15-441 Computer Networks
Homework 3

Due: April 15, 2008, 1:30 PM

Lead TA: Xin Zhang (xzhang1@cs.cmu.edu)

April 1, 2008

A How does TCP update its window

At time t, a TCP connection has a congestion window of 4000 bytes. The maximum segment size used by the connection is 1000 bytes. What is the congestion window after it sends out 4 packets and receives acks for all of them? Suppose there is one ack per packet.

(a) If the connection is in slow-start?

(b) If the connection is in congestion avoidance (linear mode)?

B Label Switching

You are trying to debug a problem with your company's virtual circuit-based network. A diagram of the network is shown below. A, B, and C are hosts attached to the network. S1, S2, and S3 are switches configured to act as label swapping virtual circuit switches.

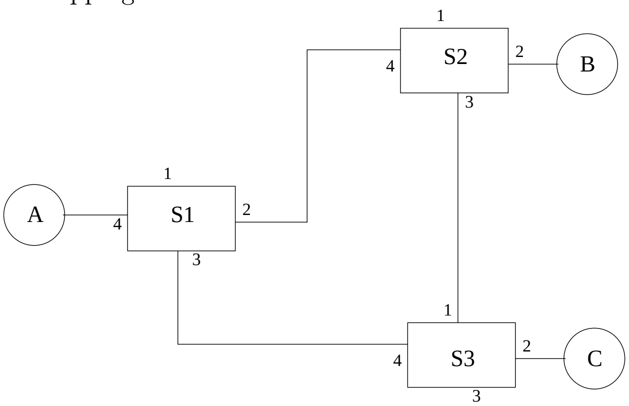

The label swapping tables for the switches are configured as follows. Some of the entries are stale and not actually in use right now.

Switch	Input Port	Input Label	Output Port	Output Label
S1	2	2	3	4
S1	4	2	3	1
S1	4	17	2	2
S2	2	19	4	2
S2	3	1	2	19
S2	3	2	2	15
S2	3	5	4	2
S2	4	2	2	1
S2	4	1	4	1
S3	2	1	1	2
S3	2	2	4	5
S3	4	1	1	1
S3	4	4	1	5

Write the sequence of (Switch, Input Port, Input Label) tuples and the destination node and label for each of these packets. We've given you the start node and starting label. The intermediate tuples should look like (S1, 1, 999) [e.g., switch S1, input port 1, label 999].

(a) Start node A, label 17. Switch tuples:

Dest node and final label:

(b) Start node A, label 2.
Switch tuples:

Dest node and final label:

(c) Start node C, label 1.
Switch tuples:

Dest node and final label:

C Why UDP vs TCP?

(a) Give **one** reason that DNS lookups are run over UDP rather than TCP:

(b) Give **one** reason that streaming multimedia is run over UDP rather than TCP:

D Pick the true choices about congestion collapse and backoff

Otto Pilot creates a new network for the 150 PC computers he mounted within his car. Each computer sends indepenent UDP query/response packets to the other computers in the car when it needs to know or do something. Requests are retried after a time out that is a fixed, small multiple of the typial response time. After running the OttoNet for a few days, Otto notices that network congestion occasionally causes a congestion collapse because too many packets are sent into the network, only to be dropped before reaching the eventual destination. These packets consume valuable resources.

Suppose each response or request can be fit into one packet. Which of the following techniques is likely to reduce the likelihood of a congestion collapse? (Circle ALL that apply)

A. Increase the size of the queue in each router from 4 packets to 8 packets. Suppose the timout value is appropriately adjusted accordingly to the queue length.

B. Use exponential backoff in the timeout mechanism while retrying queries.

C. If a query is not answered within a timeout interval, multiplicatively reduce the maximum rate at which the client application sends OttoNet query packets.

D. Use a TCP style flow control window (per session) at each receiver to prevent buffer overruns.

E Congestion Window

Consider the following plot of TCP window size as a function of time:

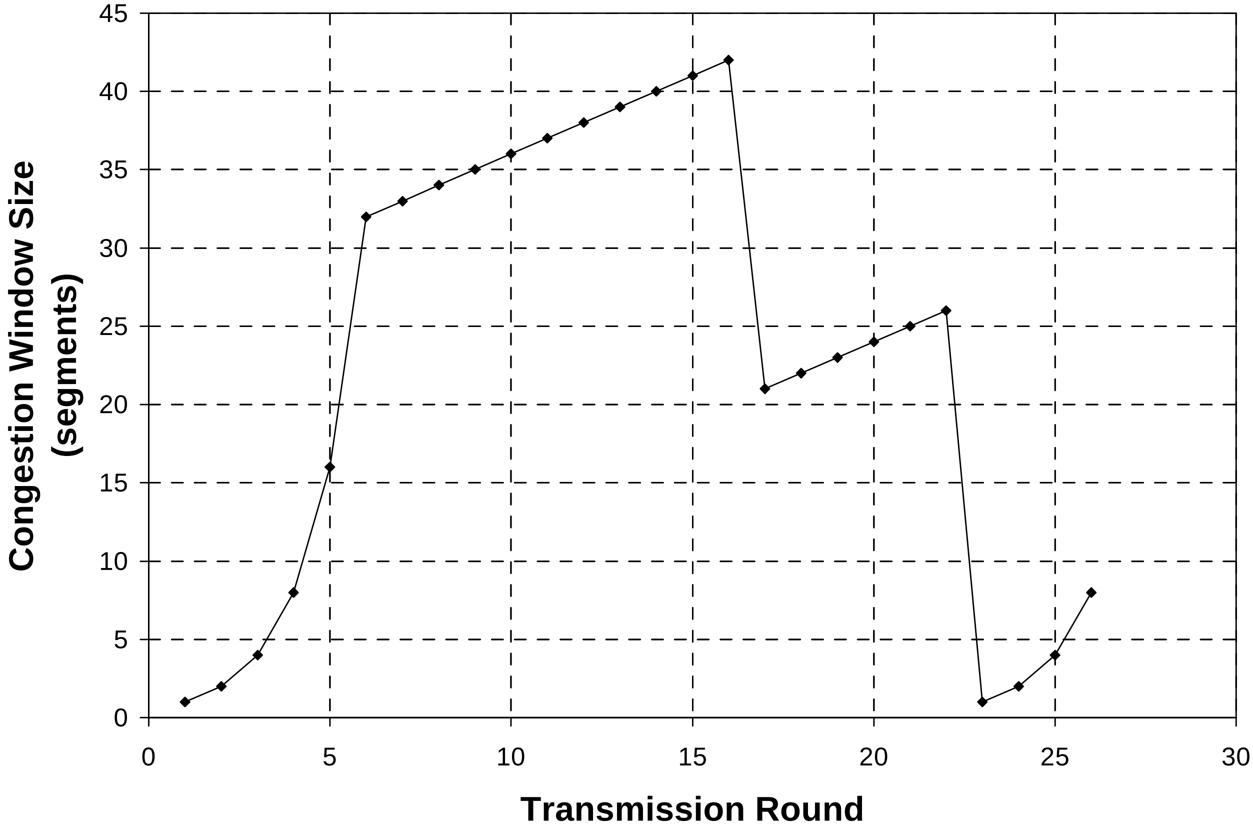

Assuming TCP Reno is the protocol experiencing the behavior shown above, answer the following questions.

(a) Identify the intervals of time when TCP slow start is operating.

(b) Identify the intervals of time when TCP congestion avoidance is operating (AIMD).

(c) After the 16th transmission round, is segment loss detected by a triple duplicate ACK or by a timeout?

(d) What is the initial value of ssthreshold at the first transmission round?

(e) What is the value of ssthreshold at the 18th transmission round?

(f) What is the value of ssthreshold at the 24th transmission round?

(g) During what transmission round is the 70th segment sent?

(h) Assuming a packet loss is detected after the 26th round by the receipt of a triple duplicate ACK, what will be the values of the congesion-window size and of ssthreshold?

F 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.

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