

Name (Print) _____

Email Address _____

ECE4110, Internetwork Programming, QUIZ 2 - PRACTICE Spring 2006

Prof. John A. Copeland
Tel.: 404-894-5177
E-Mail: copeland@ece.gatech.edu

Practice for April 11, 2006

RULES.

- i This quiz is **not** open book. One original sheet of hand-written notes may be used. Calculators are ok.
- ii Answer all questions and show all work to receive full credit. Use back of sheets only if necessary.
- iii All sub-questions have the same weight.
- iv Please do not ask the proctors any questions during the exam about exam questions. Part of the test is understanding the question, as written, without supplemental information. If you feel additional data is needed to solve the problem, make (and state) an assumption and then work the problem.
- v. This is a time-limited test. All papers must be turned in 50 minutes after the start. If you find you are taking more than 10 minutes on a particular problem, move on and come back to that problem after finishing the others. The Georgia Honor Code applies (see last page). **The actual test will be about 60% as long as this.**

Question 1 – TCP with no Congestion.

A 100 Mbps network connection has a 100 ms round trip time (RTT). A server begins sending maximum-size packets (after the initial SYN-ACK) using Slow Start. The client ACK's every packet (col. 2) or every-other (col.3), and has a receiver-window which holds 16 maximum-size packets. Show how many packets are sent in each RTT period below:

Time Period Start (ms)	No. Segments, when all ACKed	No. Segments, when 1/2 ACKed
0	1	1
100	2	2
200	4	3
300	8	3
400	16	6
500	16	9
600	16	12
700	16	16

___ 400 or 700 ms ___ How long does it take the server to ramp up to maximum transfer rate (ms).

___ 16,000 / 0.1s ___ If the maximum packet size is 1000 bytes, what is the maximum transfer rate (bytes/second).

How and when is the receiver's maximum segment size (MSS) sent to the sent to the other host?

How? ___ **TCP option** ___ When ___ **in SYN or SYN-ACK** ___

How and when is the receiver's TCP window (WIN) sent to the sent to the other host? ?

How? ___ **Field in TCP header** ___ When ___ **with every segment** ___

The last WIN received was 8000, the last ACK number was 12000, and the last TCP segment sent had a SEQ number of 14000. How many more bytes can be sent? ___ $8000 - (14000 - 12000) = 6000$ ___

Once those bytes are sent, name three conditions that will allow additional bytes to be sent:

- a. ___ Larger WIN received ___
- b. ___ larger ACK number received ___
- c. ___ "Persist" timer expires ___

A TCP host on an Ethernet LAN receives an MSS of 1460 bytes. It sends IP headers with no options (20 bytes) and TCP headers with 32 bytes (12 bytes for options). What is the maximum number of data bytes (SMSS) it can send per TCP segment? ___ (Ethernet max, 1500) - 20 - 32 = **1448 bytes** ___

Question 3 – Bandwidth-Delay Product

- a. If the round-trip time (RTT) for a dedicated 16 Mbit/s connection to California is 50 ms, what is the bandwidth-delay product of the link? 16,000,000 bits * 0.050 s = 800,000 (bits). [= 100,000 bytes]
- b. What is the minimum TCP window size that will allow this channel to be fully utilized? 100 kbytes (bytes)
- c. What will the maximum throughput be if the window is half the value in b. above. 8 Mbit/s (bits/s)
- d. What will the maximum throughput be if the window is twice the value in b. above. 16 Mbit/s (bits/s)

Question 4 – TCP Response to Congestion - Lost and Out-of-Order Packets

Once a TCP segment times out, a variable collision window-size is *cwnd* is used. For consistency, we consider *cwnd* to govern the number of segments sent during the initial Slow Start mode. SMSS (sender maximum segment size) was called *segsz* in the class slides, MSS in some of the books. Conditions at the sender may require SMSS to be smaller than the MSS value received for the other TCP host.

What places an upper limit on the maximum number of outstanding (sent, but not ACKed) bytes besides *cwnd*?
win [value received in last packet from receiver]

After a time-out, the value of *cwnd* drops to one times the maximum segment size (SMSS or *segsz*).

After a time-out, the value of *ssthresh* drops to 1/2 times the previous value of *cwnd* (**but at least 2 SMSS**).

cwnd then grows in slow-start mode until it reaches *ssthresh*, which is reduced to 1/2 times the value of min(cwnd, win) but at least 2 SMSS before the time-out.

After *cwnd* reaches *ssthresh*, *cwnd* grows linearly by adding SMSS bytes for approximately every Y number of maximum segments sent (ignore the incorrect SMSS/8 term in the slides and book).

What is this linear mode (of *cwnd* increase) called? congestion avoidance

For the case where the receiver ACKs every packet, $Y = \frac{cwnd}{SMSS}$ (in terms of *cwnd*, SMSS)

For the case where the receiver ACKs every other packet, $Y = 2 \frac{cwnd}{SMSS}$ (ditto)

What is the maximum value of *cwnd*? moot, since win used if smaller (or is it moot, see TCP Illustrated, p. 310, rule 1).

Does *cwnd* or *win* (from the receiver) govern the bytes (by sequence number) that can be sent? the smaller value

Question 5 – TCP Fast Retransmission and Fast Recovery

After a TCP sender receives 3 duplicate ACKs it retransmits the first missing segment.

How many additional bytes (at a minimum) can be sent after each additional duplicate ACK is received? SMSS
 [because *cwnd* increases by SMSS, and the ACK number does not change]

How much does *cwnd* increase after each additional duplicate ACK is received? SMSS

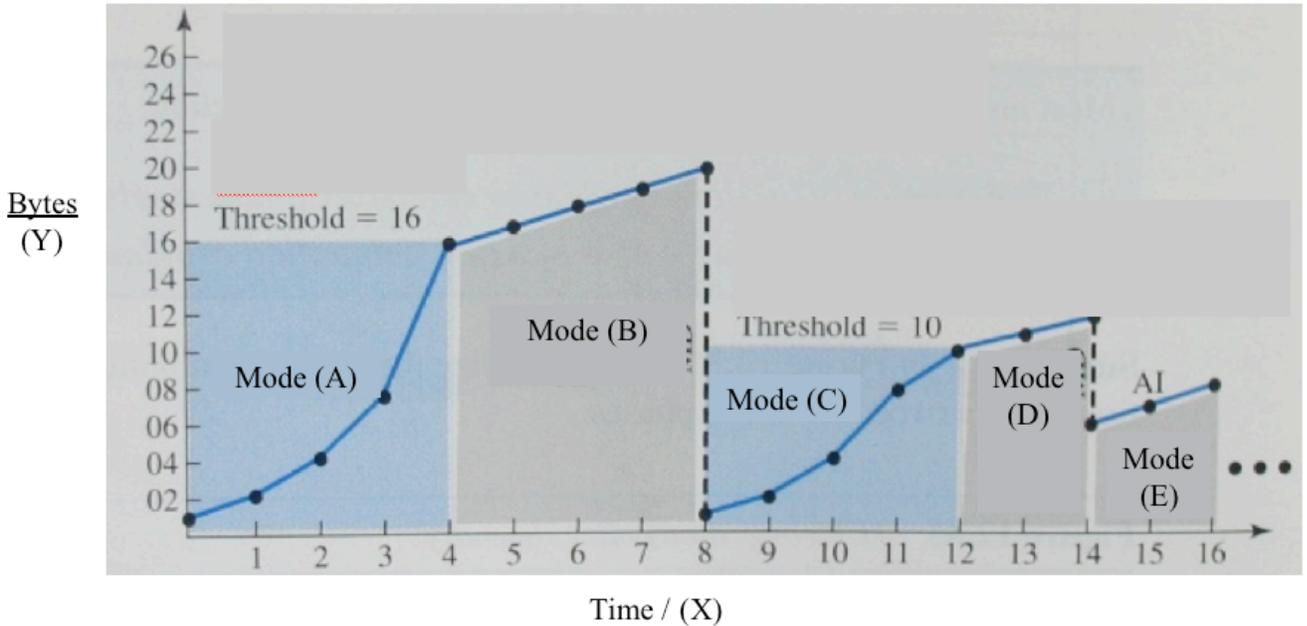
When does Fast Recovery operation start? when 1st non-duplicate ACK value received

When FR starts, *cwnd* is set equal to 1/2 x *cwnd-before-dup-ACK* + 1 SMSS.

In Fast Recovery, how many SMSS bytes are added each RTT if all segments are ACKed (do not use SMSS/8 term):
 if all segments are ACKed 1 if every other segment is ACKed? 1/2

Question 6– Window Size versus Time

SMSS =1420 bytes, MSS=1420 bytes, RTO = 50 ms, RTT = 30 ms.



What is the Time scale factor, X ? RTT
 What is the Bytes scale factor, Y ? SMSS

Identify the cwnd growth modes (Slow Start = SS, Collision Avoidance = CA, Fast Recovery=FR) ?:

- Mode A SS [not start-up since ssthresh is not 65,535 or infinite]
- Mode B CA
- Mode C SS
- Mode D CA
- Mode E FR

What causes the mode changes (Start Up = SU, Fast Retransmit = FT, Retransmit Time Out = RTO, ssthresh reached = SST) ?

- Change 0-A RTO [not start-up since ssthresh is not 65,535 or infinite]
- Change A-B SST
- Change B-C RTO
- Change C-D SST
- Change D-E FR

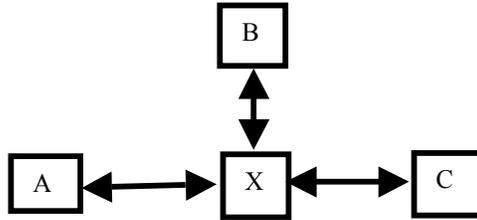
Question 7– Running Average for Calculating the Retransmit Time Out (use back of page for calculations)

New Round-Trip Time	A=Average RTT Alpha = 1/8	Deviation	D = Average Deviation Beta = 1/4	RTO = A + 4D
40	30	10	4	26
50	$30 + (50-30)/8 = \mathbf{33}$	$50 - 30 = \mathbf{20}$	$4 + (20 -4)/4 = \mathbf{8}$	$33+4*8 = \mathbf{65}$
30	$33 + (30-33)/8 = \mathbf{33}$	$30 - 33 = \mathbf{-3}$	$8 + (-3 -8)/4 = \mathbf{7}$	$33 + 4*7 = \mathbf{61}$
45	$31 + (45-33)/8 = \mathbf{35}$	$(45 - 33) = \mathbf{14}$	$7 + (12 -9)/4 = \mathbf{8}$	$35 + 4*8 = \mathbf{67}$

(ignore the 500 ms tick rounding, round results up to 1 ms) (On the quiz, only show the answers, in bold above)

Question 8– Routing, Distance Vector, RIP

The routing tables for RIP routers A, B, C are shown.



Router A Table		
Prefix	Distance	Port
128.230.	2	X
130.207.	5	N
209.196.	6	X
24.56.	9	X

Router B Table		
Prefix	Distance	Port
128.230.	2	X
130.207.	7	X
209.196.	4	M
24.56.	9	X

Router C Table		
Prefix	Distance	Port
128.230.	2	X
130.207.	7	X
209.196.	6	X
24.56.	7	P

Construct the Routing Table for Router X

Router X Table		
Prefix	Distance	Port
128.230.	1	-
130.207.	6	A
209.196.	5	B
24.56.	8	C

Using Split Horizon and Poison Reverse, construct the Updates sent from Router X to A,B,C.

Update X to A Table		
Prefix	Distance	Port
128.230.	1	-
130.207.	16	A
209.196.	5	B
24.56.	8	C

Update X to B Table		
Prefix	Distance	Port
128.230.	1	-
130.207.	6	A
209.196.	16	B
24.56.	8	C

Update X to C Table		
Prefix	Distance	Port
128.230.	1	-
130.207.	6	A
209.196.	5	B
24.56.	16	C

What 2 improvements were added to RIPv1 to make RIP v2

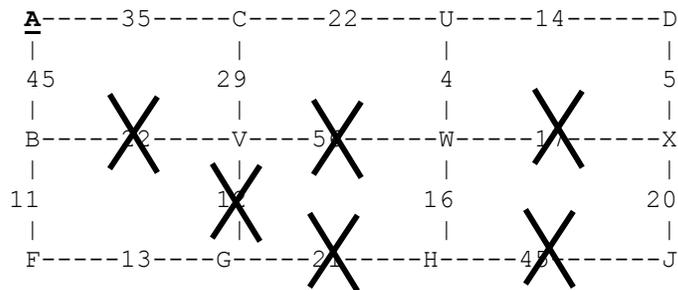
- a. _____ **Authentication**, to improve security _____
- b. _____ **Support for CIDR and VLSM** _____

Question 9– Routing, Link State, OSPF

Every router (A, B, ..., J) has advertised the costs (delays) to all the other nodes. For example, node V broadcasts the following route advertisement message: "V, B:29, G:12, U:22, W:26" (In practice, only link states that have changed since the last broadcast are included in the message.)

Based on all the advertisement messages, the network topology and link costs can be mapped. The letters below represent the nodes (routers) on the network. The numbers represent costs (delay-times) on the links between them.

These nodes are all routers (no networks) and for simplicity routes have the same cost in both directions.

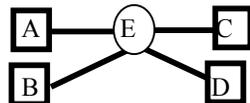


Using Dijkstra's technique, calculate the cost values of routes from node "A" to the following.

A-[0], F-[56], C-[35], G-[69], J-[96]

Cross out (X) the links that are not on the sink tree for node A.

Routers A,B,C,D are connected to Ethernet E. Show the way this would be diagrammed for an OSPF network diagram?



How are the link-states from a network node advertised? By the "designated router" What is the costs? 0

Question 9– Routing, Path Vector, BGP

The Internet is composed of interconnected Autonomous Systems (AS) that are linked by External BGP routers.

What are the three types of AS, and their defining characteristic? 1. Stub, has hosts and only one external connection 2. Multi-homed, has hosts and more than one external connection 3. Transit, has no hosts, more than one connection

A BGP routing table consists of what two items for every destination ? 1. Exit Port 2. Full path to destination

Why is it dangerous to have two border routers which both have the default route set to the other? packets could loop back and forth

What information do Autonomous Boundary Routers advertise to their exterior neighbors? costs to internal destinations

To reduce the size of routing tables, network prefixes are aggregated as much as possible.

Question 10 - Protocols for Router Updates

All of these protocols use IP for updates. Indicate multicast IP by "Multi.-" before the transport layer protocols they use.

RIP udp 520 OSPF multi- ospf BGP tcp

Question 10 - DNS, HTTP, FTP, SMTP, SPAM, SSH questions - see first and last slides of slide set 16 for topics.