CS 349/449
Internet Protocols
Final Exam
Winter 2003
12/15/2003

Name:_____

Course:_____

Instructions:
1. You have 2 hours to finish
2. Question 9 is **only** for 449 students
3. Closed books, closed notes. Write all your answers using the pages on this exam (use back pages if needed)
- Calculators allowed
- This exam has 12 pages, including this cover page

Question 1 (10 pt)

(5 pt) Explain briefly how the Reservation Protocol (RSVP) works. In particular, what are the main messages and what do end-hosts and routers do upon receiving these messages?

RSVP is a signaling protocol that, in the context of Intserv, is used for admission control and resource reservation. There are two main control messages in RSVP:

1. Path message – this message is sent by the sender to the receiver and is used to (1) carry the sender's traffic specifications, and (2) set up the path. Upon receiving this message each node stores the previous node along the path.

2. Resv message – this message is sent by the receiver to sender and is used to perform admission control and resource reservation. Upon receiving such a message a node checks whether it has enough resources to satisfy the receiver requirements. If yes, it reserves resources; if not, it marks the message to indicate that the reservation is denied. In both cases the node forwards the message to the previous node (set up by the Path message) towards the source.

(5 pt) How does RSVP handle IP route changes? Give a simple example to illustrate this point.

RSVP is a soft-state protocol, i.e., routers maintain state only as long as it is periodically refreshed; when the state timeouts it is simply removed. To refresh state on routers along the path, the sender periodically sends Path messages to the receiver and the receiver periodically sends Resv messages to the sender. RSVP handles IP route changes as follows. When the route changes, the next Path message is forwarded along the new route and a new reverse path is set up. Then, when the receiver sends the next Resv message, the message travels the new route and reserves resources on that route. Eventually, the resources on the old route are released when they timeout.

Question 2 (10 Points)

    (a) **(5 pt)** Give two advantages of RED over Drop Tail. Explain.

RED avoids (1) TCP flow synchronization and (2) large bursts of losses. The main reason behind these advantages is the fact that RED starts to drop packets randomly before the buffer is full. This allows the router to spread losses over a larger time interval, and signal the congestion to end-hosts earlier.

    (b) **(5 pt)** Give one advantage of RED over Fair Queuing and one advantage of Fair Queuing over RED. Explain.

RED can be implemented more efficiently since, unlike Fair Queuing, it does not require to maintain per flow buffers and perform per-flow scheduling; RED maintains only one queue and performs FIFO scheduling. Fair Queuing can provide per-flow isolation. While with RED, a non-adaptive flow such as an UDP flow can hurt a TCP flow if it sends traffic at more than its fair share, with Fair Queuing that flow can hurt only itself.

Question 3 (10 points)

(a) **(5 pt)** Explain how the fair max-min share is computed.

Let C be the capacity of the congested link, S be set of all flows, n be the number of flows in S, and r_i be the arrival rate of flow i. Then we can use the following iterative algorithm to compute the max-min fair share:

(1) Compute f=C/n

(2) Let S(f) be the sub-set of flows from S consisting of all flows with arrival rates that are less or equal to f. If S(f) not empty:

$$C = C - \sum_{i \in S(f)} r_i;$$
$$n = n - size(S(f));$$
$$S = S - S(f)$$

go to step 1

Otherwise, stop and f is the max-min fair share; terminate

(b) **(5 pt)** Consider 10 flows with the arrival rates of 1, 2, ..., 10 Mbps that traverse a link of 50Mbps. Compute the fair share on this link. What is the fair share if the link capacity is 60 Mbps?

(1) Iteration 1: f = 50Mbps/10 = 5Mbps; S(f) = {1, 2, 3, 4,5}. Thus, C = 50 Mbps – 15Mbps = 35Mbps, n = 5, and S = {6, 7, 8, 9, 10}

(2) Iteration 2: f = 35Mbps/5 = 7Mbps; S(f)={6,7}, C =35Mbps –13Mbps = 22Mbps, n = 3, and S = {8, 9, 10}

(3) Iteration 3: f = 22Mbps/3 = 7.33Mbps; S(f) is empty, thus max-min fair share is 7.33Mbps

For C=60 Mbps note that the sum of all arrival rates is 1+2+...+10 = 55Mbps, which means that all flows will get how much they requested. In this case the fair rate is 10, the maximum arrival rate of all flows. (You can reach the same result using the iterative algorithm above).

Question 4 (15 pt)

(a) (5 pt) Describe the Assured service model.

The assured service provides a traffic profile (usually characterized by tokenbucket) to each customer. The customer is promised that as long as it sends assured traffic without exceeding its profile, it will experience a lower packet loss than using the best-effort service. The traffic profile is defined independent of the destinations.

(b) (10 pt) Describe how RIO (Red with In and Out) buffer management scheme works and how it can be used to implement the Assured service model.

RIO is an extended version of RED that uses two thresholds. When the queue size exceeds the first threshold, RIO starts to drop the best-effort packets. RIO starts to drop assured packets only when the queue size of the assured traffic exceeds the second threshold. Since the second threshold is always larger than the first one, RIO will start to drop assured packets only as a last resort, i.e., if dropping all the best-effort packets will not eliminate the congestion.
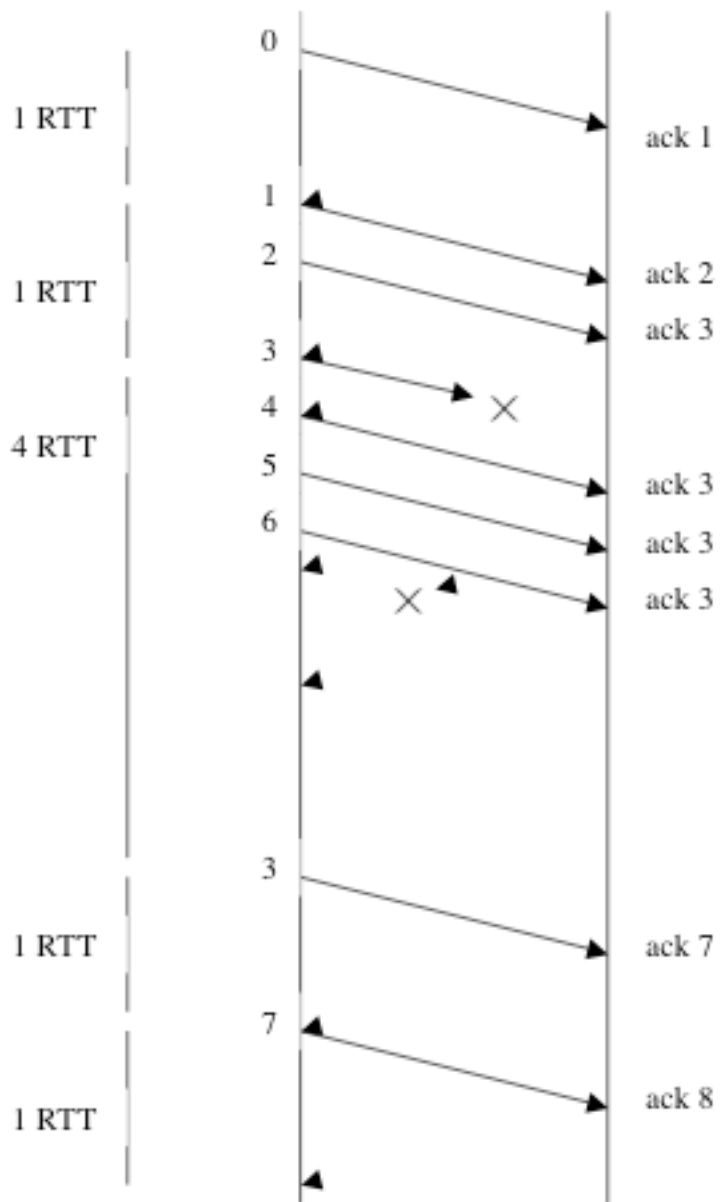
Question 5 (20 points)
Consider a source A that uses TCP to transfer a file of 8KB to destination B. Each packet has a payload of 1KB. Assume that only two packets are lost during the transfer: the 4th data packet and the acknowledgment of the 6^{th} data packet. Draw the time diagram of the entire file transfer. Assume that retransmission timeout (RTO) is 4*RTT. How many RTTs does it take to transfer the entire file?

Note: The transfer time represents the interval between the time the acknowledgement of the last packet was received by A and the time the first packet is transmitted by A. Ignore the connection setup, processing and transmission times. Fast retransmission and fast recovery are triggered after receiving 3 duplicate acknowledgements.

The answer is 8 RTTs:

Some people assumed that duplicate acknowledgements open the congestion window. TCP, as described in class, does not do this, but it is a reasonable thing for it to do, so we did not deduct points for this.

+3 points for getting the first 2RTT from 3 packets in slow start
+2 points for having the sender recognize that packet 3 and 5 are unacked
+2 points for only retransmitting packet 3
+3 for having correct RTO (in solution given above) or Fast Retransmit (in alternate solution)

1 RTT

0

ack 1

1

2

ack 2

1 RTT

ack 3

3

4

4 RTT

5

ack 3

6

ack 3

ack 3

3

1 RTT

ack 7

7

ack 8

1 RTT

Question 6 (10 Points)

Explain why Additive Increase Multiplicative Decrease **(AIMD)** divides resources fairly between two competing TCP flows sharing the same bottleneck link.

The graph shown in the slides

Question 7 (15 Points)
Suppose a very large web site wants a mechanism by which clients access whichever of multiple http servers is "closest" by some suitable measure.

(a) **(5 points)** Discuss developing a mechanism within HTTP for doing this

HTTP Redirect

(b) **(5 points)** Discuss developing a mechanism within DNS for doing this

The DNS server gives different replies based on where the client is coming from

(c) **(5 points)** Compare the two. Can either approach be made to work without upgrading the browser?

DNS does not require any upgrate. HTTP redirect is a standard but if it wasn't it would require an upgrade. Both answers are considered correct