

Name: \_\_\_\_\_

CS244a Final  
March 19, 2007

Student ID #: \_\_\_\_\_


Professor Nick McKeown

Campus/SITN-Local/SITN-Remote? \_\_\_\_\_

mGrade: MC: \_\_\_\_\_ 19: \_\_\_\_\_ 20: \_\_\_\_\_ 21: \_\_\_\_\_ 22: \_\_\_\_\_ 23: \_\_\_\_\_ 24: \_\_\_\_\_ Total:  
\_\_\_\_\_

# CS244a: An Introduction to Computer Networks

Final Exam: Monday March 19, 2007



## **You are allowed 3 hours to complete this exam.**

- (i) This exam is closed book and closed notes. However, you may refer to a sheet of 8.5"x11" paper (double-sided) of your own design.
- (ii) Write your solution directly onto this exam. Be sure to write your name and student ID clearly on the front of the exam.
- (iii) Don't panic! Be sure to start by reading the exam all the way through. Then answer the questions in whatever order you choose.
- (iv) Show your reasoning clearly. If your reasoning is correct, but your final answer is wrong, you will receive most of the credit. If you just show the answer without reasoning, and your answer is wrong, you may receive no points at all.

## **The Stanford Honor Code**

In accordance with both the letter and spirit of the Honor Code, I didn't cheat on this exam.

Signature: \_\_\_\_\_

Name: \_\_\_\_\_

CS244a Final  
March 19, 2007

Student ID #: \_\_\_\_\_









Professor Nick McKeown

Campus/SITN-Local/SITN-Remote? \_\_\_\_\_

---

### Short Multiple Choice Questions.

**Instructions:** *in the following questions, check all listed assertions that appear to be correct. There is at least one correct assertion per question, but there may be more. Each correct assertion checked will earn you one point. For each incorrect assertion you check, you will lose one point. If you don't know an answer, checking no assertion will neither earn you nor lose you any points.*

1. **General.** What transmission rate is needed to transmit a 4" x 6" photograph (uncompressed, and with a resolution of 1200 dots per inch and 24 bits per pixel) in 1 second?
  - (a.) 691,200kb/s
  - (b.) 28.8kb/s
  - (c.) 28.8kbits
  - (d.) 8.29Mb/s
  -  (e.) 829Mb/s
  
2. **Layering.** "Layering" is commonly used in computer networks because:
  - (a.) It forces all network software to be written in ANSI 'C'.
  - (b.) Encapsulation is the lowest overhead method to transmit data.
  -  (c.) It allows widespread code and implementation re-use.
  - (d.) It keeps networks warm enabling them to run faster.
  
3. **Elasticity Buffer.** An elasticity buffer is used to store bits arriving at a network interface. If the receiving station uses a 20-bit elasticity buffer and the clocks of the transmitter and receiver have a minimum frequency of 99.999MHz and a maximum frequency of 100.001MHz, which of the following statements are true:
  -  (a.) All packets have to be less than or equal to 62,500 bytes long.
  - (b.) All packets have to be less than or equal to 4500 bytes long.
  - (c.) The two clocks have a tolerance of +/- 100ppm.
  -  (d.) The two clocks have a tolerance of +/- 10ppm.
  - (e.) The transmitter's clock is always faster than the receiver's clock.
  
4. **The Internet Protocol.** IP provides:
  - (a.) a connection-oriented service.
  -  (b.) unreliable delivery.
  -  (c.) best-effort delivery.
  - (d.) in-sequence delivery.
  
5. **The Transmission Control Protocol.** TCP provides:
  -  (a.) a connection-oriented service.
  - (b.) unreliable delivery.
  - (c.) best-effort delivery.
  -  (d.) in-sequence delivery.









Name: \_\_\_\_\_

CS244a Final  
March 19, 2007

Student ID #: \_\_\_\_\_

Professor Nick McKeown

Campus/SITN-Local/SITN-Remote? \_\_\_\_\_

- 
6. **ICMP.** Which of the following applications requires the use of ICMP:
- (a.) ftp.
  -  (b.) traceroute.
  - (c.) tcpdump.
  - (d.) netstat.
7. **Coding.** Suppose a 10Mb/s adapter sends into a link an infinite stream of 1's using Manchester encoding. The signal emerging from the adapter will have how many transitions per second?
- (a.) 10 million transitions per second.
  - (b.) 5 million transitions per second.
  -  (c.) 20 million transitions per second.
  - (d.) No transitions.
  - (e.) None of the above.
8. **Link layer.** Which of the following are true?
-  (a.) An Ethernet switch can interconnect a 10Mb/s Ethernet network and a 1Gb/s Ethernet.
  - (b.) An Ethernet hub can interconnect a 10Mb/s Ethernet network and a 1Gb/s Ethernet.
  - (c.) An Ethernet network cannot detect collisions until it has computed a checksum over the frame.
  -  (d.) 4B/5B is considered more efficient than Manchester encoding because more user data is transmitted in same amount of time.
  -  (e.) The 802.11b wireless protocol incorporates a link-layer ACK not present in regular Ethernet.
9. **Multicast routing.** In reverse path broadcast, which of the following are true:
-  (a.) Packets reach every router in the network.
  - (b.) Packets are guaranteed to follow the lowest-cost path from source to destinations.
  -  (c.) Unless a packet arrives to the router on the interface that is on the shortest path to the source, the packet is dropped.
  - (d.) Every router receives exactly one copy of each packet.
10. **CSMA/CD.** In a CSMA/CD network we require  $TRANSP > 2PROP$  because:
- (a.)  $PROP$  is the round-trip time from a source to the destination and back again. Therefore, it must be at least twice the one way propagation delay.
  - (b.) Otherwise, the signal would degrade too much along the wire making it difficult to detect collisions.
  -  (c.) The sender needs to unambiguously determine that a packet encountered a collision before it finishes transmitting the packet.
  - (d.) In any network (regardless of whether we use CSMA/CD or not) the transmission time of a packet is a function of both the data rate, and speed of propagation along the wire.

Name: \_\_\_\_\_

CS244a Final  
March 19, 2007


Student ID #: \_\_\_\_\_

Professor Nick McKeown

Campus/SITN-Local/SITN-Remote? \_\_\_\_\_


---

11. **Ethernet.** In an Ethernet network, which of the following are true:


- (a.) Ethernet switches (a.k.a bridges) learn addresses by looking at the destination address of packets as they pass by.
- (b.) Ethernet hubs and repeaters learn addresses by looking at the addresses of packets as they pass by.
- (c.) A correctly operating Ethernet switch never sends a packet to the wrong outgoing port.
-  (d.) Ethernet switches (a.k.a bridges) learn addresses by looking at the source address of packets as they pass by.
- (e.) Collisions occur less on a switched Ethernet network because links run faster.

**For questions 12-18.** It has been predicted that the capacity of the public Internet will grow significantly because of a rapid increase in the amount of video traffic delivered to people's homes. In the following questions, we'll make some crude assumptions to try and estimate how much total network capacity might be needed in the future.

12. Assuming that there are 100 million households in the US, and that each household has two HDTVs, and that each HDTV is used to watch 4 hours of video per day. We'll assume that each (compressed) video stream runs at 9Mb/s. If all households are watching TV at the same time, but all watching video on-demand (i.e. data is delivered unicast to every household), then which of the following is the closest approximation of the total peak aggregate data rate delivered to all homes?

- (a.) 100 Tb/s ( $100 \times 10^{12}$  b/s)
-  (b.) 2 Pb/s ( $2 \times 10^{15}$  b/s)
- (c.) 1 Pb/s
- (d.) 100 Gb/s
- (e.) 2 Tb/s

13. Now let's compare our answer above with an estimate of the future capacity needed to carry web traffic. We'll assume 100 million people download an average of 100 web pages per day, each containing 1Mbyte of data. If the peak traffic is five times the average, then which of the following is the closest estimate of the capacity needed by the network to deliver this traffic?

- (a.) 10Gb/s
- (b.) 500Gb/s
- (c.) 1 Tb/s
-  (d.) 5 Tb/s
- (e.) 10 Tb/s

Name: \_\_\_\_\_

CS244a Final  
March 19, 2007


Student ID #: \_\_\_\_\_

Professor Nick McKeown


Campus/SITN-Local/SITN-Remote? \_\_\_\_\_

---


14. Having (hopefully) convinced you that video traffic could be huge, let's look at a specific design example for a video distribution network. We'll consider what might happen if well-known DVD-by-mail company, Ixflent, were to deliver all their videos over the Internet instead. We'll work with today's numbers, even though – of course – we can expect the numbers to grow in time: More subscribers, and eventually all DVDs will be HD-quality. For now, we'll assume they ship 1 million DVDs per day, and each carries 10Gbits of video data. Approximately how many bits are delivered per day?

- (a.)  $10^{15}$
-  (b.)  $10^{16}$
- (c.)  $10^{17}$
- (d.)  $2 \times 10^{15}$
- (e.)  $2 \times 10^{14}$


15. Assume that all the videos are delivered from a single nationwide server, over a single connection to the public Internet. What is the approximate data rate of the connection if we assume that it is provisioned to carry five times the average data rate?

- (a.) 100Gb/s
- (b.) 1Gb/s
- (c.) 385Mb/s
- (d.) 1Tb/s
-  (e.) 500Gb/s

16. Let's estimate how much it costs Ixflent to send traffic to their customers each month. We'll assume they rent connections to two independent Tier-1 network operators (for redundancy) and provision the network to carry five times the average data rate. Both connections are large enough to carry all their traffic so that service is not interrupted when one network fails. We'll assume the network operators charge \$10 for a 1 Mb/s connection per month, and that the price scales linearly with the connection rate. Roughly how much does it cost Ixflent per year to deliver data?

- (a.) \$1M/yr
- (b.) \$8M/yr
-  (c.) \$140M/yr
- (d.) \$4B/yr
- (e.) \$1B/yr

17. About how much is that per delivered video?

-  (a.) \$0.40
- (b.) \$4
- (c.) \$0.04
- (d.) \$1.40
- (e.) \$14

Name: \_\_\_\_\_


CS244a Final  
March 19, 2007

Student ID #: \_\_\_\_\_

Professor Nick McKeown

Campus/SITN-Local/SITN-Remote? \_\_\_\_\_

18. We can think of the network as a spanning tree, with the video server at the root, and the one million subscribers at the leaves. Assuming that the tree has degree four (i.e. each router in the tree connects to four routers closer to the leaves), then roughly how many routers does a packet pass through from the root to each subscriber?

- (a.) 2
- (b.) 15
-  (c.) 10
- (d.) 20
- (e.) 40

## Longer questions

19. **(3 points) Fairness.** Three users are sharing a common link of 1 Mb/s. User  $A$  is downloading large files and is connected to the common link via a slow access link at  $x$  Mb/s, user  $B$  is connected via a 100 Mb/s link, but the application she is using requires at most  $x$  Mb/s, and finally user  $C$  is connected via a 1 Gb/s link and is downloading a movie that can take up any amount of bandwidth available to it. What is the max-min fair allocation for these three flows at the common link?

*Two Cases: 1) If  $x < 1/3$  Mb/s then  $A$  and  $B$  get  $x$  Mb/s,  $C$  gets  $(1-2x)$  Mb/s, otherwise  $A$ ,  $B$ , and  $C$  get  $1/3$  Mb/s*

20. **(18 points) Spoofing.** If we want to “break-in” to an ongoing TCP connection, and pretend to be one of the end-hosts, we need to send a sequence number that will convince the other end that we are legitimate. For example, imagine that a source host,  $A$ , tries to create a TCP connection with host  $B$ . When host  $A$  sends the SYN message, host  $C$  (masquerading as  $B$ ) sends a SYN+ACK message to host  $A$  with the correct sequence number. Even though  $C$  didn’t see the original SYN message, it gets lucky and guesses the right sequence number to send back so as to fool host  $A$ .

- (a.) If Host  $C$  were to guess the value from among all possible 32-bit sequence numbers, on average how many guesses would it take to get it right?

$$2^{31}$$

Name: \_\_\_\_\_

CS244a Final  
March 19, 2007

Student ID #: \_\_\_\_\_

Professor Nick McKeown

Campus/SITN-Local/SITN-Remote? \_\_\_\_\_

---

If Host C already knew 12 of the bits (e.g. because it knew something about the pseudo-random ISN generator), on average how many guesses would it take to get it right?

$$2^{19}$$

If Host C wanted to send a spoofed email to A, it would take just one “guessed” packet. Assuming the e-mail body and packet headers were 1500bytes, given a 1Gb/s link (and assuming that C already knew 12 bits of the sequence number), on average how long would it take to send one e-mail with a spoofed source?

$$2^{19} * 1500\text{bytes} * 8\text{bits/byte} * 1/10^9 = 6.3\text{seconds}$$

(b.) What simple change could you add to the SMTP protocol to defend against this attack?

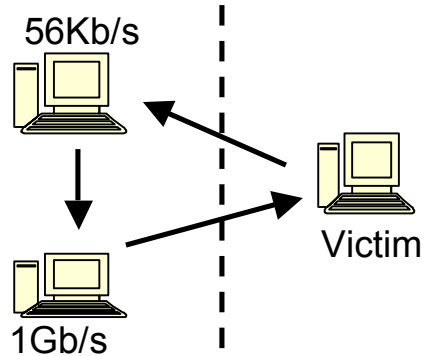
*Send an application level challenge and wait for a response before sending the message.*

*(Multiple answers were accepted for this. Solutions requiring secrets or keys were docked points because this is a public mail system.)*

(c.) (See Figure on next page). A variant of spoofing is used to protect a high-bandwidth (and thus valuable) machine by forging the source of a lower bandwidth (less valuable) machine also under the attacker’s control. The victim’s return packets are then forwarded by the low-bandwidth host back to the high-bandwidth host which then can retrieve the correct sequence number. Assuming the low bandwidth host is connected to a 56Kb/s link and the high bandwidth host has a 1Gb/s connection, how many messages per second can the high bandwidth server send (again use single packet e-mail messages of 1500bytes)?

*Acks on the 56k link are the bottleneck in this attack. Since an ACK is generally 64 bytes*

$$56k/(64*8\text{bits}) = 109\text{messages/sec}$$

**Attacker Controlled**

(d.) Does your fix for part (d) above also apply to this form of spoofing?

*In general, any solution that relied on the attacker not receiving the initial SYN|ACK does not work in this setup.*

**21. (20 points) Congestion Control.** Physicists at Stanford Linear Accelerator Center (SLAC) generate a few terabytes of data per day, which they would like to share with their colleagues in CERN, Geneva. SLAC and CERN are connected with 10 Gb/s links end-to-end and have a round-trip time of 200 milliseconds. They use standard TCP in transferring the bulk file, but quickly discover they are not able to achieve a sustained 10 Gb/s from end to end. They consult the local networking expert, Alice, who conjectures that the problem is the way in which TCP responds to losses in the network. Let's first find out if Alice's conjecture is true.

(a.) Assuming a simplified model of TCP (ignore Slow-Start and assume TCP is in the AIMD mode) derive an approximate expression for TCP's average window-size as a function of the loss-probability,  $p$ .

The area under the sawtooth is  $\left(\frac{1}{2} \times \frac{\hat{W}}{2} \times \frac{\hat{W}}{2}\right) + \left(\frac{\hat{W}}{2} \times \frac{\hat{W}}{2}\right) = \frac{3}{8} \hat{W}^2$ . The

probability that a packet will be dropped is  $\frac{1}{\frac{3}{8} \hat{W}^2}$  (1 drop per sawtooth).

Solving for  $\hat{W}$ , we see  $\hat{W} = \sqrt{\frac{8}{3p}}$ . The average window is equal to 3/4 of  $\hat{W}$ ,

which is  $\sqrt{\frac{3}{2p}}$ .

(b.) Using your answer to part (a), what is the average number of round-trip times between loss events that TCP can tolerate so as to sustain an average window-



size of  $w$  packets. Observe that the time interval between losses needs to increase as the "pipe" size increases.

Since the average window,  $w_a$ , is equal to  $3/4$  of  $\hat{W}$ , then we can see that  $\hat{W}$  is equal to  $4/3$  of  $w_a$ .  $\frac{\hat{W} \times RTT}{2}$  is the time interval between loss events.

Therefore, we need  $\frac{2}{3} w_a \times RTTs$  between loss events.

- (c.) Let's see what it takes to fill the pipe (i.e. achieve 10Gb/s). To achieve 10Gb/s, we need an average window size of  $10Gb/s \times 200ms \div 1500bytes$  (we are assuming the packets are 1500bytes long). How low a loss-rate do we need for this to happen? Comment on your answer.

The average window size is

$$\frac{10 \times 10^9 \times 200 \times 10^{-3}}{1500 \times 8} = \frac{5 \times 10^5}{3} = 166666.67 \text{ packets. Using the equation from}$$

part (a), we can see that  $\sqrt{\frac{3}{2p}} = 166666.67$ . Solving for  $p$ , we see that we need a  $5.4 \times 10^{-11}$  loss rate.

- (d.) Alice recommends Alice-TCP: an alternative to AIMD – and proposes that TCP be modified to use MIMD instead (Multiplicative Increase Multiplicative Decrease). She believes it will increase the throughput for bulk transfers. In particular, she proposes that when an acknowledgement is correctly received, Alice-TCP will increase the window size by additive constant  $a$  (i.e.  $w \rightarrow w+a$ , so that it increases by a *multiplicative* amount when the whole window has been acknowledged); when there is a loss, the window size is multiplied by factor  $(1-b)$ , where  $b < 1$ . Assuming that exactly one packet is lost each time the window size reaches  $\hat{W}$ , show that the average window size is given by:  $w = \left(\frac{2-b}{2}\right)\hat{W}$ .

The average window per cycle is equal to the sum of one half of the smallest window size and largest window size. The largest window equals  $\hat{W}$  and the smallest window is equal the drop factor multiplied by  $\hat{W}$ . Therefore,

$$w = \frac{1}{2} \left( (1-b)\hat{W} + \hat{W} \right) = \left( \frac{(1-b)\hat{W} + \hat{W}}{2} \right) = \left( \frac{2-b}{2} \right) \hat{W}.$$

- (e.) Alice argues that Alice-TCP is better. She says that if we want to sustain a particular average window size, the number of round-trip times between losses is independent of the pipe size. Is she correct? Explain your answer.

Show that for MIMD TCP, the following holds true:

$$pw_a = \frac{a}{b} \left(1 - \frac{b}{2}\right) =: p$$

So, the number of round-trips times between losses for a sustainable average window size,  $w_a$ , does not increase with the pipe size!


One way to do this would be to write a differential equation for rate change and equate it to 0 for the equilibrium.

$$\dot{w}_i = \frac{aw_i(t)}{T_i} - \frac{2b}{2-b} x_i(t) p_i(t) w_i(t)$$

where  $x_i(t) := w_i(t)/T_i$ . Now equate equilibrium rate change  $\dot{x}_i(t)$  to zero.

An alternative way to do it would be to follow the procedure like Question 12 of the 2004 final.

22. (16 points) **Routing.** In the following topology, assume all the links are bidirectional and the cost is the same for both directions.



(a.) Use the Bellman Ford algorithm to find the lowest cost paths to R1 for all the nodes. For each step before convergence, show the current lowest cost to reach R1 and the next hop router to get there for each node.

Note: You may not have to use the whole table.

Router	Hop	Dist	Hop	Dist	Hop	Dist	Hop	Dist	Hop	Dist
R <sub>2</sub>	-	∞	R <sub>1</sub>	5						

Name: \_\_\_\_\_

CS244a Final  
March 19, 2007

Student ID #: \_\_\_\_\_

Professor Nick McKeown

Campus/SITN-Local/SITN-Remote? \_\_\_\_\_

R <sub>3</sub>	-	∞	R <sub>1</sub>	10						
R <sub>4</sub>	-	∞			R <sub>2</sub>	15				
R <sub>5</sub>	-	∞			R <sub>3</sub>	20				
R <sub>6</sub>	-	∞					R <sub>4</sub>	20		
R <sub>7</sub>	-	∞			R <sub>3</sub>	35	R <sub>5</sub>	30		
R <sub>8</sub>	-	∞			R <sub>2</sub>	45	R <sub>6</sub>	39	R <sub>6</sub>	30
R <sub>9</sub>	-	∞					R <sub>7</sub>	40	R <sub>7</sub>	35

(b.) After the minimum cost spanning tree is formed, the link between R4 and R6 fails. For each step before re-convergence, and for each node, show the current lowest cost to reach R1 and the next hop router to get there. Assume split-horizon with poison reverse. Why do we have to assume it?

Note: You may not have to use the whole table.

Router	Hop	Dist	Hop	Dist	Hop	Dist	Hop	Dist	Hop	Dist
R <sub>2</sub>	R <sub>1</sub>	5								
R <sub>3</sub>	R <sub>1</sub>	10								
R <sub>4</sub>	R <sub>2</sub>	15								
R <sub>5</sub>	R <sub>3</sub>	20								
R <sub>6</sub>	R <sub>4</sub>	20	-	∞	R <sub>5</sub>	30				
R <sub>7</sub>	R <sub>5</sub>	30								
R <sub>8</sub>	R <sub>6</sub>	30			-	∞	R <sub>7</sub>	34		
R <sub>9</sub>	R <sub>7</sub>	35								

Name: \_\_\_\_\_

CS244a Final  
March 19, 2007

Student ID #: \_\_\_\_\_

Professor Nick McKeown

Campus/SITN-Local/SITN-Remote? \_\_\_\_\_

---

23. **(12 points) Leaky Bucket Constrained Traffic.** A router buffers arriving packets in separate queues – one per flow. Assume that the queues are all empty at time 0, and that the packets of each flow have a fixed size. Assume that flow  $A$ 's traffic is confined to a leaky-bucket constraint with  $\sigma = 5$  packets, and  $\rho = 10$  packets per second. The queues are all served using a WFQ scheduler, so as to meet an end-to-end delay guarantee.

(a.) What is the maximum backlog (in packets) at the queue corresponding to flow  $A$ ?

*5 packets*

*Packets of flow  $B$  follow a modified leaky-bucket model: The number of arrivals until time  $t$  is given by  $a(0, t) \leq 5 + \sqrt{t}$ .*

(b.) What is the maximum backlog at the queue corresponding to flow  $B$ ?

*5 packets*

*Packets of user  $C$  enter the router periodically. All we know is that in each 10 seconds, no more than 1000 packets arrive.*

(c.) Can this traffic be modeled as leaky-bucket traffic? If so, what are the corresponding  $\sigma$  and  $\rho$  parameters? If not, explain why.

*Yes.  $\sigma=1000$  and  $\rho=100$*

(d.) At a given time, assume packet  $p$  of flow  $A$  is scheduled to leave the router before packet  $q$  of flow  $B$ . Is it possible that a future arriving packet changes the departure order of these packets? Why?

*When any two packets have been scheduled, they have been placed in order of increasing Finishing Time (as determined by the WFQ scheduler). They then leave in order of Finishing Time. Therefore, once scheduled, their departure order is fixed.*

24. **(9 points) Leaky-bucket constrained traffic.** In this question we'll assume that flows consist of a stream of bits rather than packets. In class, we saw how a leaky bucket regulator can constrain a flow so that the number of bits sent in an interval of duration  $\tau$  is given by:  $A(t, t + \tau) \leq \sigma + \rho\tau$ . If you think about it, this leaky bucket would allow a burst of up to  $\sigma$  bits to depart from the source in zero time (i.e. at an infinite rate)! This is clearly not possible; so in practice, the source is constrained by the data rate of its outgoing link. We'll assume that the flow is constrained by a leaky bucket at the source, and that a burst can depart no faster than  $R$  bits/second (the data rate of the link to which the source is connected). The figure below shows how the traffic is constrained. The flow passes through ten routers, all of which use bit-by-bit weighted fair queuing to serve

Name: \_\_\_\_\_

CS244a Final  
March 19, 2007


Student ID #: \_\_\_\_\_

Professor Nick McKeown

Campus/SITN-Local/SITN-Remote? \_\_\_\_\_

---

multiple flows.



- (a.) If the first router allocates the flow a service rate greater than or equal to  $c$  (where  $\rho \leq c$ ), write down expressions for the maximum delay of a bit in the router, and the size of the buffer needed to prevent overflows. (Write your answers in terms of  $\sigma$ ,  $\rho$  and  $R$ , but not  $c$ ).

$$\text{Maximum buffer size} = \sigma - \frac{\sigma\rho}{R}, \text{ Delay} = \frac{\sigma - \frac{\sigma\rho}{R}}{\rho} = \frac{\sigma}{\rho} - \frac{\sigma}{R}$$

- (b.) What happens to your answer in (a) when  $R = \rho$ ? Explain your answer.

*No buffer is needed*

- (c.) Why do you think it is customary to assume that a burst can leave a leaky bucket regulator at infinite rate?

*R is usually much greater than rho so the answer changes very little when using R.*