Name: ____________________________ **ID:** _____________

EECS 122 Final Exam

May 22, 2004

- **Please write your name on top of each page.**
- **If required, use the blank side of each page.**

Question	Topic	Total Points	Points Received
1	Modeling	10	
2	Shannon's Theorem	6	
3	Wi-Fi	12	
4	Multiple Choice	20	
5	Scheduling	12	
6	TCP Flow Control	10	
7	Internetworking	20	
8	Sockets, etc.	10	
Total		**100**	

Name: ______________________________

Problem 1: Modeling

(1) Consider the logical representation in the figure below.

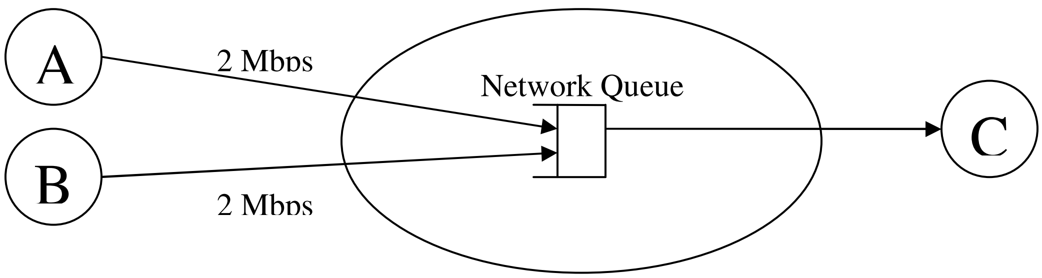

The sources A and B are transmitting data at the average rate of 2 Mbps each to the destination C across a network. The data consists of 1000 bit packets. Assume that all the contribution to the end-to-end delay comes from the store-and-forward network queue shown in the figure, and that there are no packet losses. This queue has an average occupancy of 12 packets. What's the average end-to-end delay experienced by each packet? **[3 points]**

Using Little's Law, [Avg Occupancy] = [Avg arrival rate] x [Avg Delay]
12 x 10^3 bits = 4 x 10^6 bits/ sec x [Avg Delay]
Avg Delay = 3 ms

Problem 1: Modeling (Continued)

(2) Now consider the modified network as represented in the figure below.

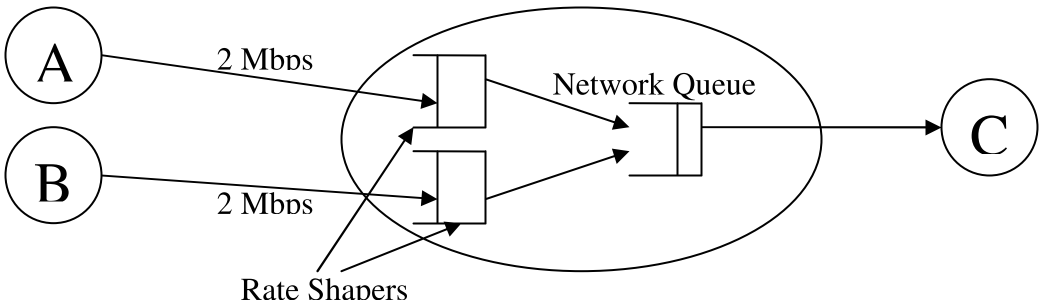

Here each of the data streams considered in the part (1) is first passed through a rate shaper queue, and then fed into the network queue. A rate shaper queue puts out the "shaped" traffic such that consecutive outgoing packets have a minimum separation of a specified duration. (Knowledge of the precise mechanism used by the rate shapers is not required to solve this problem.) Assume that all the contributions to the end-to-end delay only come from the rate shaper and the network queues, and that there are no packet losses. Each rate shaper has an average occupancy of 12 packets, and the network queue has the average occupancy of 1 packet. What's the average end-to-end delay experienced by each packet? **[4 points]**

Using Little's Law, [Avg Occupancy] = [Avg arrival rate] x [Avg Delay]
25 x 10^3 bits = 4 x 10^6 bits/ sec x [Avg Delay]
Avg Delay = 6.25 ms

(3) In light of the end-to-end delays determined in the parts (1) and (2), describe reasons that make the rate shapers desirable. **[3 points]**

Network buffers are expensive. If one flow is overactive or misbehaving, rate shaping provides protection for the well behaving flow.

Name: ________________________________

Problem 2: Shannon's Theorem

For phone line, the available channel spans 300 Hz to 3.3 KHz. These lines typically have signal-to-noise ratio of 30 dB.

(1) What is the channel capacity of such a phone line (in Kbps)? (To simplify the calculations, you can approximate $\log_2(1001)$ as 10.) **[3 points]**

$C = B \log_2(1 + S/N) = 3000 \text{ Hz} \log_2(1 + 1000) \sim= 30\text{Kbps}$

(2) In light of your answer for the part (1), what advances have made 56 Kbps modems possible? **[3 points]**

Improved S/N ratio, increased bandwidth.

Problem 3: Wi-Fi

Suppose there are only two nodes, a source and a destination, equipped with 802.11b wireless LAN radios that are configured to use RTS/CTS for packets of all sizes. The source node has a large amount of data to send to the destination node. Also, the nodes are separated by 750 meters, and have powerful enough radios to communicate at this range. No other nodes operate in the area. What will be the data throughput between source and destination assuming that packets carry 1100 bytes of data? **[12 points]**

Important Parameters:

- Propagation speed is 3×10^8 meters/sec. (Make sure not to neglect propagation delay!)
- DIFS – the first component of the delay between noticing the medium being idle and beginning a new RTS is 50µs.
- The contention window (CW), the additional delay after DIFS until a new packet begins transmission, is chosen randomly for each packet. When there is no contention, it is uniformly likely to be anywhere from 0 to 31 slots long, where one slot is 20µs.
- The preamble, the physical layer header, the MAC header and trailer take a combined 200 µs per packet to transmit.
- The Data is transmitted at an 11 Mbps rate.
- An ACK, RTS, and CTS each have a 200 µs transmission time.
- SIFS has a value 10 µs. Recall that SIFS is used as follows:
    - The delay a node waits after receiving an RTS to reply with a CTS
    - The delay a node waits after receiving a packet before sending a data packet
    - The delay a node waits after receiving a packet before replying with an ACK

Total time = DIFS + CW + RTS + SIFS + CTS + SIFS + PREAMBLE + TX + SIFS + ACK + 4xPROP

TX = Data packet / Transmission Rate = 1100 x 8 / 11e6 = 800us
4 x PROP= 4 x 750m / 3e8 = 10us
DIFS = 50us
CW = 15.5 x 20us = 310us
RTS + CTS + ACK = 3 x 200 = 600us
PREAMBLE = 200us
SIFS x 3 = 10us x 3 = 30us

Total time = 800us + 10us + 50us + 310us + 600us + 200us + 30us = 2000us
Efficiency = 800/2000 = 0.4
Throughput = 0.4 * 11Mbps = 4.4Mbps

Name: ______________________________

Problem 4: Multiple Choice Questions

Select all items that apply. [Each question 2 points]

1. Packet switching, compared to circuit switching

(a) Offers a more predictable quality of service

(b) Exploits the gap between average and peak traffic rate

(c) Requires a longer connection setup phase

(d) Is not suitable for telephone services

(e) None of the above

2. In the Internet, Layer 3 is used

(a) to speed up transmissions, because routers are faster than switches

(b) for scalability – in particular, the topology-based structure of layer 3 addresses reduces the size of routing tables

(c) to interconnect networks with different technologies

(d) to eliminate the need for the address resolution protocol

(e) None of the above

3. CIDR is

(a) an addressing scheme designed for multicast

(b) a method to differentiate classes of service

(c) a more efficient way to use the IP address space

(d) incompatible with Ethernet

(e) None of the above

4. The capacity of a communication channel is

(a) the realized bit rate of the link

(b) the maximum transmission rate of the link with a given BER

(c) equal to twice the maximum frequency that the channel transmits

(d) the maximum possible reliable transmission rate of any link that uses that channel

(e) None of the above

5. One advantage of a path vector protocol over distance vector and link state is that

(a) it requires fewer messages between the routers

(b) it is guaranteed to converge to the shortest paths

(c) it enables the implementation of policies that may not correspond to simple metrics

(d) it enables load balancing across parallel paths

(e) None of the above

Name: _______________________________

Problem 4: Multiple Choice Questions (Continued)

6. One key advantage of 802.11 (Wi-Fi) over wired Ethernet is that
- (a) it uses carrier sensing
- (b) it does not need to back off exponentially
- (c) it supports multiple classes of service
- (d) it can hide exposed terminals
- **(e) None of the above**

7. TCP tries to
- (a) minimize the number of retransmissions of erroneous packets
- **(b) share the bandwidth of bottleneck links among different connections**
- (c) acknowledge packets link by link instead of end to end
- (d) correct packets so that they do not have to be retransmitted the third time (using the 3-dup ack)
- (e) None of the above

8. IP Multicast
- (a) automatically sets up a one-to-one connection between the source and each destination
- (b) supervises the reliable transmission from one source to many destinations
- (c) speeds up TCP by creating multiple parallel transmissions from the server to the client
- (d) uses clients to retransmit erroneous packets to other subscribers to the multicast group
- **(e) None of the above**

9. Select the correct statements, if any:
- (a) NAT translates network addresses into layer 2 addresses
- (b) NAT translates IPv4 addresses into IPv6 addresses
- (c) DHCP configures the host dynamically based on the link bandwidth
- **(d) DHCP provides a temporary IP address to a host**
- **(e) NAT and DHCP increase the efficiency of the use of IP addresses**
- **(f) A server should not use DHCP because clients would not know how to find it**
- (g) None of the above

10. Fibers are used today for higher bandwidth links because
- (a) light travels faster than electricity
- **(b) their attenuation at high frequencies is smaller than that of copper wires and cables**
- (c) they simplify timing recovery
- (d) one can use optical switches instead of routers
- (e) None of the above

Name: ______________________________

Problem 5: Scheduling

For the purpose of this problem, you can consider the idealized Generalized Processor Sharing (GPS) as being equivalent to Weighted Fair Queueing (WFQ).

(1) Consider a system of four queues being serviced according to a WFQ scheduling policy. The weights given to the four queues (A, B, C, D) are 4, 1, 3, and 2 respectively. They are being serviced by a processor at the rate of 10 Mbps.

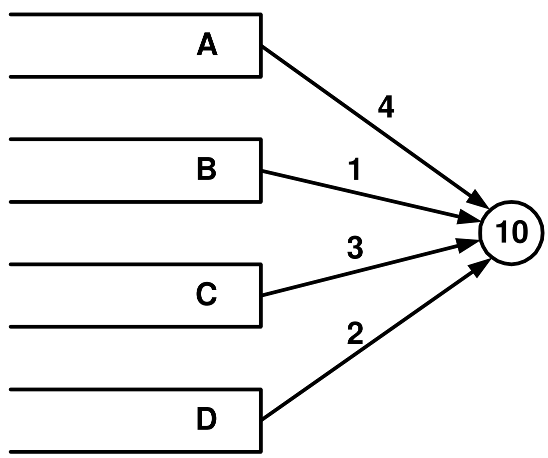

Part (a)

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 these four queues. We have filled in the first two rows to get you started! **[Each row 2 points]**

INPUT RATES				RECEIVED RATES			
A	B	C	D	A	B	C	D
1	1	1	1	1	1	1	1
10	10	10	10	4	1	3	2
6	6	2	2	4.8	1.2	2	2
8	0	0	8	6.67	0	0	3.33
1	5	3	5	1	2	3	4

Problem 5: Scheduling (Continued)

(2) This time we still have the same four queues, but the processors are split in two levels and are based on priority queueing, WFQ, or Time Division Multiplexing (TDM).

Each of the two second level processors X and Y have a weight of 5. The policy implemented between these two processors is that of TDM, i.e. each of the processors gets a fixed share of the bandwidth.

Processor X implements a WFQ scheduling between its two queues A and B. A has a weight of 3 and B has a weight of 2.

Processor Y implements a strict priority between its two queues C and D. Queue C has strictly higher priority than queue D.

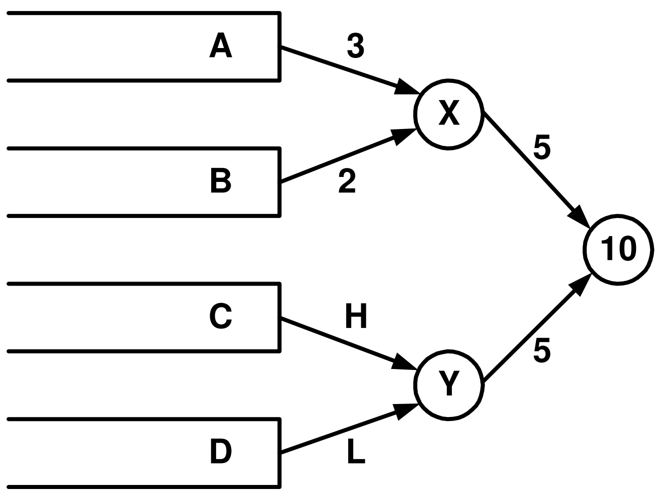

Part (b)

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 these four queues. **[Each row 2 points]**

INPUT RATES				RECEIVED RATES			
A	B	C	D	A	B	C	D
6	6	2	2	3	2	2	2
2	6	2	6	2	3	2	3
6	2	6	2	3	2	5	0

Name: ______________________________

Problem 6: TCP Flow Control

(1) TCP congestion window size (cwnd) is adjusted differently in the two separate phases of the congestion window evolution. In one phase the window size evolves as

(a) cwnd = cwnd + 1, every time an ack is received,

and in the other phase it evolves as

(b) cwnd = cwnd + 1/cwnd, every time an ack is received.

Indicate which phase corresponds to the Slow Start and which phase corresponds to the Congestion Avoidance. **[2 points]**

(a) corresponds to Slow Start, and (b) to Congestion Avoidance

Assuming no packet losses, show how the congestion window size evolves differently in the two phases as a result of this difference in the cwnd update mechanism. **[3 points]**

Let the window size at the beginning of the RTT period be W. And assume there are no losses in the RTT that we are considering.

During Slow Start, each ack updates the window size by 1. The source will receive W acks, thus increasing its window size by W + W = 2W. This leads to an exponential increase.

During Congestion Avoidance, the window size increases by 1/W. So the size at the end of the RTT will be approximately W + (1/W) x W = W+1. This gives a linear increase.

Problem 6: TCP Flow Control (Continued)

(2) Consider two TCP flows A and B:

- The RTT for flow A is 100 ms while the RTT for flow B is 200 ms.
- Recall that ssthresh denotes the threshold at which the congestion window size evolution switches over from the Slow Start phase to the Congestion Avoidance phase. Both flows have the ssthresh value of 8.

At time T seconds, both flows have just had a timeout, and so their window size is set to 1. Calculate how many packets each flow is able to send in the next 1 second (including at time T+1 seconds). Assume that each packet can be transmitted in 0 time, and that there are no dropped packets for either flow during the interval [T, T+1] seconds. **[5 points]**

1 + 2 + 4 + 8 + 9 + 10 + 11 + 12 + 13 + 14 + 15 = 99 for flow A
1 + 2 + 4 + 8 + 9 + 10 = 34 for flow B

Name: ______________________________

Problem 7: Internetworking

Consider the following network:

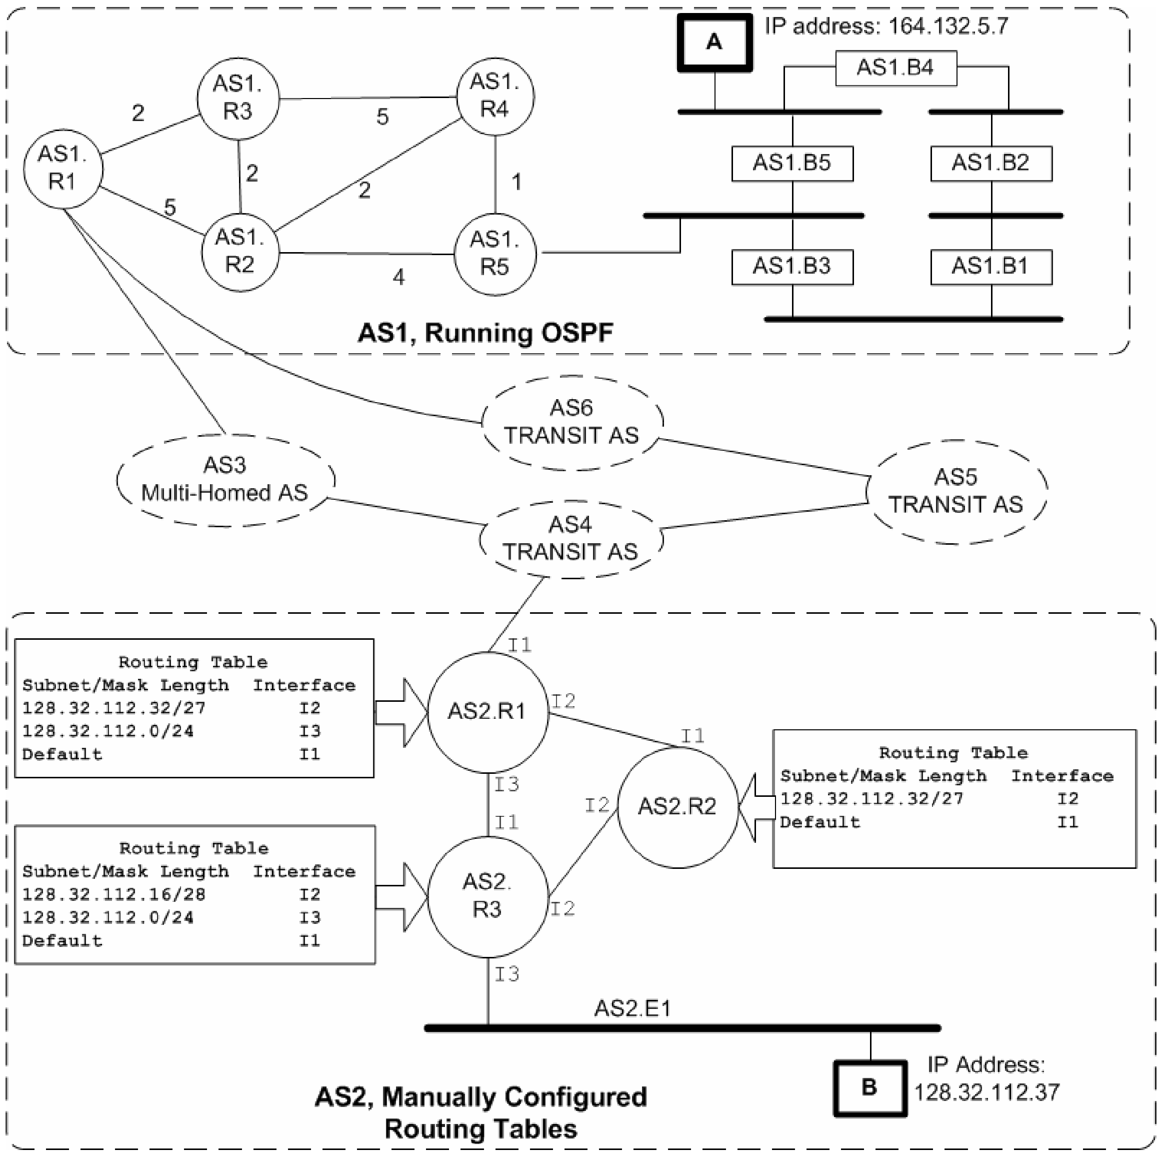

Name: ________________________________

Problem 7: Internetworking (Continued)

(1) Host A, located in AS1, sends a packet to Host B located in AS2. Host A's default gateway router is AS1.R5.

i. The Ethernet bridges in AS1 (AS1.B1 – AS1.B5) are learning bridges that run the spanning tree algorithm. Assuming the bridges have been on for a long enough time to reach a steady state, specify the path a packet takes from Host A to AS1.R5 as a sequence of bridges. Note that the bridges' names (e.g. AS1.B1) are chosen to reflect the order of the bridge ID's:

ID(AS1.B1) < ID(AS1.B2) < … < ID(AS1.B5) **[2 points]**

A --- AS1.B4 --- AS1.B2 --- AS1.B1 --- AS1.B3 --- AS1.R5

ii. 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. **[2 points]**

AS1.R5 --- AS1.R4 --- AS1.R2 --- AS1.R3 --- AS1.R1

iii. The inter-network uses BGP routing between autonomous systems. Specify the path a packet takes from AS1.R1 to AS2.R1 as a sequence of ASs. Note that AS4, AS5, and AS6 are transit ASs, while AS3 is a multi-homed AS. **[2 points]**

AS1 --- AS6 --- AS5 --- AS4 --- AS2

iv. AS2 uses statically configured routing tables. The routers use **longest prefix matching**. Specify the path that a packet with destination IP address 128.32.112.37 (Host B's Address) would take from AS2.R1 as a sequence of routers. **[2 points]**

AS2.R1 --- AS2.R2 --- AS2.R3 --- B

Name: ______________________________

Problem 7: Internetworking (Continued)

(2) Host B replies to A by sending a packet, call it packet **P**, to IP address 164.132.5.7 (Host A's address.)

i. Specify the path that packet **P** takes from B to AS2.R1 as a sequence of routers. **[2 points]**

B --- AS2.R3 --- AS2.R1

ii. Suppose that when **P** arrives at AS1.R5, that AS1.R5 has no record of A's MAC address. However, AS1.R5 does know it should route packets destined for A's IP address to the bridged Extended LAN shown in the diagram. What will AS1.R5 do next? **[2 points]**

AS1.R5 will do an ARP request on the LAN with A's IP address, in order to find out the MAC.

iii. Adding to the scenario described in the part (2-ii), suppose that all the bridges in the bridged Extended LAN have no record of A's MAC address when **P** arrives at AS1.R5. Assuming AS1.R5 does what you said it would in the part (2-ii), will the bridges learn A's MAC address during the events that follow? If so, describe how they will learn it. Also, will **P** make it to its destination? **[2 points]**

An ARP request will be broadcast on the LAN. A will hear it and reply to it.
Bridges that see the ARP reply will learn of A's MAC address.
P will make it to the destination.

Problem 7: Internetworking (Continued)

(3) The routing tables of AS2 are misconfigured, and will result in packets with a particular range of IP addresses to loop between routers of AS2 indefinitely (in reality, until their TTLs expire).

i. Specify the full range of IP addresses that will loop, and specify which routers the looping packets would traverse. **[4 points]**

128.32.112.16 to 128.32.112.31
Loop is R1 --- R3 --- R2 --- R1

ii. You are allowed to delete one line of the routing tables in one of the routers. Your objective is to eliminate the loop, while making sure that hosts with IP addresses 128.32.112.0 through 128.32.112.255 located on AS2.E1 have connectivity both to and from other ASs. Which line would you delete? **[2 points]**

Delete the line 128.32.112.16/28 from AS2.R3

Name: ______________________________

Problem 8: Sockets, etc.

Please answer each question briefly.

(1) A student wrote a client-server application program, and is now running his program on two communicating hosts. The student's program works perfectly when he uses only 1 byte for the packet sequence number in his custom-designed packet header, but works intermittently when he changes the program (on both hosts) to use a 2 byte sequence number in his header. What information can you ascertain from this, with almost certain probability? **[2 points]**
[Select all items that apply]

(a) His client and server are running on different operating systems
(b) His client and server are running on different hardware architectures
(c) His client and server are being compiled with different compilers
(d) His client and server are multi-threaded

(2) We know that IP packets can be fragmented based on the MTU of the underlying network. We know that UDP packets do not have any scheme for reassembling packets in a certain order. If so, how is it possible that UDP/IP works? Limit your explanation to ~3 lines. **[2 points]**

Fragmentation reassembly is handled by the IP layer. The IP header has a fragment offset field that the IP layer uses to reassemble the packet.

(3) A connection is defined by the IP-port pair of both the client and server. Thus, any connection to a server is unique as long as the origination IP/port pair is unique => a webserver running on port 80 will not run out of ports if it gets multiple client requests. So then why is it desirable for a webserver to maintain short term connections instead of persistent connections? Limit your explanation to ~3 lines. **[2 points]**

The server needs to keep state information about each active TCP connection. With long lived TCP connections, more will be active at a time, and the server would have to keep much more state information. Managing all the state information might slow down a busy server.

Name: ________________________________

Problem 8: Sockets, etc. (Continued)

(4) A student designed a client to receive data on two UDP sockets, `sock_a` and `sock_b`, and then store the arriving data to two separate files with file pointers `file_a` and `file_b`, respectively. Data is arriving at to `sock_a` at 1 Kbps, and to `sock_b` at 10 Kbps. Yet the student notices that `file_b` is being written to at a rate of only 1 Kbps. Consider a snippet of his code below, and explain what he did wrong. **[4 points]**

- Assume that he/she initialized his variables correctly, created his sockets correctly, bound his sockets to distinct port numbers, opened his files correctly, made `buffer` large enough, and used the correct syntax for his `recvfrom()` and `fwrite()` calls. Also, assume that both the UDP flows are transferring segments of the same fixed size.

```
while (1)
{
   num_bytes_received = recvfrom (sock_a, buffer, buffer_size, 0,
                                  (struct sockaddr *) &sender, &length);
   if (num_bytes_received > 0)
          fwrite (buffer, 1, num_bytes_received, file_a);

   num_bytes_received = recvfrom (sock_b, buffer, buffer_size, 0,
                          (struct sockaddr *) &sender, &length);
   if (num_bytes_received > 0)
          fwrite (buffer, 1, num_bytes_received, file_b);
}
```

`recvfrom` for **`sock_a`** is blocking. `recvfrom` blocks until a UDP segment is available, and then removes just single segment. Should have used the `select()` function.