COE 431 – Computer Networks

Welcome to Exam II
Thursday May 09, 2013

Instructor: Dr. Wissam F. Fawaz

Name: __Solution Key____________________

Student ID: _________________

<u>Instructions:</u>

1. This exam is **Closed Book**. Please do not forget to write your name and ID on the first page.
2. You have exactly **85 minutes** to complete the 6 required problems.
3. Read each problem carefully. If something appears ambiguous, please write your assumptions.
4. Do not get bogged-down on any one problem, you will have to work fast to complete this exam.
5. Put your answers in the space provided only. No other spaces will be graded or even looked at.

Good Luck!!

<u>Problem I:</u> Multiple choice questions (**20 minutes**) [16 Points]

1. After a loss event is detected by a **triple duplicate ACK**, TCP Reno
   a. Cuts its congestion window to 1 MSS.
   b. **Enters a fast recovery phase.**
   c. Enters a slow-start phase.
   d. Both (a) and (c)
2. Which of the following is correct about the **congestion control** service in TCP?
   a. The receiver advertises its buffer size
   b. The sender does not overflow the receiver's buffer by transmitting too many segments
   c. Both of the above
   d. **None of the above**
3. A TCP transmitter has received a segment with a **sequence number** equal to 80. This means that:
   a. **The receiver has received byte 80**
   b. The receiver has received all the bytes preceding byte 80
   c. The receiver can accept up to 80 bytes without overflow of its buffer
   d. The sender should send 80 bytes in the next segment
4. Consider two hosts that are interconnected by a channel with a transmission rate *R* of 1 Gbps. Assume that the round trip delay between these two end systems is approximately 30 ms. Suppose further that the two hosts communicate with each other via a stop-and-wait reliable data transfer protocol by exchanging packets with an average size *L* of 1000 bytes. Assuming for simplicity that the ACK packets that the receiver sends are extremely small, what is the effective throughput offered by the sender-to-receiver channel?
   a. 270 kbps
   b. 265 kbps
   c. 300 kbps
   d. **Other (provide your answer): 267 kbps**
5. Over a TCP connection, suppose host A sends two segments to host B, host B sends an acknowledgment for each segment, the first acknowledgement is lost, but the second acknowledgment arrives before the timer for the first segment expires. In this case,
   a. Host A will retransmit the second segment
   b. **Host A will retransmit neither segments**
   c. Host A will retransmit the first segment
   d. None of the above
6. Which of the following is true about TCP?
   a. The size of the TCP `RcvWindow` never changes throughout the duration of the connection
   b. Suppose Host A sends one segment with sequence number 38 and 4 bytes of data over a TCP connection to Host B. In this segment the acknowledgment number is necessarily 42.
   c. Suppose Host A is sending a large file to Host B over a TCP connection. If the sequence number for a segment of this connection is m, then the sequence number for the subsequent segment will necessarily be m+1.
   d. **Suppose Host A is sending Host B a large file over a TCP connection. The number of unacknowledged bytes that A sends cannot exceed the size of the receive buffer.**
7. Which of the following protocols does not run on top of UDP?
   a. DNS
   b. SNMP
   c. TFTP
   d. **None of the above**

8. Which of the following is false about POP3?
    a. It uses TCP as its transport protocol
    b. It guarantees the authentication of the user
    c. **It allows a user to create remote folders and assign messages to folders**
    d. None of the above
9. Which of the following parameters is used to calculate the timeout interval in TCP?
    a. Sample Round Trip Time (RTT)
    b. RTT variation
    c. **All of the above**
    d. None of the above
10. In a reliable data transfer protocol, duplicate ACK
    a. Is a technique to increase the probability of successful ACK transmission
    b. **Can be used instead of NAK packets**
    c. Is used to detect packet loss before timeout
    d. Both (b) and (c)
11. By replacing a stop-and-wait protocol with a pipelined protocol, a sender increases its link utilization by a factor of k. How many packets is the sender allowed to transmit without receiving acknowledgments?
    a. k/2
    b. 2k
    c. **K (correct)**
    d. None of the above
12. In a TCP connection, the timeout interval is 1 second long. What will be the new length of the timeout interval when the timer expires?
    a. 1.5 seconds
    b. **2 seconds**
    c. It depends on `SampleRTT` and `EstimatedRTT` values
    d. None of the above
13. In a Go-Back-N (GBN) protocol, if the last correctly received (and delivered to the upper layer) segment has a sequence number of 5 and the receiver next receives a segment with a sequence number 7, it does the following
    a. Buffers segment 7 and sends an ACK for segment 5
    b. Buffers segment 7 and sends an ACK for segment 7
    c. Discards segment 7 without sending any ACK
    d. **Discards segment 7 and sends an ACK for segment 5**
14. A user requests a Webpage that consists of some text and 3 images. The browser's cache is empty. For this page, the client's browser:
    a. Sends 1 http request message and receives 1 http response message
    b. Sends 3 http request messages and receives 3 http response messages
    c. Sends 1 http request message and receives 4 http response messages
    d. **None of the above**
15. Which of the following limits a TCP sender's sending rate?
    a. Receive window
    b. Congestion window
    c. **Both of the above**
    d. None of the above
16. Which of the following can be used to detect packet loss?
    a. Checksum
    b. Timer
    c. Sequence number
    d. **Both (b) and (c)**

<u>Problem II:</u> Comparing terminologies (**10 minutes**) [10 Points]

What is the difference between each of the following pairs of concepts?

1. Go-Back-N and Selective Repeat error recovery protocols

In a GoBackN protocol, the sender is allowed to transmit multiple packets without waiting for an acknowledgment, but is constrained to have no more than N unacknowledged packets in the pipeline. In GoBackN, the receiver discards out-of-order packets. If a timeout occurs, the sender resends all packets that have been previously sent but not yet been acknowledged. With selective repeat, the sender retransmits only those packets that were received in error at the receiver. Retransmission in this case requires that the receiver individually acknowledge correctly received packets whether or not they are in-order.

2. Stop-and-wait and pipelined reliable data transfer protocols

With stop-and-wait, the sender will not send a new piece of data until it is sure that the receiver has correctly received the current packet. In pipelined data transfer, the sender is allowed to transmit multiple packets without having to wait for an acknowledgment.

3. Client-server file distribution and peer-to-peer file distribution

In client-server file distribution, the server must send a copy of the file to each of the peers. In p2p file distribution, each peer can redistribute any portion of the file that it has received to any other peers, thereby assisting the server in the distribution process.

4. TCP sequence number and TCP acknowledgment number

TCP sequence number is the number of the first byte in a TCP segment. The acknowledgment number is the sequence number of the next byte the receiver is expecting from the sender.

5. Flow control and congestion control

Congestion control is related to network carriage capacity while flow control is related to the receiver’s capacity.

<u>Problem III:</u> Network delays (**15 minutes**) [20 Points]

In modern packet-switched networks the source hosts segments long application layer messages (for example, an image or a music file) into smaller packets and sends the packets into the network. The receiver then reassembles the packet back into the original message. We refer to this process as message segmentation. The figure given below illustrates the end-to-end transport of a message with and without message segmentation. Consider a message that is $8 * 10^6$ bits long that is to be sent from source to destination. Suppose each link has a maximum capacity of 2 Mbps. Ignore propagation, queueing and processing delays.

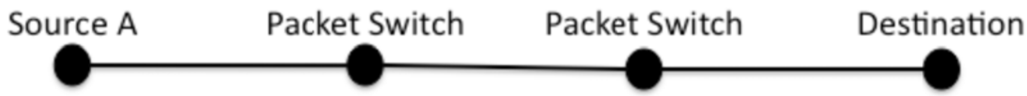

1. Consider sending the message from source to destination without message segmentation. How long does it take to move the message from the source host to the first packet switch? Keeping in mind that each switch uses store-and-forward packet switching, what is the total time to move the message from source host to destination host?

Time to send message from source host to first packet switch = ——— .

With store-and-forward switching, the total time to move message from source host to destination host = 4 sec x 3 hops = 12 sec.

2. Now suppose that the message is segmented into 4000 packets, with each packet being 2000 bits long. How long does it take to move the first packet from the source to the first switch? When the first packet is being sent from the first switch to the second switch, the second packet is being sent from the source to the first switch. At what time will the second packet be fully received at the first switch?

Time to send 1^{st} packet from source host to first packet switch = ——— .

Time at which 2^{nd} packet is received at the first switch = time at which 1^{st} packet is received at second switch = 2 x 1 msec = 2 msec.

3. How long does it take to move the entire file from source host to destination host when message segmentation is used? Compare this result with your answer in part (a) and comment.

Time at which 1^{st} packet is received at the destination host = 1 msec x 3 hops = 3 msec. After this, every 1 msec one packet will be received; thus time at which last packet is received = 3 msec + 3999 x 1 msec = 4.002 sec. It can be seen that delay in using message segmentation is significantly less.

<u>Problem IV:</u> P2P file distribution (**15 minutes**) [20 Points]

Consider distributing a file of F=15 Gbits to N peers. The server has an upload rate of u_s = 30 Mbps, and each peer has a download rate of d_i = 2 Mbps and an upload rate of u. Give the minimum distribution time for each of the combinations of N and u listed in the tables below for both client-server distribution and P2P distribution. Show your work.

<u>Client-server file distribution</u>:

Client-server file distribution	N=10	N=100	N=1000
u=300 kbps	7680	51200	512000
u=700 kbps	7680	51200	512000
u=2 Mbps	7680	51200	512000

<u>P2P file distribution</u>:

P2P file distribution	N=10	N=100	N=1000
u=300 kbps	7680	25904	47559
u=700 kbps	7680	15616	21525
u=2 Mbps	7680	7680	7680

Problem V: GBN, SR, and TCP (**15 minutes**) [18 Points]

Compare GBN, SR, and TCP (no delayed ACK). Assume that the timeout values for all three protocols are sufficiently long such that 5 consecutive data segments and their corresponding ACKs can be received (if not lost in the channel) by the receiving host (Host B) and the sending host (Host A) respectively. Suppose Host A sends 5 data segments to Host B, and the second segment (sent from A) is lost. In the end, all 5 data segments have been correctly received by Host B.

1. How many segments has Host A sent in total and how many ACKs has Host B sent in total? What are their sequence numbers? Answer this question for all three protocols.

GoBackN:

A sends 9 segments in total. They are initially sent as segments 1, 2, 3, 4, 5 and later resent as 2, 3, 4, and 5. B sends 8 ACKs. They are 4 ACKs with ACK number 1, and 4 ACKs with ACK numbers 2, 3, 4, and 5.

Selective Repeat:

A sends 6 segments in total. They are initially sent as segments 1, 2, 3, 4, 5 and later resent as 2. B sends 5 ACKs. They are 4 ACKs with ACK numbers 1, 3, 4, 5, and there is one ACK with ACK number 2.

TCP:

A sends 6 segments in total. They are initially sent as segments 1, 2, 3, 4, 5 and later resent as 2. B sends 5 ACKs. They are 4 ACKs with ACK number 2. There is one ACK with sequence number 6. Note that TCP always sends an ACK with expected sequence number.

<u>Problem VI:</u> TCP congestion control (**10 minutes**) [16 Points]

Recall the macroscopic description of TCP throughput. In the period of time from when the connection's rate varies from **W / (2 x RTT)** to **W / RTT**, only one packet is lost (at the very end of the period).

a. Show that the loss rate (fraction of packet lost) is equal to:

$$L = \frac{1}{\frac{3}{8}W^2 + \frac{3}{4}W}$$

The loss rate, L, is the ratio of the number of packets lost over the number of packets sent. In a cycle, 1 packet is lost. The number of packets sent in a cycle is

$$\frac{W}{2} + \left(\frac{W}{2}+1\right) + \cdots + W = \sum_{n=0}^{W/2}\left(\frac{W}{2}+n\right)$$

$$= \left(\frac{W}{2}+1\right)\frac{W}{2} + \sum_{n=0}^{W/2} n$$

$$= \left(\frac{W}{2}+1\right)\frac{W}{2} + \frac{W/2(W/2+1)}{2}$$

$$= \frac{W^2}{4} + \frac{W}{2} + \frac{W^2}{8} + \frac{W}{4}$$

$$= \frac{3}{8}W^2 + \frac{3}{4}W$$

Thus the loss rate is

$$L = \frac{1}{\frac{3}{8}W^2 + \frac{3}{4}W}$$

b. Use the result above to show that if a connection has a loss rate of L, then its average rate is approximately given by:

$$rate = \frac{1.22 \cdot MSS}{RTT \cdot \sqrt{L}}$$

For W large, $\frac{3}{8}W^2 >> \frac{3}{4}W$. Thus $L \approx 8/3W^2$ or $W \approx \sqrt{\frac{8}{3L}}$. From the text, we therefore have

$$\text{average throughput } = \frac{3}{4}\sqrt{\frac{8}{3L}} \cdot \frac{MSS}{RTT}$$