EE122 MIDTERM EXAM: 2003-10-13 (Fall 2003 Midterm 1)
Scott Shenker, Ion Stoica

This is a closed-book exam. This booklet has 10 numbered pages including the cover page, and contains 4 questions. You have 1.5 hours to complete the exam. Good luck.

Question - Max Score
1 - 18
2 - 20
3 - 20
4 - 42

Question 1 (18 pts)
Node A (IP address: 111.111.111.111) connects to router R (IP address: 222.222.222.222). The store-and-forward router R connects to node B (IP address: 123.123.123.123). All links have 40kbps bandwidth and have propagation delay 0.75ms. A sends packets, each of which is 50 bytes, to B. B sends an ACK to A each time it receives a packet. The size of the ACK packet is 5 bytes. The protocol is stop-and-go. (i.e. The sender sends a new packet only after the previous packet has been acked.)

You can assume that there is no loss during the transmission and the connection A to B is the only connection for the links. You can ignore the processing delay and queuing delay of a packet, but not the transmission time.

A sends a packet to B through R.

(4 pts) (a) What are the source IP address and destination IP address of the header of the packet when node A sends the packet to the router?

(4 pts) (b) What are the source IP address and destination IP address of the header of the packet when the router sends the packet to node B?

(5 pts) (c) What is the throughput of the flow from A to B?

(5 pts) (d) Are the terms bandwidth and throughput the same? If not, why not?

Question 2 (20 pts)
Consider two routers A,B that are connected by a 10 Mbps link. Suppose 4 TCP flows and a constant UDP flow of 5 Mbps traverse this link, and that this link is the bottleneck for all these flows. Also assume all the TCP flows have the same RTT and each of them can use up to 5 Mbps alone.

(5 pts) (a) What will be the average throughput of each of the TCP flows? What will be the average throughput of the UDP flow?

(5 pts) (b) Suppose that router A implements the round-robin scheduling discipline, and that the packet size of TCP-1 and TCP-2 is 500 bytes, the packet size of both TCP-3 and TCP-4 is 1500 bytes, and the packet size of the UDP flow is 1000 bytes. What is the average throughput of each flow?

(5 pts) (c) Suppose FQ is implemented at router A (FQ implements max-min fairness). What will be the average throughput of each of the TCP flows? What will be the average throughput of the UDP flow?

(5 pts) (d) Now, suppose that the maximum possible rates for each of the TCP flows are as follows: TCP-1 (1 Mbps), TCP-2 (2 Mbps), TCP-3 (2.5 Mbps), TCP-4 (3 Mbps). As before, the UDP flow has a constant rate of 5 Mbps. If the router A uses FQ, what is the average throughput of each of the TCP flows and the UDP flow?

Question 3 (20 pts)
Compute the routing tables for each node in the graph below using the distance vector method. Note that:

- The routing tables reflect the routing state in each router at the start of each time unit.
- Whenever an entry in the routing table of a node changes, that node propagates its routing table to each of its neighbors.
- At the end of each time unit, each node updates its routing table based on the routing tables received from its neighbors. A router updates its routing table instantaneously, i.e., you can neglect the time taken by the update operation.
- Propagation time along all links takes 1 time unit (ie. sending it at the start of the interval causes the information to reach the other end of the link at the end of the time unit, in time for updating), except for the link between nodes A and C, which takes 2 time units.

Give the routing tables at the start of each time unit until the algorithm converges. The first set of tables have been done for you.

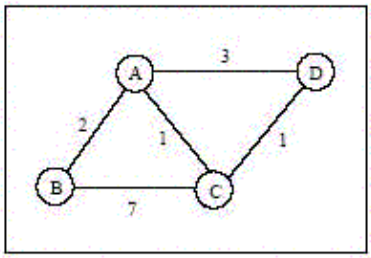

Time 0

Node A

Destination	Cost	Next Hop
B	2	B
C	1	C
D	3	D

Node B

Destination	Cost	Next Hop
A	2	A
C	7	C
D	inf	–

Node C

Destination	Cost	Next Hop
A	1	A
B	7	B
D	1	D

Node D

Destination	Cost	Next Hop
A	3	A
B	inf	–
C	1	C

Time 1

Node A

Destination	Cost	Next Hop
B	2	B
C	1	C
D	3	D

Node B

Destination	Cost	Next Hop
A	2	A
C	3	A
D	5	A

Node C

Destination	Cost	Next Hop
A	1	A
B	7	B
D	1	D

Node D

Destination	Cost	Next Hop
A	2	C
B	5	A
C	1	C

Time 2

Node A

Destination	Cost	Next Hop
B	2	B
C	1	C
D	2	C

Node B

Destination	Cost	Next Hop
A	2	A
C	3	A
D	5	A

Node C

Destination	Cost	Next Hop
A	1	A
B	3	A
D	1	D

Node D

Destination	Cost	Next Hop
A	2	C
B	5	A
C	1	C

Time 3

Node A

Destination	Cost	Next Hop
B	2	B
C	1	C
D	2	C

Node B

Destination	Cost	Next Hop
A	2	A
C	3	A
D	4	A

Node C

Destination	Cost	Next Hop
A	1	A
B	3	A
D	1	D

Node D

Destination	Cost	Next Hop
A	2	C
B	4	C
C	1	C

Question 4 (42 pts)

(5 pts) (a) What is the purpose of the *congestion window* (cwnd) in TCP?

(5 pts) (b) What is the *slow-start* phase in TCP? What is its purpose?

(5 pts) (c) What is the *congestion avoidance* phase in TCP? What is its purpose?

(12 pts) (d) Suppose a source S and a destination D are connected via a single link with round trip time RTT. Suppose S starts sending to D at time $t = 0$ using TCP Tahoe, and suppose the 14th segment sent is lost. Note that TCP's cwnd is initially set to 1. Assume TCP's RTO is equal to 3 RTTs (ie, TCP will detect loss when the RTO expires). Fill in the diagram below showing how the congestion window grows over time. Label which periods of time are *slow-start* and which are *congestion avoidance*.

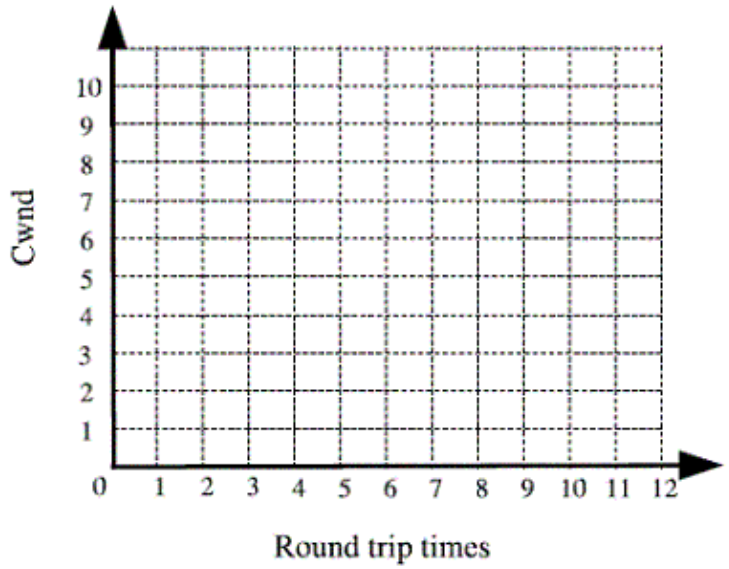

(10 pts) (e) Repeat part (d), but this time assume the 12th segment is lost, and assume TCP Reno is used. Note that TCP Reno also detects loss upon receiving 3 duplicate acks.

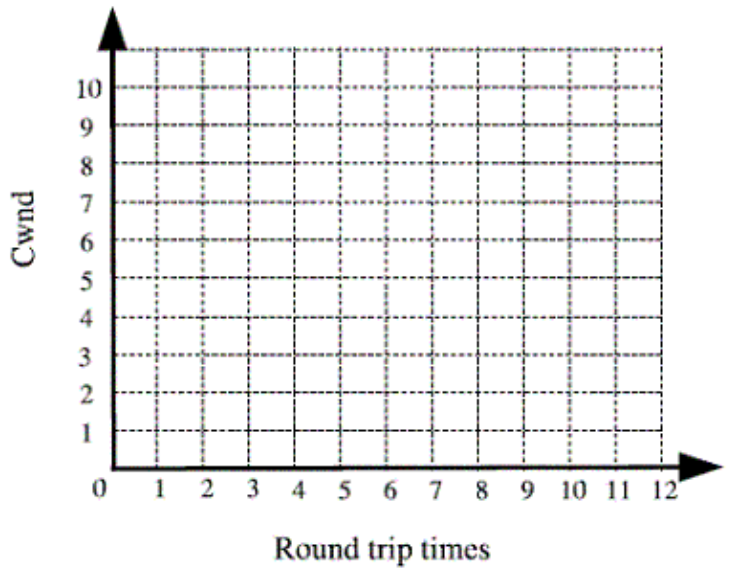

(5 pts) (f) In part (d), after 10 RTTs, how many segments has D received?