the figure, for each of the labelled regions, I, II, III, and IV, draw vectors indicating the direction in which the system will move next i.e., after the AIAD response to the packet loss.

2. In each of the following regions, does the response after the packet loss increase, reduce or keep fairness the same?

   - Region I:
   - Region II:
   - Region III:
   - Region IV:
   - On the 45 degree line through the origin:

3. Does Harry's scheme converge to a fair allocation? Explain your answer briefly.

4. Does Harry's scheme converge to an efficient allocation? Explain your answer briefly.

6 Multimedia

1. Harry Bovik decides that for his video application to recover from losses over the Internet he should set the video buffer playout time equal to the current RTT to the sender. He runs his multimedia application playing videos from all around the world. He notices that for larger RTT values his video works fine while for shorter RTT values his video misses several frames. What is the mistake that Harry is making?

2. Why do you think TCP is not used for streaming multimedia (audio, video) files?

3. Harry Bovik says that since there are negligible losses within the US, the streaming multimedia applications do not need to have a playback buffer in US. What would be your response?

7 Reliability

For all parts of this question, refer to figure 2. The graph shows the packets sent and ACKs received by a TCP Reno sender. The x-axis shows the time in seconds and the y-axis shows the sequence number of the packet (or ACK) being

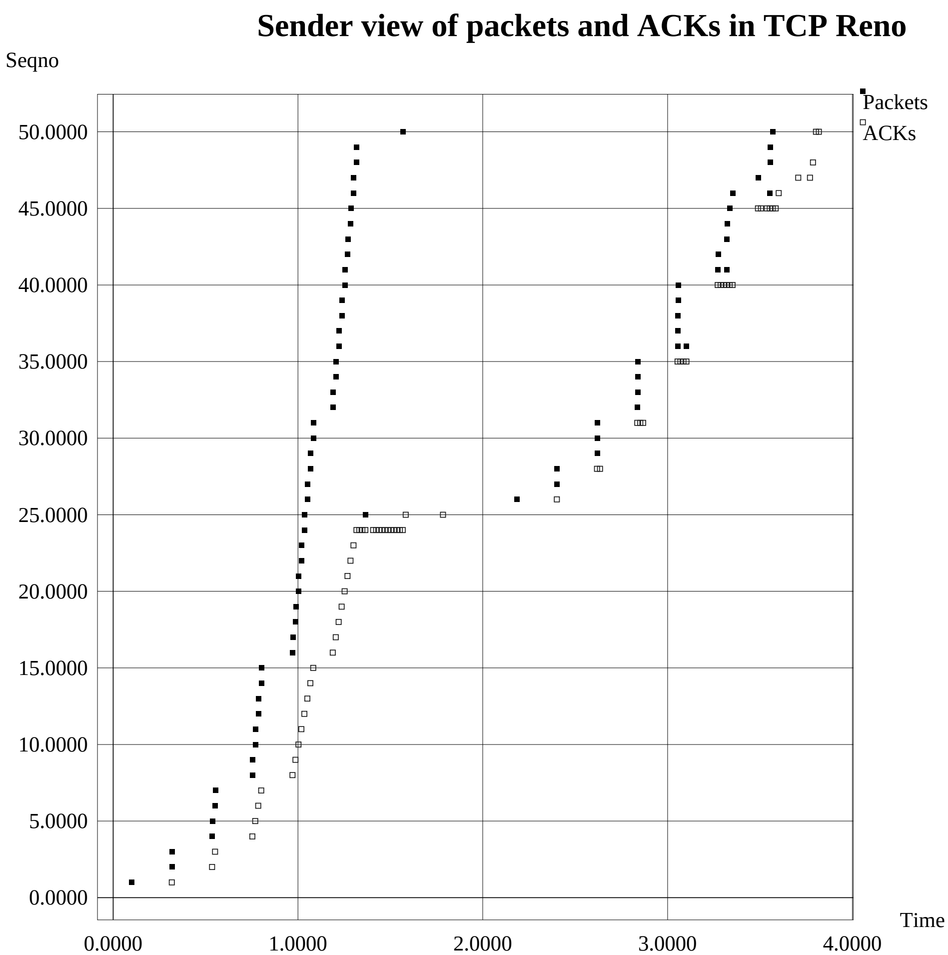

Figure 2: Reno : Fast Retransmission

sent (or received). Assume that all packets that are not lost will arrive in order at both the sender and receiver. Also assume that receiver keeps all the out of order received packets.

1. Mark on the x-axis of the graph, the time interval when the connection is in slow-start.
2. Mark the packets (if any) with × that are lost.
3. Mark the duplicate packets received by the receiver (if any) with a ◯.
4. Mark the timed out packet (if any) with a T.
5. Draw on the graph an approximate curve showing how a Tahoe TCP Sender would behave.