Carnegie Mellon
Computer Science Department.
15-744 Fall 2006 Problem Set 1

This problem set has ten questions. Answer them as clearly and concisely as possible. You may discuss ideas with others in the class, but your solutions and writeup must be your own. If you do discuss at length with others, please mention in your solution for the problem who you collaborated with. Do not look at anyone else's solutions or copy them from anywhere.

This assignment is due by **5:00pm, Wednesday, February 7th** either in class or to the course secretary in Wean Hall 8018.

A Background: Ethernet

1. This is a background question about ethernet that you should be able to answer.

   Suppose you set up a 100 Mbps Ethernet with a maximum packet size of 1500 bytes and a diameter of 1 km. You measure the worst case RTT in your Ethernet and find it to be 20.48 usec. Also, the resistance of the wire is 50 ohms.

   You have heard that when using Ethernet you should make sure that all packets have a certain minimum packet size. You want to avoid unnecessary overhead (stuffing packets to reach the minimum packet size) by choosing a small minimum packet size. How would you choose the minimum packet size?

B Multiplexing

In this problem, you will compare statistical multiplexing to time-divison multiplexing.[1]

In this problem, there are a number of communications *sessions*. In the statistical multiplexing scheme, the packets from all sessions are placed in a single queue and transmitted on a first-come, first-served (FCFS) basis. In the TDM scheme, packets from each session are queued separately. Time is divided into slots and each queue is assigned a time slot in round-robin fashion. Assume that the size of a slot is much *smaller* than the time required to transmit an entire packet on the link; each packet would be transmitted over many time slots.

A switch is termed *work conserving* if the only time it is idle is when there are no frames (in any session) waiting for service.

2. **Work conserving.** Is the TDM scheme work conserving? Is the statistical multiplexing scheme?

3. **Queueing delays.** Suppose that there are N concurrent sessions. Each session sends traffic with a Poisson distribution with rate λ packets/second. The switch can service packets at a rate of μ packets per second. Assume that the switch is *not* completely overloaded, i.e., $\mu > N\lambda$. What is the average delay observed by a frame under TDM? Under statistical multiplexing? Briefly explain what your result means and why it occurs.

4. **Utilization vs. sessions.** Assume that all sessions send frames at a simple constant rate and that the switch is configured to handle a maximum of Max sessions. Plot a graph of utilization vs the number of currently active connections N from 0 to Max for both statistical multiplexing and for TDM. Your plot should have axes like the template below:

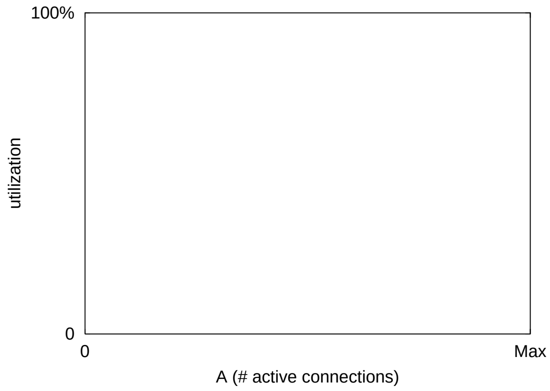

TDM

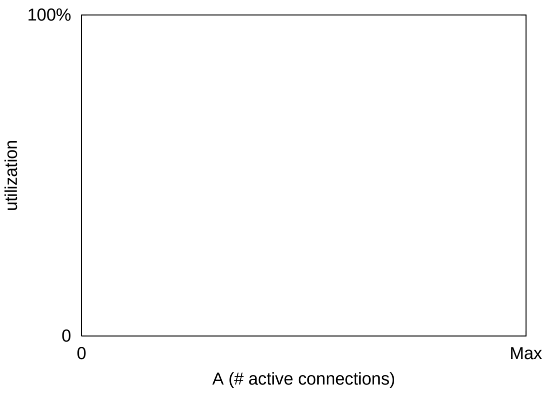

Statistical Multiplexing

C TCP Timeout Estimation: Now and Then

5. Suppose that TCP is consistently observing RTTs of 1.0 second, with a mean deviation of 0.1 second. Suddenly, the RTT jumps to 5.0 seconds with no deviation. Compare the behavior of the original RTO estimator with the Jacobson/Karels algorithm:
   (a) How many timeouts are encountered with each algorithm?
   (b) What is the largest value of RTO calcuated?

   Assume that α, the srtt EWMA parameter, $= \frac{1}{8}$, and that γ, the sdev EWMA parameter, $= \frac{1}{4}$.

   [This is based on problem 5.29 from Peterson & Davie, 3rd edition]

D Buffers, Losses, and TCP

Harry Bovik is given the responsibility of configuring the packet queuing component of a new router. The link speed of the router is 100 Mbit/s and he expects the average Internet round-trip time of connections through the router to be 80ms. Harry realizes that he needs to size the buffers appropriately.

You should assume the following:

- You're dealing with exactly one TCP connection.
- The source is a long-running TCP connection implementing additive-increase (increase window size by 1 packet after an entire window has been transmitted) and multiplicative-decrease (factor-of-two window reduction on congestion).
- The advertised window is always much larger than the congestion window.
- The loss recovery is perfect and has no impact on performance.
- The overhead due to headers can be ignored.

  Harry argues that because the average RTT is 80ms, the average one-way delay is 40ms. Therefore, the amount of buffering he needs for high link utilization is 100 Mbit/s * 40 ms or 500 KBytes.

6. Approximately what bandwidth will TCP achieve with this buffering?

   Note: There are two approaches to solving this problem. The easiest way to solve it is to assume that queueing delay does *not* increase the RTT. A more complete analysis will take queueing delay when the router buffers are over-full into account. The non-queueing-delay analysis will receive 80% of the points for this problem, so if you're stuck, it's much better to do the linear analysis than nothing, but we encourage you to do the full analysis. It's got some summations, but doesn't actually require computing an integral

E Riding a Sawtooth to Reno

In this problem, you will examine the behavior differences between the original TCP congestion control (Tahoe) and the improvements that were made in adding fast recovery to Reno.

First, assume that:

- The MSS is 1000 bytes
- The initial ssthresh is 8*MSS (for our convenience)
- The initial cwnd is 1 MSS
- Ignore the setup (SYN) and teardown exchanges. Just look at data packets.

Consider the transfer of a 51,000 byte file (assume it fits in exactly 51 packets)

7. Draw a "sequence number" vs time plot for these packets, assuming that they were all successfully received. Mark the RTT boundaries on the X axis, and for the start of each RTT, indicate the sender's congestion window. Indicate the point at which the connection switches from slow start to congestion avoidance. (There are no losses, so Reno and Tahoe behave the same). There's a figure for you to fill in on the next page to make it a bit simpler.

file transfer

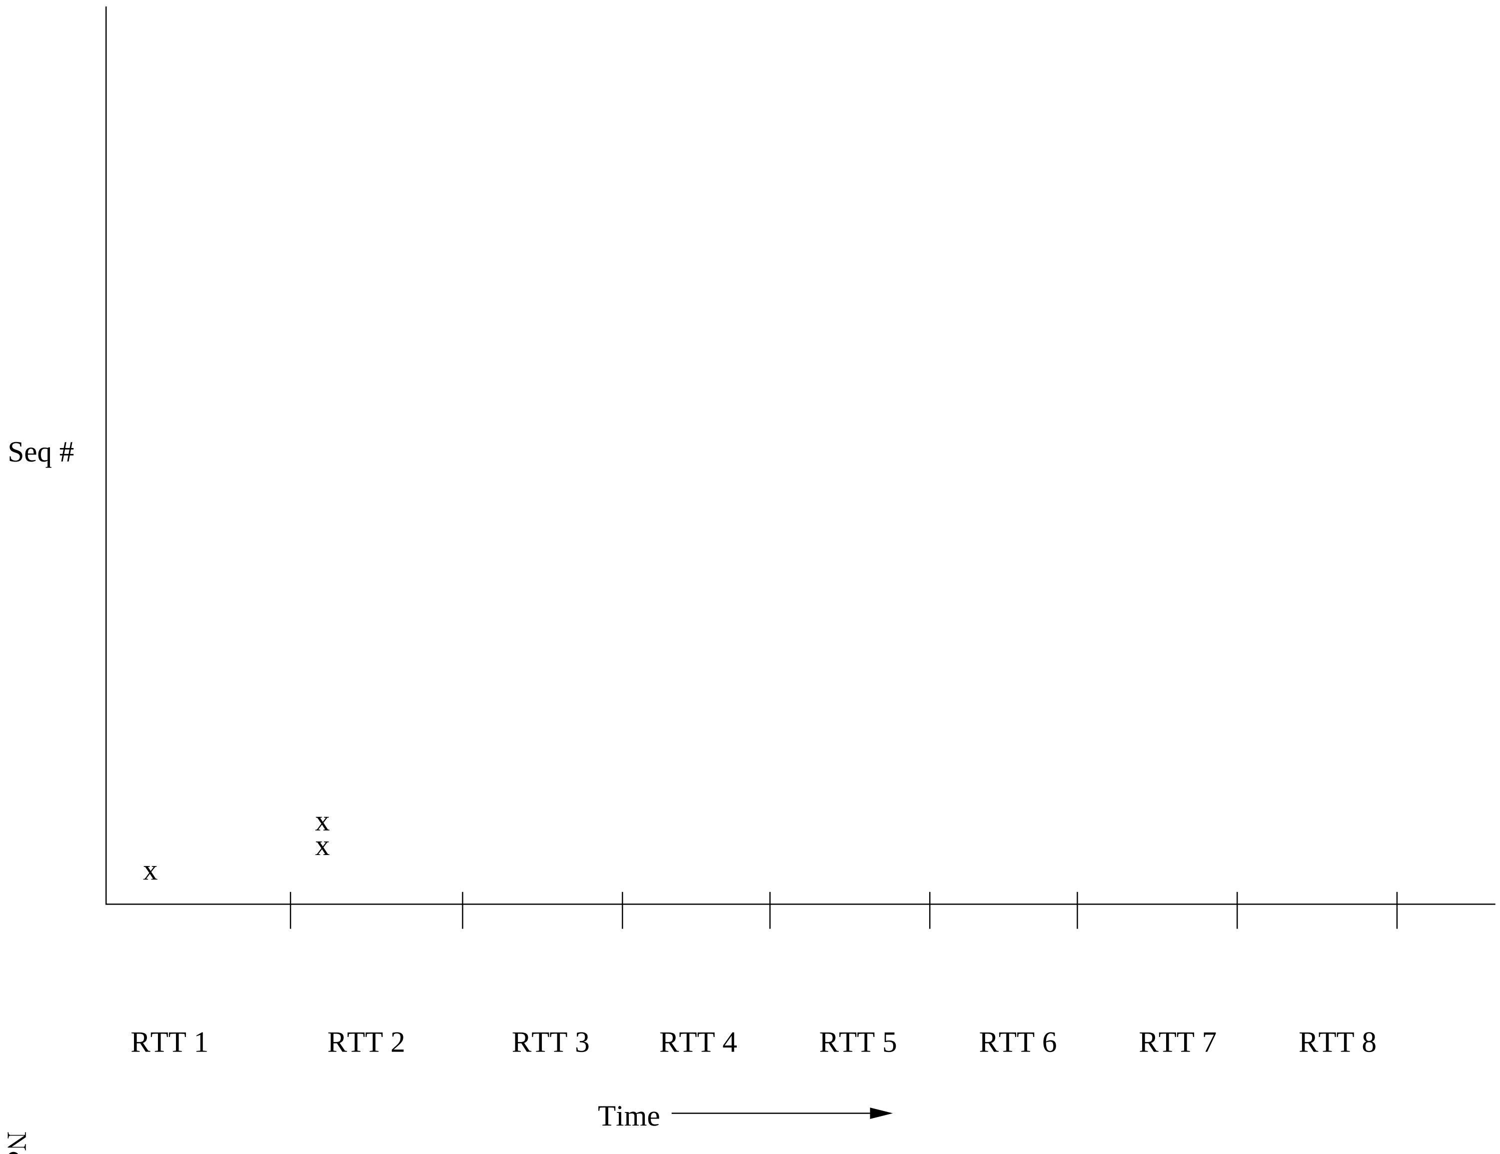

No loss

8. Now assume that there was a single loss of the 17th packet (seq 16001:17000). For the RTT in which the packet loss occurred and the following three RTTs, draw the sequence number trace that would happen under **Tahoe**. Then draw the sequence number trace that would happen with fast recovery as implemented in **Reno**.

9. Things on the network are really starting to get bad. This time, not only was the 17th packet lost, but so was the 23rd. Draw the sequence number trace that would happen with **Reno**. (Note: For these last two questions, you *only* need to draw the sequence number trace for the RTT in which the loss occurred and the following three RTTs. Don't redraw the whole thing!

F BGP

10. Consider the AS graph below in parts (a) and (d):

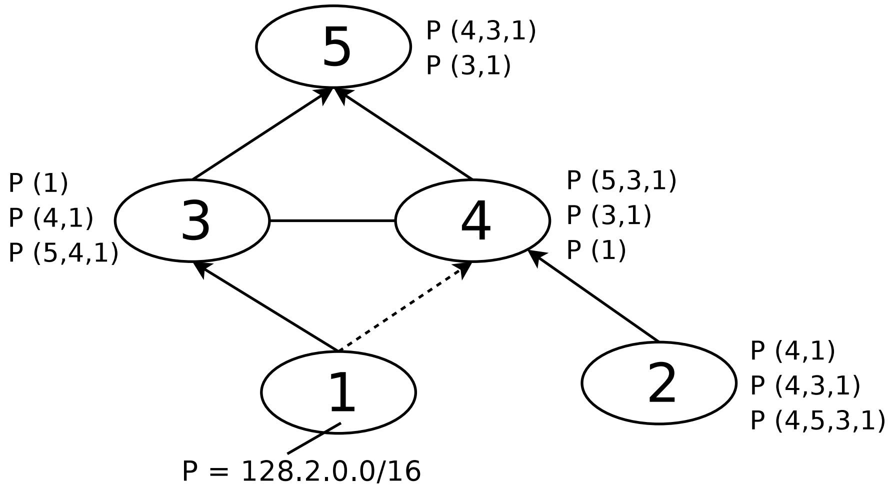

The vertices are inidividual ASes and edges are links between them.

(a) Now suppose the arrows represent customer-provider relationships where the customer points to its provider. An edge without arrows represents a link between peers. The dashed edge is a backup link. Suppose all ASes follow the export rules proposed by Gao and Rexford. List the routes in each AS's local-pref list shown above that can never be received by that AS.

(b) Suppose AS X thinks that AS Y drops too many packets. Using only BGP, is it possible for AS X to implement a policy stating that "traffic outbound from my AS should not cross Y?" Why or why not?

(c) Now suppose AS X thinks that AS Y generates a lot of illegal file sharing traffic. Using only BGP, is it possible for AS X to implement a policy stating that, "I don't want to carry traffic from Y to my customers?" Why or why not? Assume that AS X does not want to deny transit to traffic from any other AS.

(d) **Bonus question**: Shown next to each AS is its local-pref ranking of possible routes to P, a prefix in AS 1. Ignore the arrows for now and consider all edges equally valid. Is this AS graph gauranteed to converge? Why or why not? (If you read Gao and Rexford early because it was still on the syllabus, this question is for you!)