Name: ______________________________

CS/ECpE 5516 **Final Exam** **Spring 1998**

In answering this exam, you may use any written references you wish, including the Bertsekas and Gallager text and written notes.

Problem	Score
Part 1	
1	/5
2	/5
3	/5
4	/4
5	/6
6	/6
7	/4
8	/5
Part 2	
1	/20
2	/20
3	/20
Total	/100

Part 1: Short problems [All problems, unless noted otherwise, are worth 5 points each.]

1. Consider a two-story office building, in which an Ethernet connects computers in offices throughout the building. You are the network troubleshooter. One week ago, the second story office occupants begin complaining that the delay of messages sent between their computers noticeably increased. You begin investigating the source of the problem. You discover that one week ago, the first story office occupants begin using a teleconferencing application amongst themselves that frequently sends video and audio over the network. What piece of equipment should you purchase and install to isolate the second story users from the video traffic, but still allow any host to reach any other host in the building?

Equipment name: ___________________________

Justification:

2. In Bertsekas and Gallager Chapter 2, an assumption was made that a network uses two types of packets: data and control. The control packets are used for acknowledgments and for the initiation/disconnection protocol. How does TCP manage to use just a single type of packet?

3. Why does a token ring have higher average latency than ALOHA, when the offered load is nearly zero?

Name: ________________________________

4. [4 points] List the two reasons that the length of a slot may be less than one unit in CSMA/CD slotted ALOHA.
#1)

#2)

5. [6 points] Give 2 reasons why the IP protocol (version 4, the one used universally today) cannot provide quality of service guarantees.
#1)

#2)

3 [6 points] List 2 reasons why a frame or datagram can be lost in the Internet.
#1)

#2)

7. [4 points] Give two reasons why the ATM protocol designers probably chose a small, fixed length (46 byte) data field length in ATM frames.
#1)

#2)

8. Sprint and MCI are two of the backbone carriers for the Internet in the United States. Suppose Sprint didn't want to route any of its traffic through MCI's backbone. What protocol could be used to achieve this?

Protocol name: ______________________

Justification:

Name: ________________________________

Part 2: Long problems [All problems are worth 20 points each.]

1. Refer to "Summary of Post-1987 TCP Specification" for the variable names used in this problem. Consider the use of TCP to send a message, which is partitioned into seven segments (S_0, S_1, S_2, ..., S_6). Assume that the initial values of TCP's variables are:

$RTT_0 = 0$
$DEV_0 = 0$
$RTO_0 = 0.5$

Also assume that the connection is first opened immediately before time 0. Use a separate timer for each segment. Using δ=1/8, calculate the value of

(a) RTT_i,
(b) CWND, and
(c) RTO_i

just after time 2.5 in the following scenario:

Time	Event
0	A sends S0
1	Ack for all packets up to S0 arrives
2	A sends S1
2.1	A sends S2
2.5	Ack for all packets up to S2 arrives

The scenario does *not* show segment retransmissions. It *does* show the first time each segment is sent and *all* acknowledgements received up to time 2.5.

Show your work by completing the table below. To get at least partial credit, write the formulas used to fill each cell's value. Write "N/C" in a cell if there is no change in the value from the cell immediately above.

Time	Event	DIFF	RTT_i	DEV_i	RTO_i	CWND	Timers Set
0-	Initial	-	0	0	0.5	1	none

Final values:

RTT_i = ________________________

CWND = ________________________

RTO_i = ________________________

Name: ________________________________

2. An Ethernet frame consists of the following fields:

64 bit preamble
48 bit source address
48 bit destination address
16 bit frame type
358 to 12,000 bits of frame data
32 bit CRC

Consider two hosts, A and B, connected to a 10Mb/sec Ethernet, with a round trip delay of 1 millisecond. What is the maximum throughput at the service interface to the transport protocol of a file transfer from host A to host B using TCP/IP, assuming that no traffic other than the file transfer is present on the Ethernet? (In other words, what throughput will layer 5 see when it uses TCP in this environment?) State any assumptions that you require. Recall that the maximum TCP window size is 2^{16} bytes, and that the TCP and IP headers each have a minimum length of 20 bytes. (Note: If you do not have a calculator, you may leave the final answer in an unevaluated expression, or you may round off the numbers to get a "ballpark" answer.)

Maximum throughput = ____

Assumptions:

Justification (or how you derived above answer):

3. Can you use an M/M/1 queue to compute the average time that a customer spends in a fast food restaurant (e.g., Burger King, McDonalds)?

____ YES (There are no violations of the assumptions)
____ MAYBE (There are minor violations of the assumptions)
___ NO (There are major violations of the assumptions)

In the table below, list *each* assumption that is required to apply M/M/1, and *for each* assumption whether the fast food restaurant meets the assumptions. (There may be more table rows than assumptions!)

Assumption	**Met by fast food restaurant?**

Solution

Short problems:

1. Install a bridge between the two floors.
A bridge filters traffic between the floors. A router will do this, but it is more expensive and/or lower performance. A repeater will not do this. A hub broadcasts incoming traffic on all outgoing ports, so a hub will also not do this.

2. The control packets in TCP consist of zero data length packets with one of the code bits set in the header (ACK, RST, SYN, FIN).

3. A host on an idle the token ring must wait for the token before transmitting, whereas a host on an idle Ethernet transmits immediately.

4. #1: Collision detect; #2: Carrier sense

5. Here are some possible reasons, but not an exhaustive list:
- IP doesn't reserve resources
- IP does "best effort delivery"
- IP is connectionless, while quality of service is normally associated with a connection. A connectionless protocol has trouble guaranteeing quality of service parameters that depend on the relationship of two packets, such as jitter.
- IP may route different packets between the same source/destination pair over different routes, and the IP routing algorithm does not use quality of service measures (the Interior Gateway Protocol uses Open Shortest Path First routing).

6.
- layer 2 misses the end of a frame
- at layer 3, IP's buffer is full when a datagram arrives
- time to live field in IP header may reach zero due to a loop in routing tables.
- a node may crash that is buffering the datagram (note that a link crash may not be enough - if the link happens to be carrying the datagram at the time it crashes, or if all outgoing links of a node that buffers a datagram are down, then the datagram would be lost)
- IP header gets corrupted, so datagram destination is unknown.

7.
- permits voice traffic without too much echo
- allows reduced variance in latency
- pipelining effect: ACK of first packet can occur before last packet is transmitted in paths with many hops
- avoids slow truck phenomena
- simpler switch implementation, due to fixed length format (e.g., framing is easier)

Note: The switch is *not* faster with a small packet length. In fact, a *disadvantage* of ATM is that long messages require many packets, which causes a switching action to be performed many times. So you need a higher performance switch with short packets than long ones.

8. Exterior Gateway Protocol (EGP)
Justification: The Path field in EGP's *update* message list the Atomic Systems through which a route passes. Sprint and MCI operate different

Atomic Systems (this must be the case, since one organization manages an AS). Thus the AS list in the Path field can be used to avoid MCI's backbone. (See slide 31 in the lecture notes on Routing.)

Long problems:

1.

1.

Time	Event	DIFF	RTT_i	DEV_i	RTO_i	CWND	Timers Set
0-	Initial	-	0	0	0.5	1	none
0	Send S0	-	N/C	N/C	N/C	N/C	S0@0 + 2^0*0.5 = 0.5
0.5	Resend S0	-	N/C	N/C	N/C	N/C	S0@0.5 + 2^1*0.5 = 1.5
1	Ack up to S0	1-0=1	0+.125*1 = 0.125	0 + .125*(1-0) = .125	.125 + 2*.125 = .375	2	none
2	Send S1	N/C	N/C	N/C	N/C	N/C	S1@2.375
2.1	Send S2	N/C	N/C	N/C	N/C	N/C	S2@2.475
2.375	Resend S1	N/C	N/C	N/C	N/C	N/C	S1@ 2.375+ 2*.375 = 3.125, S2@2.475
2.475	Resend S2	N/C	N/C	N/C	N/C	N/C	S1@3.125, S2@ 2.475 + 2*.375 = 3.225
2.5	Ack up to S2	N/C	N/C	N/C	N/C	N/C	none

Note that at time 2.5 the value of RTT_i, DEV_i, and RTO_i are *not* updated because segments S1 and S2 were both retransmitted.

Solution 2:

Recalculating the values of DIFF, RTT, DEV, RTO only upon an acknowledgement without retransmission.

Time	Event	DIFF	RTT_i	DEV_i	RTO_i	CWND	Timers Set
0-	Initial	-	0	0	0.5	1	none
0	Send S0	-	N/C	N/C	N/C	N/C	S0@0 + 2^0*0.5 = 0.5
0.5	Resend S0	-	N/C	N/C	N/C	1	S0@0.5 + 2^1 *0.5= 1.5
1	Ack up to S0	N/C	N/C	N/C	N/C	N/C	none
2	Send S1	N/C	N/C	N/C	N/C	N/C	S1@2.0+0.5 = 2.5

2.1	Send S2	N/C	N/C	N/C	N/C	N/C	S1@2.5 S2@2.1+0.5 = 2.6
2.5	Ack up to S2	0.4-0=0.4	0.125*0.4 = 0.05	0.125*0.4 = 0.05	0.05+2*0.05 = 0.15	1+2 = 3	none

2.
There are several approaches to solving the problem. One is shown below. (A reasonable approach gets most of the credit for the problem.)

MAX THROUGHPUT = 9.57 Mbit/sec

ASSUMPTIONS:

The only assumption required by the solution below is that the delay required for processing in the protocol stack is zero. (Zero gives an upper bound on throughput, which is what we want.)

JUSTIFICATION:

There are two things that could limit the transport-level throughput:
- TCP's window will limit the rate at which the sender can transmit
- The 10Mbit/sec data rate of Ethernet will limit the rate at which the sender can transmit

So the *maximum throughput* is the smaller of the two.

TCP window's limit on transport-level throughput:

Use Little's Law: $N=\lambda T$

Because TCP max window is 2^{16} bytes, $N \leq 2^{16}$ bytes.
We know $T=10^{-3}$ sec.
Thus $\lambda \leq N/T$, or $\lambda \leq 2^{16}/10^{-3}$, or $\lambda \leq 65.5$ Mbytes/sec

10Mbit/sec Ethernet's limit on transport-level throughput:

The 10Mbit/sec figure is the *data-link* level throughput. To obtain the maximum *transport* service interface throughput, we must compute the amount of overhead due to layers 2, 3, and 4.

Layer 2 overhead for each frame sent is 8+6+6+2+4=26 bytes, using the sizes for preamble, source address, etc., given in the problem.

Layer 3 overhead for each datagram is 20 bytes.

Layer 4 overhead for each segment is 20 bytes.

Thus total overhead is 26+20+20=66 bytes.

The maximum transport service interface throughput will be obtained if TCP sends the maximum data size permitted by Ethernet, which is 1460 bytes of layer 5 data. (The problem states that an Ethernet frame can contain as much as 12,000 bits, or 1500 bytes, of data. 40 bytes are needed for TCP and IP headers, leaving 1500=1460 bytes of layer 5 data.)

So for every 1460 bytes sent, 66 additional bytes of overhead bytes are sent. Thus an upper bound on the fraction of Ethernet that is not available for Ethernet is 66/(1460+66)*10Mbit/sec, or 4.3%. Thus the limit on transport service interface throughput is 9.57 Mbit/sec.

The minimum of the two throughputs, 65.5 Mbytes/sec and 9.57 Mbit/sec, is the answer. That is <u>9.57 Mbit/sec</u>.

3.
NO

You must check 3 assumptions (corresponding to the “M” / “M” / “1” parts):

Assumption	**Met by fast food restaurant?**
1. Single server	YES – if there is only one person behind the counter NO : if there are two or more people behind the counter
2. Exponential interarrival time	MAYBE – customers arrive one at a time, and can have an arbitrarily small time between arrivals. It’s unlikely to have a long time between arrivals (say days), so the distribution’s tail must be long, like the exponential distribution. The distribution probably obeys the memoryless distribution, if arriving customers are unrelated. (The time you arrive isn’t correlated to the time I arrive.) On the other hand, the distribution might not be memoryless when a family arrives (N people have a correlated arrival time).
3. Exponential service time	PROBABLY NOT – the time it takes to fill a customer order probably <u>*does not*</u> obey the memoryless property. (Food will burn if you use a memoryless property to decide how long to cook it!)