EECS489 Computer Networks,
Midterm (Winter 2007)
SOLUTION

Instructions: You are allowed to use one sheet of notes (front and back). This exam is closed book, no computers are allowed. You can use a calculator. Read the entire exam through before you begin working. Work on those problems you find easiest first. Read each question carefully, and note all that is required of you. Please keep your answers clear and concise, and state all of your assumptions carefully.
Please write out the details of how you reach your final answer in order to get partial credit.

You are to abide by the University of Michigan/Engineering honor code.
Please sign below to signify that you have kept the honor code pledge.
Honor code pledge: I have neither given nor received aid on this exam.

Name: ____

Signature: ____

Unique name: ____

Problem 1: Router Queue Management – 20 points

[4 pts] (1)
Explain why there can be packet drops in packet-switched networks, but usually no packet loss for circuit-switched networks.

Answer: Packet-switched networks require no reservation of resources, so packets can be sent into the network at a faster rate than the network can process. In circuit-switched networks, a circuit is established for each communication session, guaranteeing resource needed for the session.
The router architecture consists mainly of the input ports, the switching fabric, the output ports, and the routing processor. Given n input ports each with K Gbps link speed. What is the minimum switching fabric speed to avoid packet loss at input ports?

Answer: $n \times K$.
Assuming for a given router, the switching fabric is fast enough to process packets from all input ports, does this guarantee no packet loss at this router? Why or why not?

Answer: No, it doesn't. Packet drop can still occur at the output ports due to slow output link speed.
Give an example of a packet once received by a router would require the routing processor to do some additional work.

Answer: IP option processing, TTL TIME EXCEEDED replies (from traceroute packets).
[12 pts] (2) Consider a system of four queues being serviced according to a WFQ (Weighted Fair Queuing) scheduling policy. The weights given to the four queues (A, B, C, D) are 3, 2, 4, and 1 respectively. They are being serviced by a processor at the rate of 10 Mbps.
The table below gives a list of different input traffic rates (in Mbps) at the four input queues. Fill in the resultant output rates for each of the four queues. The total bandwidth is 10. [Each row 2 points]

If TDM (time division multiplexing) is used instead, where each queue obtains a fixed share of the bandwidth based on the same weights as WFQ, which of the above row(s) will not suffer from resource waste.

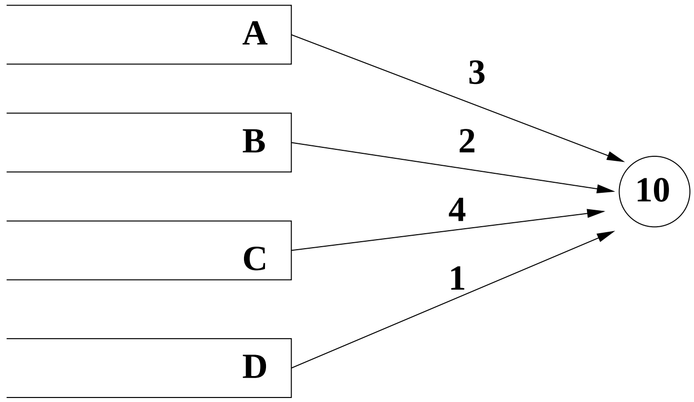

INPUT RATES				OUTPUT RATES			
A	B	C	D	A	B	C	D
2	2	1	1	2	2	1	1
15	10	10	10	3	2	4	1
6	6	1	2	4.5	3	1	1.5
8	0	0	8	7.5	0	0	2.5
5	1	3	5	4.5	1	3	1.5

(No resource waste is defined as whenever there is a demand for the resource, i.e., bandwidth, the available resource is always used to meet the demand.)

Answer: The first two rows: when all input rates are all smaller than its fair share, or all larger than its fair share.
Generalize under which circumstances TDM and WFQ are equivalent.
Answer: see above.

[4 pts] Besides WFQ and TDM, we introduce another type of queuing discipline called strict priority queuing (SPQ) between two queues: the queue with higher priority is always scheduled before the queue with lower priority. (Obviously, this can lead to starvation!)

There are four classes of traffic: U, V, W, X. Traffic class U is the control traffic (e.g., routing data). Traffic classes V, W, and X are data traffic (e.g., HTTP, SMTP). Class V should get 5 times the amount of bandwidth of class X. Class W should get 4 times the amount of bandwidth of class X. Given these requirements, assign each traffic class into one queue and draw the diagram to describe the queuing discipline. (Hint: you may need a two level scheduling).
Answer: Two level scheduling: first level involves WFQ for V,W,X using weights of 5:4:1 respectively, the second level is SPQ with U having strictly higher priority than other traffic.

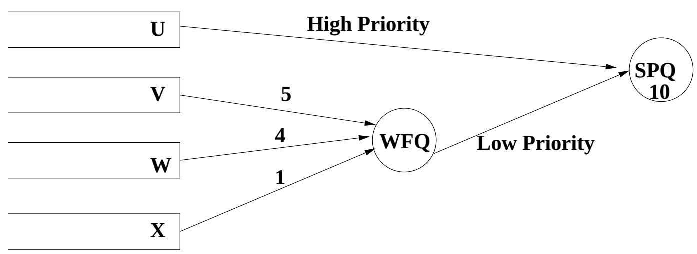

Please also fill in the table below based on your diagram.

INPUT RATES				OUTPUT RATES			
U	V	W	X	U	V	W	X
2	6	2	1	2	5	2	1
1	6	10	6	1	4.5	3.6	0.9
2	1	1	1	2	1	1	1

Problem 2: Packet switching – 20 points

[10 pts] (a) We would like to transfer 20KB file across a network from node A to node F. Packets have a length of 1 KB (neglecting header). The path from node A to node F passes through 5 links, and 4 intermediate nodes. Each of the links is a 10km optical fiber with a rate of 10 Mbps. The 4 intermediate nodes are store-and-forward devices and each intermediate node must perform a 50 μsec routing table look up after receiving a packet before it can begin sending it on the outgoing link. How long does it take to send the entire file across the network? (Describe individual delay components.) (Assume the propagation speed through optical fiber is $2 \times 10^8 m/sec$).

Answer:

Number of packets to transfer is $20KB/1KB = 20$ packets.

The per packet transmission time is: $TRANSP = \frac{8\times 2^{10} bits/packet}{1\times 10^7 bits/sec} = 819.2\mu$sec per packet.

The per link propagation delay is $PROP = \frac{10^4 m}{2\times 10^8 m/sec} = 50\mu$ sec

The total transfer time is $20TRANSP + 4TRANSP + 4PROC + 5PROP = 20.11ms$.

[10pts] (b) [Please ignore all the assumptions in part (a)] Suppose we would like to transfer a file of K bits from node A to node C. The path from node A to node C passes through two links and one intermediate node, B, which is a store-and-forward device. The two links are of rate R bps. The packets contain P bits of data and a 6 byte header. What value of P minimizes the time it takes to transfer the file from A to C? (Hint:to find the value that minimizes a given equation, take the first derivative with respect to that value and solve when it is set to zero).

Answer:

Number of packets to transmit the file: $\frac{Kbits/file}{Pbits/packet} = \frac{K}{P}$ packets.

(You can ignore the header size)

The transmission time of a packet on each link accounting for the 48 bit header is: $TRANSP = \frac{(P+48)bits/packet}{Rbits/sec} = \frac{P+48}{R}$sec per packet.

The overall expression for Transfer Time is $TT = (\#packets + 1store\&forward)TRANSP + 2PROP = (\frac{K}{P} + 1)\frac{P+48}{R} + 2PROP$.

Taking the first derivative with respect to the packet size P:

$\frac{d}{dP}TT = \frac{1}{R} - \frac{48K}{R}P^{-2}$.

Setting it to zero and solve: $P = 4\sqrt{3K}$.

Problem 3: Internet Routing – 20 points

[3 pts] (a) Describe the two dominant AS relationships on today's Internet. Give an example of an AS path that violates AS relationships and provide intuition why it is an violation.

Answer:
Peer to Peer, customer-Provider. Going from a provider to customer to another provider, customers should not provide transit service for its providers.

[9pts](b) Consider the following network:

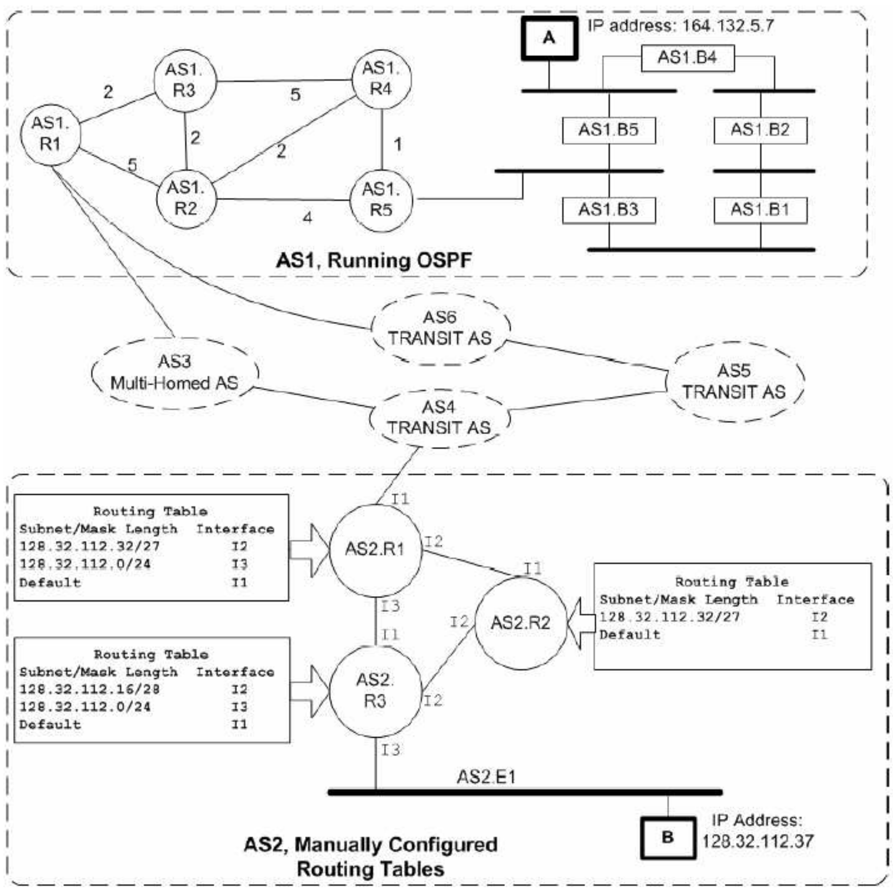

[2pts](i) AS1 uses OSPF routing with the link weights shown on the diagram. Specify the path the packet takes from AS1.R5 to AS1.R1 as a sequence of routers.

Answer:
AS1.R5 – AS1.R4 – AS1.R2 – AS1.R3 – AS1.R1
[2pts](ii). The inter-network use BGP routing between Autonomous systems. Specify the path a packet takes from AS1.R1 to AS2.R1 as a sequence of ASes. Note that AS4, AS5, and AS6 are transit ASes, while AS3 is a multi-homed AS.

Answer:
AS1-AS6-AS5-AS4-AS2
[2pts](iii). AS2 uses statically configured routing tables. The routers uses longest prefix matching. Specify the path that a packet with destination address 128.32.112.37 (Host B) would take from AS2.R1 as a sequence of routers.

Answer:
AS2.R1 – AS2.R2 – AS2.R3 – B
[3pts](iv). Host B replies to A by sending a packet, call it packet **P** to IP address 164.132.5.7 (Host A). Specify the path that packet **P** takes from B to AS2.R1 as a sequence of routers.

Answer:
B–AS2.R3 – AS2.R1

Run Dijkstra's algorithm on the following network to determine the routing table for Node 3, please show all intermediate steps.
[8 pts]

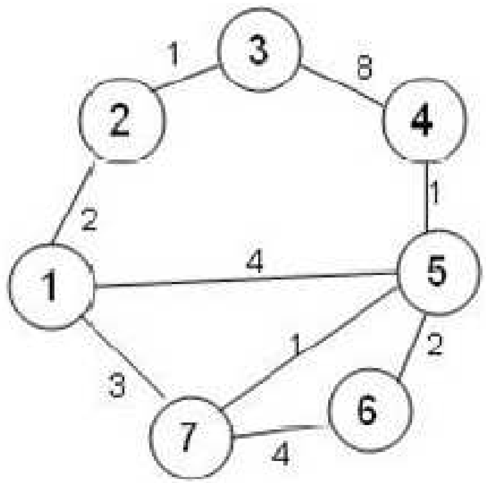

Answer:
It takes 7 steps as there are seven nodes all together. We show each step from node 3's perspective.

Step	Set S	1	2	3	4	5	6	7
0	{}	inf	inf	0	inf	inf	inf	inf
1	{3}	inf	1	0	8	inf	inf	inf
2	{3,2}	3	1	0	8	inf	inf	inf
3	{3,2,1}	3	1	0	8	7	inf	6
4	{3,2,1,7}	3	1	0	8	7	10	6
5	{3,2,1,7,5}	3	1	0	8	7	9	6
6	{3,2,1,7,5,4}	3	1	0	8	7	9	6
7	{3,2,1,7,5,4,6}	3	1	0	8	7	9	6

Problem 4: The Web Server Project – 20 points

Part A

```
/* reuse local address so that bind doesn't complain
      of address already in use. */
reuse = 1;
err = setsockopt(sd, SOL_SOCKET, SO_REUSEADDR, (char *) &reuse, sizeof(int));
tin_check_err(err, "setsockopt");
```

didn’t check bind return value

self.sin_port = ntohs((unsigned short) port);

Non-blocking IO, didn’t check the returned error code errno;

Didn’t close the socket before abort

Part B

Didn’t initialize the write set and the exception set

The reset operation should be inside the while loop.

The first parameter of select is the highest file descriptor in all given sets plus one, not MAX_CONNECTIONS.

Didn’t call select to protect recv

```
if (FD_ISSET(socks[i].sock, &rfds))
{ recv()
}
```

didn’t put newly accept sockets into the select set

```
// sending sockets
if (FD_ISSET(socks[i].sock, &wfds))
{
        rval=send(socks[i].sock,(char *)socks[i].buf+socks[i].buf_pos,to_send,0);

        if (rval==-1)
        {nconn -= close_connection(socks[i],ERRDISCNT);
                break;}
                socks[i].buf_pos+=rval;
}
```

Part C

a)

Number of bytes is greater than low-water mark for socket receive buffer

Read half of the connection is closed(FIN,EOF)

Listen socket, incoming connection;

Socket error pending

b) Messages from Server to Client (display)

Msg name	purpose
PASSWD	Ask user for PIN (password)
OK	last requested operation (PASSWD, WITHDRAWL) OK
ERR	last requested operation (PASSWD, WITHDRAWL) in ERROR
REGISTERED_COURSE	sent in response to REGISTERED request, list the registered courses
BYE	user done, display welcome screen at ATM

Messages from Client to Server

Msg name	purpose
HELO <unique name>	Let server know that there is a student registering in the machine Client transmits user unique name to Server
PASSWD <passwd>	User enters PIN, which is sent to server
REGISTERED	User requests for a list of registered courses
WITHDRAWL <course id>	User asks to withdraw course ID
REGISTER <course id>	User asks to register course ID
BYE	user all done

Interaction:

client server

HELO (userid) --------------> (check if valid userid)

<------------- PASSWD

PASSWD <passwd> --------------> (check password)

<------------- OK (password is OK)

REGISTERED -------------->

<------------- REGISTERED_COURSE

WITHDRAWL <id > --------------> check if the course is registered

If (yes)

<------------- OK

If (no)

```
                 <-------------      ERR
REGISTER <id >--------------> check if there is a time conflict
                 If (yes)
                 <-------------      OK
                 If (no)
                 <-------------      ERR

BYE               -------------->
                 <-------------    BYE
```

Problem 5: Transport Layer: TCP and UDP – 20 points

[3 pts] (a) TCP Reno uses a feature called *fast recovery* during slow start phase. Describe what this means.

Answer:
Fast recovery is tied with fast retransmit which is triggered when three more more duplicate ACKs are received, instead of waiting for retransmission timer to expire, a retransmission for the missing segment is performed. After fast retransmit sends what appears to be the missing segment, congestion avoidance, but not slow start is performed. This is the fast recovery algorithm. It is an improvement that allows high throughput under moderate congestion, especially for large windows.
[17pts] (b) Consider the following plot of TCP window size as a function of time. Assuming TCP Reno is the protocol used, answer these questions. Please justify your answers.

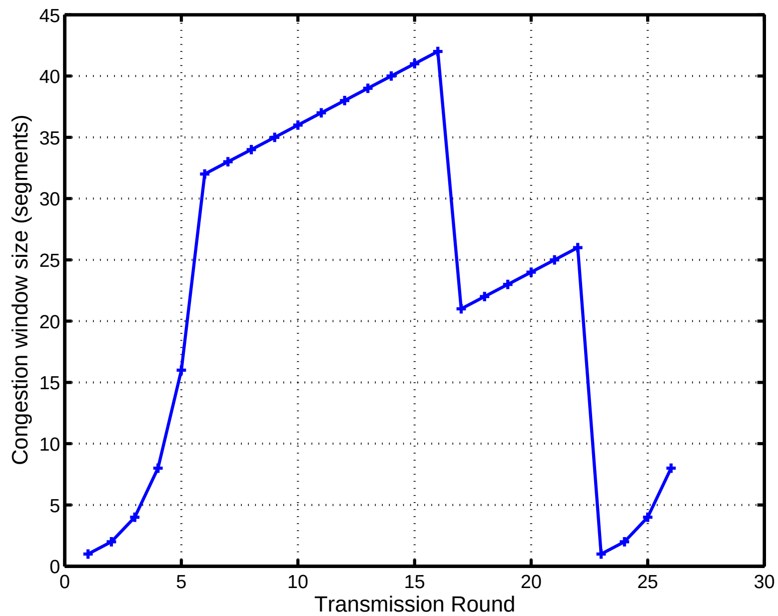

For your convenience, the data points in the plot are in the table below.
[1pt](a) Identify the intervals of time when TCP slow start is operating.
Answer:
1-6, 23-26

transmission round	congestion window size (segments)
1	1
2	2
3	4
4	8
5	16
6	32
7	33
8	34
9	35
10	36
11	37
12	38
13	39
14	40
15	41
16	42
17	21
18	22
19	23
20	24
21	25
22	26
23	1
24	2
25	4
26	8

[2pt](b) What is the name of the other TCP phase, identify its intervals.

Answer:
Congestion avoidance: 6-16, 17-22

[1pt](c) After the 16th transmission round, is segment loss detected by a triple duplicate ACK or by a timeout?

Answer:

Packet loss is recognized by a triple dup ACK. If there was a timeout the congestion window size would have dropped to 1.

[1pt](d) After the 22nd transmission round, is segment loss detected by a triple duplicate ACK or by a timeout?

Answer:
Due to time out, that's why the window size is 1.
[2pt](e) What is the initial value of `Threshold` at the first transmission round?

Answer:
It's initially 32, as this is window size when slow-start stops and congestion avoidance begins.

[2pt](f) What is the value of `Threshold` at the 18th transmission round?

Answer:
it is set to half of the value of the congestion window when packet loss is detected. When loss is detected during round 16, the congestion window size is 42, so the threshold is 21 during 18's round.

[2pt](g) What is the value of `Threshold` at the 24th transmission round?

Answer:
It is set to 13.
[2pt](h) During what transmission round is the 32th segment sent?

Answer:
round 1, packet 1 is sent; round 2: packets 2-3 are sent, round 3: packets 4-7 are sent; round 4, packets 8-15 are sent; round 5, packets 15-31 are sent. So at round 6, packets 32-63 are sent.
[4pt](i) 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 conges-

tion window size and of `Threshold`? **Answer:**
the congestion window and threshold are set to half of the current value of the congestion window when loss occurred. So the new values will be 4 for both.

Extra Credit: 10 points

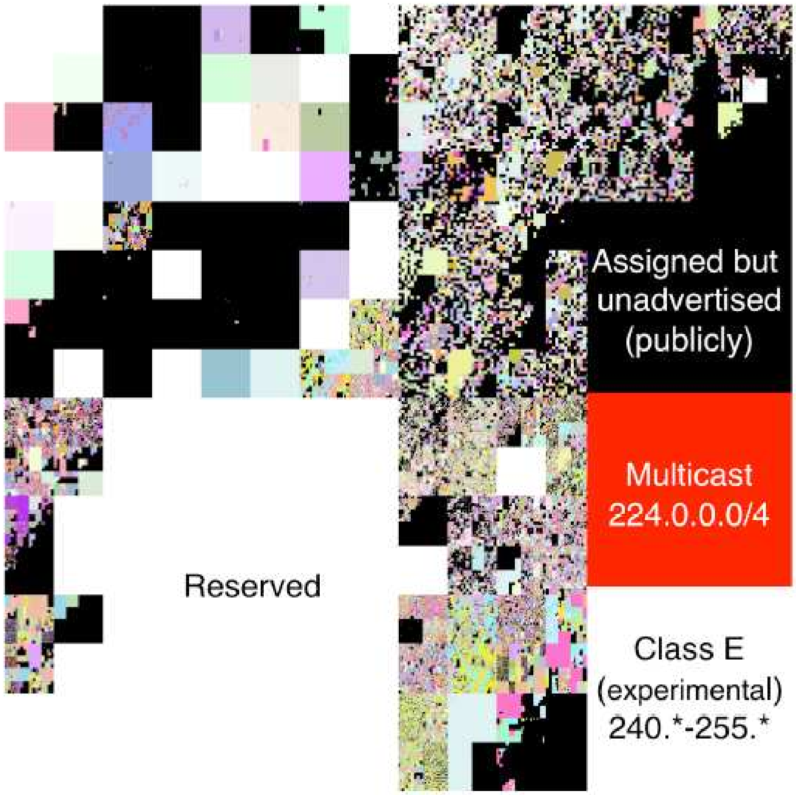

[6pts] Describe the algorithm of assigning IPV4 addresses to the square area shown, based on the example assignments. The upper left corner is 0.0.0.0.

[4pts] Given the problem with "pharming" (i.e., a hacker's attack aiming to redirect a website's traffic to another bogus website), describe any practical techniques to detect such problems.