CS144: Intro to Computer Networks
Practice Midterm
Answer Key Provided Tuesday, Nov 1st

Your Name: _____

SUNet ID: _____ @stanford.edu

In accordance with both the letter and the spirit of the Stanford Honor Code, I will not distribute this homework to future classes.

Signature: _____

- Some questions may be much harder than others.

- All questions require you to justify your answer to receive full credit, even multiple choice questions for which you circle the correct answer(s).

- Keep your answers concise. We will deduct points for a correct answer that also includes incorrect or irrelevant information.

1	/10
2	/5
3	/5
4	/5
5	/5
6	/15
7	/10
8	/5
9	/10
10	/10
Total	/80

I Stop-n-Wait-n-Answer

1. **[10 points]:**

   An application needs to send 100KB of data using a stop-and-wait reliable protocol. The protocol splits the data into segments that have a 1KB application data payload. Each segment fits in a single IP packet. The RTT is 50ms, there is no packetization delay, and no queueing delay. The protocol uses a fixed retransmission timeout of 200ms and has no retransmission limit.

   (i) How long will the transmission take, in seconds, if the network does not drop, duplicate or corrupt any packets? You may assume the connection is established when you start your measurement, so there is no additional latency from connection setup and consists only of data transmissions. Your answer must be accurate to two decimal points (10ms).

   _____ seconds

   (ii) Let us now suppose that the network drops each segment with a probability of 10%, independently from segment to segment. The network drops both data and acknowledgements. What is the *expected* duration of the transmission? Show your calculations. Your answer must be accurate to two decimal points (10ms).

   _____ seconds

II   Without TCP The Internet Ain't Nuthin'

There are devices and services in the Internet, such as proxy servers, that "split" TCP connections. Suppose a host A wants to open a connection to a host C. A device somewhere along the path, B, can terminate A's connection at itself, and open a connection to C. So in this case there are now two TCP connections, A to B and B to C. A thinks it's sending data to C, but B is processing the TCP segments itself and sending acknowledgments back to A, spoofed from B's IP address. Simultaneously, B opens a TCP connection to C, pretending to be A.
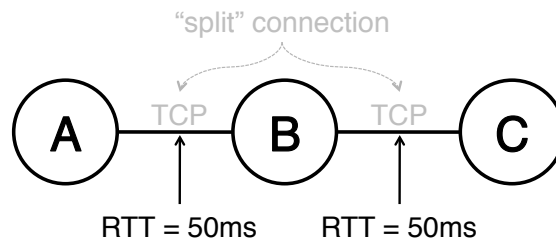
Suppose you have the network above, where the RTT from A to B is 50ms, the RTT from B to C is 50ms, and there is no packetization, queueing, or processing delay, such that the RTT from A to C is 100ms. The maximum segment size is 1400 bytes. A is sending an infinite stream of bytes, such that every segment is the maximum segment size. Recall that a TCP flow's throughput can be approximated as

$$\text{MSS} \cdot \sqrt{\frac{3}{2}} \cdot \frac{1}{\text{RTT}\sqrt{p}}$$

where p is the packet drop rate.

Please write out answers numerically and do not leave radicals or variables in your solutions. You may leave fractions. If you do not have a calculator, you may approximate with the following values:

$MSS \cdot \sqrt{\frac{3}{2}} = 13,717$ bits

$\sqrt{0.1} = 0.32$

$\sqrt{0.19} = 0.44$

$\sqrt{0.2} = 0.45$

$\sqrt{0.21} = 0.46$

2. **[5 points]:**

   Suppose that B does not split the TCP connection, such that packets flow directly from A to C, through B. The route between A and B drops 10% of data segments and does not drop acknowledgments, while the route between B and C does not drop any packet. What will the TCP throughput from A to C be?

   _____ kbps

3. **[5 points]:**

   Suppose that B does split the connection, such that packets flow from A to B, terminate at B, then are forwarded in separate flow from B to C. The route between A and B drops 10% of data segements and drops no acknowledgments, while the route between B and C does not drop any packet. What will the throughput from A to C be?

   _____ kbps

4. **[5 points]:**

   Suppose that B does split the connection, such that packets flow from A to B, terminate at B, then are forwarded in separate flow from B to C. The route between A and B drops 10% of packets, and the route between B and C also drops 10% of packets. What will the throughput from A to C be?

   _____ kbps

5. **[5 points]:**

   Finally, suppose that B does not split the connection, such that packets flow from A to B, passing through but not terminating at B. The route between A and B drops 10% of data segments, and the route between B and C also drops 10% of data segments. What will the throughput from A to C be?

   _____ kbps

III Put It All Together, Now

6. **[15 points]:**

   You type the following URL into your web browser:

   `http://gradadmissions.stanford.edu/inquiry/onlineinq.htm`
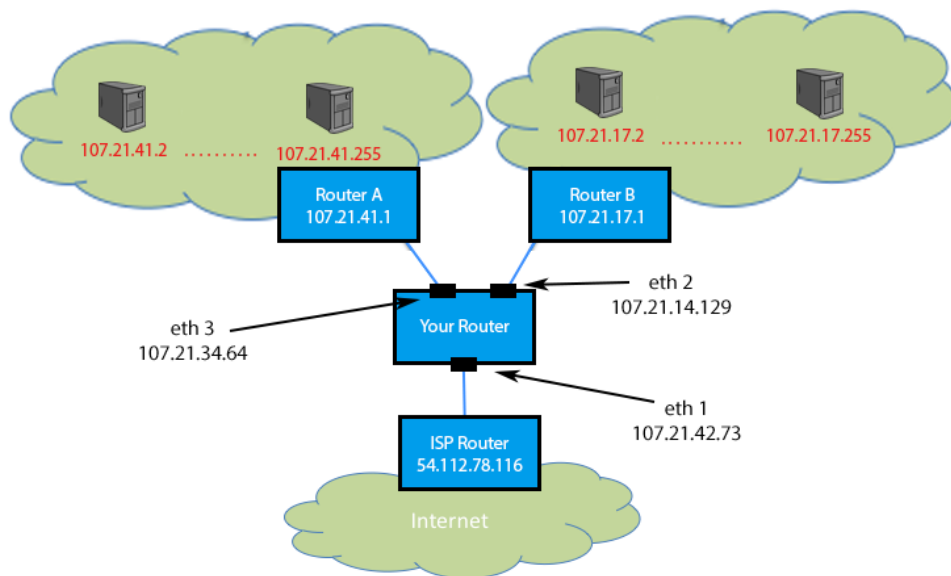
   Assuming that

   - your DNS resolver is 171.64.7.77,
   - neither your host nor your DNS resolver have any cached DNS entries,
   - DNS never needs to fail over to TCP, and
   - the HTML response returns 200 OK with a web page,
   - the HTML request and response each fit in a single segment, and
   - the web page requires loading no additional resources,

   write down the series of packet exchanges that will occur for your host to receive the web page. Include packets sent by your DNS server as well as control packets for TCP connection setup and teardown. You need not include any ARP packets, and you do not need to write down message formats. Simple descriptions such as "X sends a UDP segment to the HTML server on the HTTP port" are sufficient. In the case of the HTTP request, clearly state the path of the file requested in the GET.

IV   Here, There, Anywhere

7. **[10 points]:**

A web site hosted on a single server becomes extremely popular. The administrators decide to replace the server with two sets of 200 servers (i.e. 400 servers in total), each set of servers connected to the Internet via a different router (for fault tolerance). The figure shows the topology. The two routers (Router A and Router B) connect to the Internet via a router that you manage ("Your Router"). The administrator of Router A assigns 200 host IP addresses in the range `107.21.41.2–107.21.41.255`, and the administrator of Router B assigns 200 host IP addresses in the range `107.21.17.2–107.21.17.255`. You decide to manually insert routing table entries into "Your Router" to correctly route packets between the servers and the rest of the Internet.

107.21.41.2 .... 107.21.41.255

107.21.17.2 .... 107.21.17.255

Router A
107.21.41.1

Router B
107.21.17.1

eth 2
107.21.14.129

eth 3
107.21.34.64

Your Router

eth 1
107.21.42.73

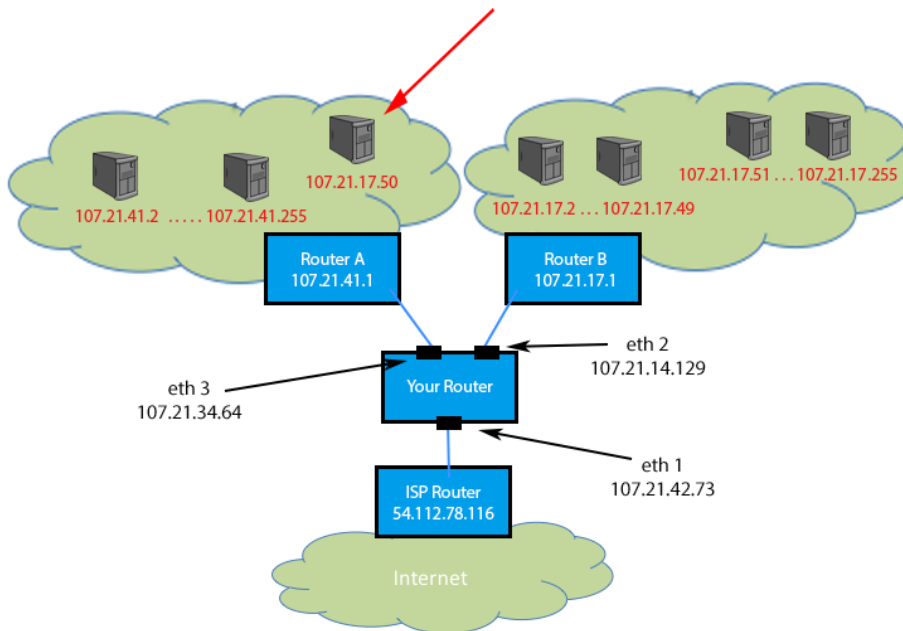ISP Router
54.112.78.116

Internet

Write down the routing table entries for "Your Router" to correctly route packets between the Internet and the four hundred servers. You need not include routes to "Your Router"'s interfaces. Use as few table entries as possible.

destination prefix (e.g. 10.0.0.0)	net mask (e.g. 255.0.0.0)	next hop (e.g. 10.0.0.1)	interface (e.g. eth0)

8. **[5 points]:**

The administrators decide to move the server with IP address `107.21.17.50` from the network behind Router B to the network behind Router A. Assuming that they change no server IP addresses (the moved server keeps the *same* IP address), how do you need to modify the routing tables of "Your Router" to correctly route packets between all servers and the Internet? You may assume the administrators will update Router A and Router B properly to support your change. The new topology looks like this:

107.21.17.50

107.21.41.2 ..... 107.21.41.255

107.21.17.51 ... 107.21.17.255

107.21.17.2 ... 107.21.17.49

Router A
107.21.41.1

Router B
107.21.17.1

eth 2
107.21.14.129

eth 3
107.21.34.64

Your Router

eth 1
107.21.42.73

ISP Router
54.112.78.116

Internet

Write down any new or modified entries you need:

destination prefix	net mask	next hop	interface

Explain your answer:

V Playback Buffers and Queues

9. **[10 points]:**

Two endpoints in a voice-over-IP session are connected by a path of 4 routers. All links are running at 1Mb/s and the hosts are separated by 3000km. All packets are of size 1500Bytes. Assume the bit propagation speed is $2\text{x}10^8$ m/s. Note that 1KB of data is 1024Bytes, but 1Mb/s is 10^6 bits/s.

(i) What is the minimum round trip time (RTT), assuming there is no queueing delay and assuming processing time at each host is negligible?

_____ milliseconds

(ii) For this part only let us assume that one router on the path has a steady queue occupancy of 5 packets. What is the end-to-end delay (one way, not round trip) in this case?

_____ milliseconds

(iii) Now let us assume the maximum queue occupancy for *every* router queues is 5 packets. What is the maximum end-to-end delay?

_____ milliseconds.

(iv) For part (iii), how long should the playback buffer be at the destination voice-over-IP client be if each packet arrives successfully, without being dropped? Express your answer in bytes.

_____ bytes.

(v) For part (iii), how long should the playback buffer be at the destination voice-over-IP client if each packet either arrives successfully when first transmitted, or is dropped the first time and then arrives successfully on the second attempt? Assume the sender retransmits a packet only after it finds out that the packet was dropped, which is after twice the maximum end to end delay. Express your answer in bits.

_____ bits.

10. **[10 points]:**

Recall that a (σ, ρ) leaky-bucket traffic regulator ensures that the number of bits departing the regulator in any interval of time (t_1, t_2) is bounded by $B(t_1, t_2) \le \sigma + \rho(t_2 - t_1)$, where σ and ρ are non-negative constants.

Now, consider a queue which drains at rate R. The queue has an infinite buffer size. Answer the following questions. **Although it's a multiple choice question, don't forget to show your work below!**

(i) Which of the following is correct?
   Circle all that apply. (There may be multiple answers.)

   A If the source feeding the queue is leaky-bucket constrained, then the delay of a bit through the queue is always less than 2σ, regardless of the source's average sending rate, ρ.

   B The delay of any bit through the queue is finite only when $\sigma < R$.

   C If several leaky-bucket constrained sources (σ_i, ρ_i) feed the queue, and when combined $\sum_i \sigma_i = \sigma$ and $\sum_i \rho_i = \rho < R$, then the delay through the queue will be bounded and finite. Assume σ is finite.

   D If several leaky-bucket constrained sources (σ_i, ρ_i) feed the queue, and $\sum_i \sigma_i = 10$, $\sum_i \rho_i = 500\text{b/s}$, and $R = 600\text{b/s}$, then no bit will ever experience a delay through the queue of more than 100ms.
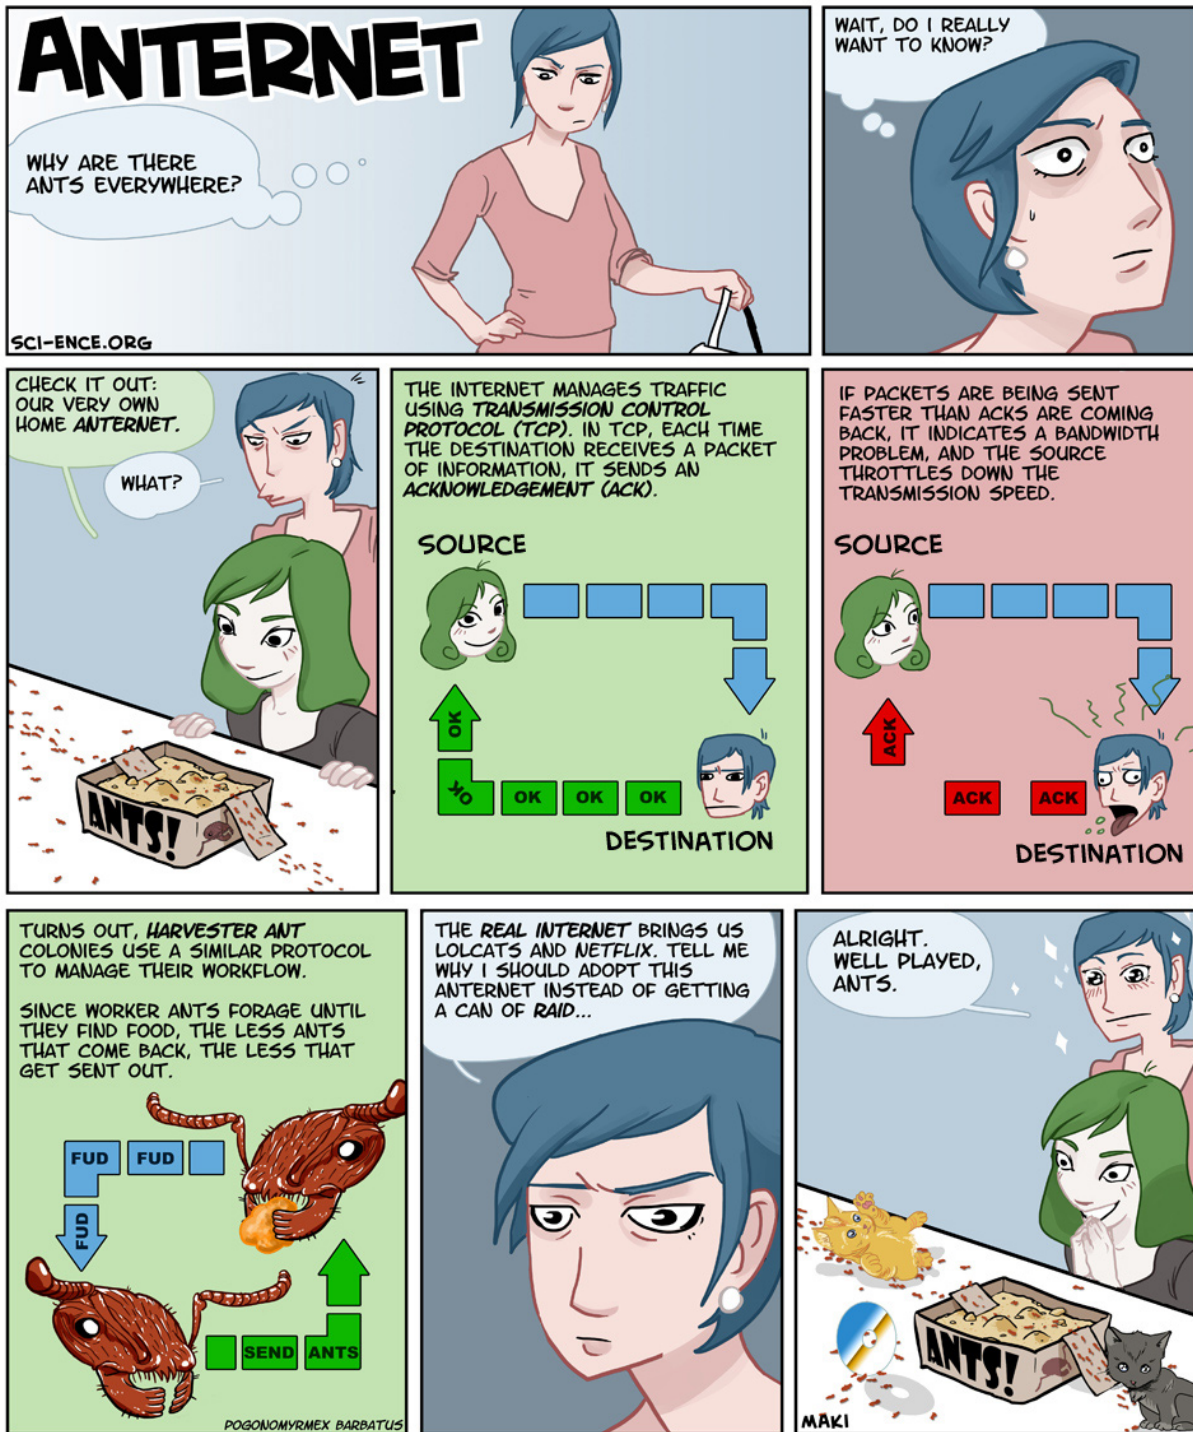
Figure 1: Cartoon based off research from Stanford by student Katherine N. Dektar, prof. Balaji Prabhakar (EE/CS), and prof. Deborah M. Gordon (biologist).