Name_______________ Student ID____________________ Department/Year___________

Midterm Examination

Introduction to Computer Networks
Class#: EE 4020, Session#: 01
Fall 2019

10:20-12:10 Thursday
November 14, 2019

Prohibited

1. You are not allowed to write down the answers using pencils. Use only black- or blue-inked pens.
2. You are not allowed to read books or any references not on the question sheets.
3. You are not allowed to use calculators or electronic devices in any form.
4. You are not allowed to use your own sheets of papers. Ask the exam administrators for extra blank sheets
5. You are not allowed to have any oral, visual, gesture exchange about the exam questions or answers during the exam.

Cautions

1. Check if you get 20 pages (including this title page), 9 questions.
2. Write your **name in Chinese**, student ID, and department/year down on top of the **first** page.
3. There are in total 100 points to earn. You have 100 minutes to answer the questions. Skim through all questions and start from the questions you are more confident with.
4. Use only English to answer the questions. Misspelling and grammar errors will be tolerated, but you want to make sure with these errors your answers will still make sense.
5. If you have any extra-exam emergency or problem regarding the exam questions, raise your hand quietly. The exam administrator will approach you and deal with the problem.

1. (Overview) Consider a micro-Internet consisting of 5 subnets, namely V, W, X, Y and Z as labeled below.

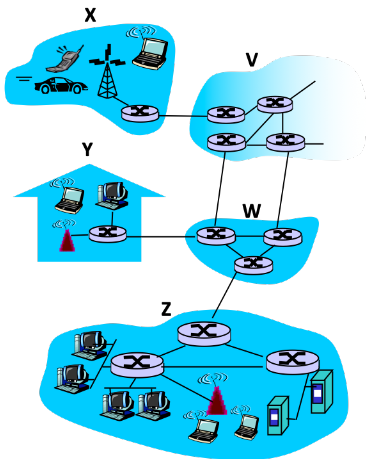

(1) Which of the subnets is classified as an institution network in the lecture? (2%)
(2) Which of the subnets is classified as a residential network in the lecture? (2%)
(3) Which of the subnets is classified as a mobile network in the lecture? (2%)

Sample Solution:

(1) Z
(2) Y
(3) X

2. (Overview) Consider the 3 figures below that highlight different parts of the micro-Internet from Question 1.

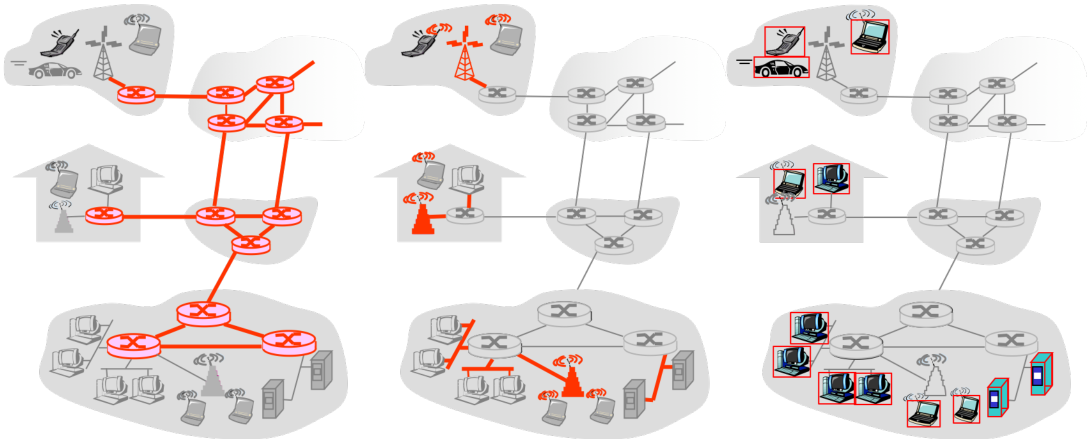

(a) (b) (c)

(1) Which highlights the core part of the micro-Internet (2%)?
(2) Which highlights the edge part of the micro-Internet (2%)?
(3) Which highlights the access network part of the micro-Internet (2%)?

Sample Solution:

(1) (a)
(2) (c)
(3) (b)

3. (Overview) Based on your understanding of packet switching and circuit switching principle, select the keywords that fit better the characteristics of a circuit switching network. (6%)
   (a) Contention
   (b) Congestion
   (c) Idle resource
   (d) Reservation
   (e) Call setup
   (f) Delay guarantee

Sample Solution:
(c)(d)(e)(f)

4. (Overview) ping is a commonly used network testing tool. What it does is to send a probing message periodically to a remote host. The remote host, upon receiving the probing message, sends back a reply message. From the reply message, the probing host calculates the round-trip time (RTT) between the two. By tracking the number of probing messages sent and the reply messages received, the probing host calculates the loss rate as well.

   Below is the result of a ping test to www.eurocom.fr from Polly's office desktop PC. The command is interrupted manually at the 196th test. One can see that the RTT for the 1st probe is 409.086 ms. One can also see in the ping statistics that the average RTT is 345.159 ms.

```
pollyss-iMac:~ pollyoffice$ ping www.eurocom.fr
PING eurocom.fr (217.70.191.35): 56 data bytes
64 bytes from 217.70.191.35: icmp_seq=0 ttl=40 time=409.086 ms
64 bytes from 217.70.191.35: icmp_seq=1 ttl=40 time=428.191 ms
…
64 bytes from 217.70.191.35: icmp_seq=195 ttl=40 time=315.882 ms
64 bytes from 217.70.191.35: icmp_seq=196 ttl=40 time=328.620 ms
^C
--- eurocom.fr ping statistics ---
197 packets transmitted, 195 packets received, 1.0% packet loss
round-trip min/avg/max/stddev = 281.737/345.159/468.015/47.067 ms
pollyss-iMac:~ pollyoffice$
```

   (1) Find the (a) packet loss rate, the (b) minimum, (c) maximum, and (d) standard deviation of the RTT. (4%)
   (2) The RTT observed by the 2nd probe is longer than the 1st probe. State at least 1 reason that might have contributed to the difference. (2%)
   (3) Do you think 217.70.191.35 is in Taiwan? Why or why not? (2%)
   (4) If you need to estimate the average RTT from your host to another host, would you prefer to use traceroute or ping? And why? (4%)
   (5) If you need to find out which router on the way from your host to another host contributes the most to the RTT, would you prefer to use traceroute or ping? And why? (4%)

Sample Solution:

(1) 1.0%, 281.737 ms, 468.015 ms, 47.067 ms
(2) An increasing amount of cross traffic, a longer traversal route, or any reason that makes sense
(3) Take your pick and justify
(4) As long as the pick and justification make sense
(5) As long as the pick and justification make sense

5. (Application) As shown in the figure below, the requesting host is asking for the IP address of www.amazon.com using an iterative DNS query. In particular, the local DNS server did not cache TLD DNS server, nor previous hostname-IP answers, so the delay will be the time to pass a DNS message (query or reply) through all steps of a recursive query. Let's say the time to pass a query or reply message through link i is ti. The overall DNS query delay without caching is then t1+t2+t3+t4+t5+t6+t7+t8.

   Now consider 2 caching strategies at the local DNS: (a) Caching the TLD DNS servers' IP addresses vs. (b) Caching the previous hostname-IP answers.

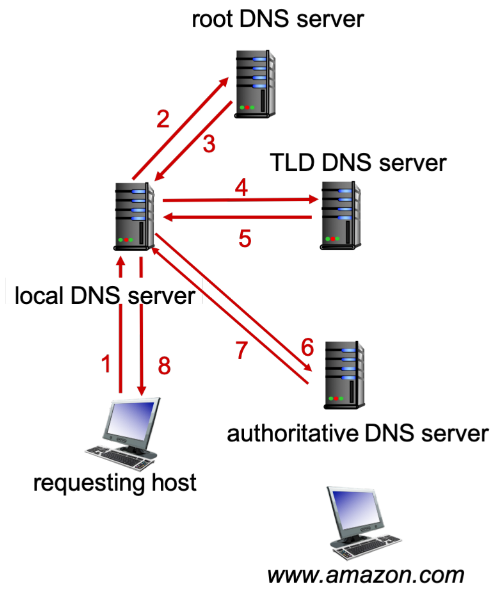

   (1) If the local DNS server implements just strategy (a), what is the delay the 2nd time the requesting host querying the IP of www.amazon.com? (2%)
   (2) Continue from (1). What is the delay querying www.microsoft.com next? (2%)
   (3) If the local DNS server implements just strategy (b), what is the delay the 2nd time the requesting host querying the IP of www.amazon.com? (2%)
   (4) Continue from (3). What is the delay querying www.microsoft.com next? (2%)
   (5) Suppose the requesting host queries often IP addresses of various something.com. Which caching strategy would you recommend? And Why? (4%)
   (6) Suppose the requesting host queries often IP addresses of the same something.com. Which caching strategy would you recommend? And Why? (4%)

Sample Solution:

(1) t1+t4+t5+t6+t7+t8
(2) t1+t4+t5+t6+t7+t8
(3) t1+t8
(4) t1+t2+t3+t4+t5+t6+t7+t8
(5) Take your pick and justify
(6) Take your pick and justify

6. (Application) Below is approximately how Netflix's CDN looks like. In the scenario, 7 of the nodes in the CDN store the movie that the customer is looking for. The client program implements (a) a CDN node selection algorithm and (b) the DASH protocol. In (a), the algorithm selects the nearest CDN node. In (b), the client program selects the best stream whose bitrate is within the available bandwidth to the selected node.

   Suppose the measured delay from the customer's host to node A, B, C, D, E, F, G are 10, 40, 60, 80, 50, 30, 20 ms respectively. The measured available bandwidth from node A, B, C, D, E, F, G to the customer's host are 1, 10, 25, 30, 20, 7, 4 Mbps respectively. The movie is encoded into 4 resolutions 1080p, 720p, 480p, and 360p whose bitrates are 12, 5, 2.5, and 0.8 Mbps respectively.

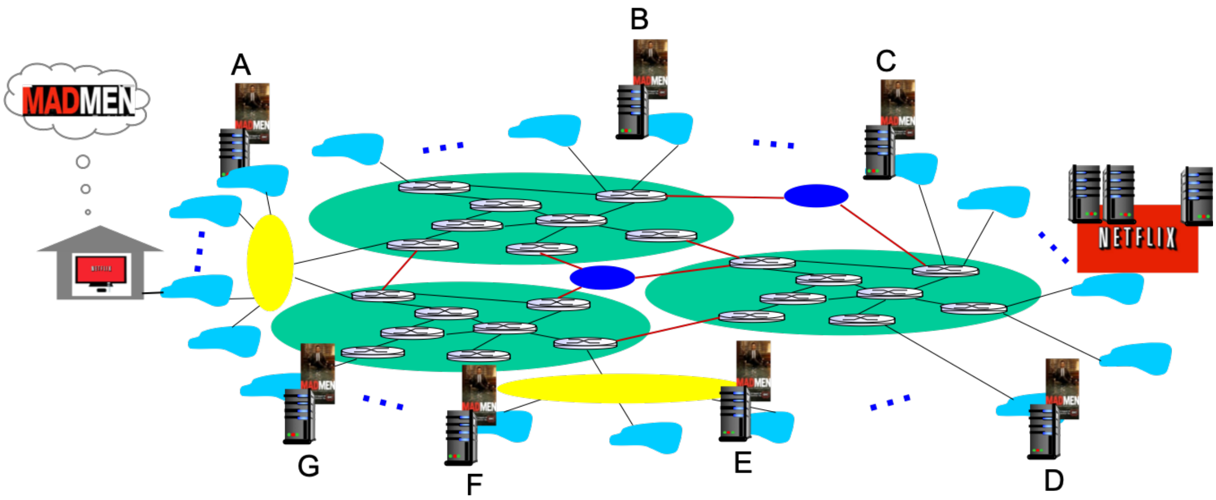

   (1) If the client program runs exactly the design of (a) and (b) above, which node will be selected and which resolution video will be streamed? (4%)
   (2) Suppose the client does not mind an initial RTT delay for up to 40 ms and therefore we can relax the node selection strategy to include all nodes with RTT <= 40ms. Which CDN nodes will remain in consideration using the relaxed node selection strategy? (2%)
   (3) Continue from (2). Suppose DASH is also generalized to select among multiple nodes such that the best resolution video within the bandwidth constraint is streamed. Which resolution video will be streamed and from which node? (4%)
   (4) Continue from (3). If the client does not mind an initial RTT delay for up to 100 ms, which resolution video will be streamed and from which node? (4%)

Sample Solution:

(1) Node A. 360p.
(2) If the RTT limit is 40ms, only node A and G are considered.
(3) Node G. 480p. A's got 1Mbps and G 4Mbps available bandwidth. Best resolution downloadable from A is 360p and from G 480p. So, 480p video will be streamed from G.
(4) Node E 1080p. A, B, E, F, G are considered. A can stream 360p the best, B 720p, E 1080p, F 720p, G 480p. So, the result should be 1080p streamed from E.

7. (Transport) The two figures below illustrate the packet interaction of a connectionless demux and a connection-oriented demux.

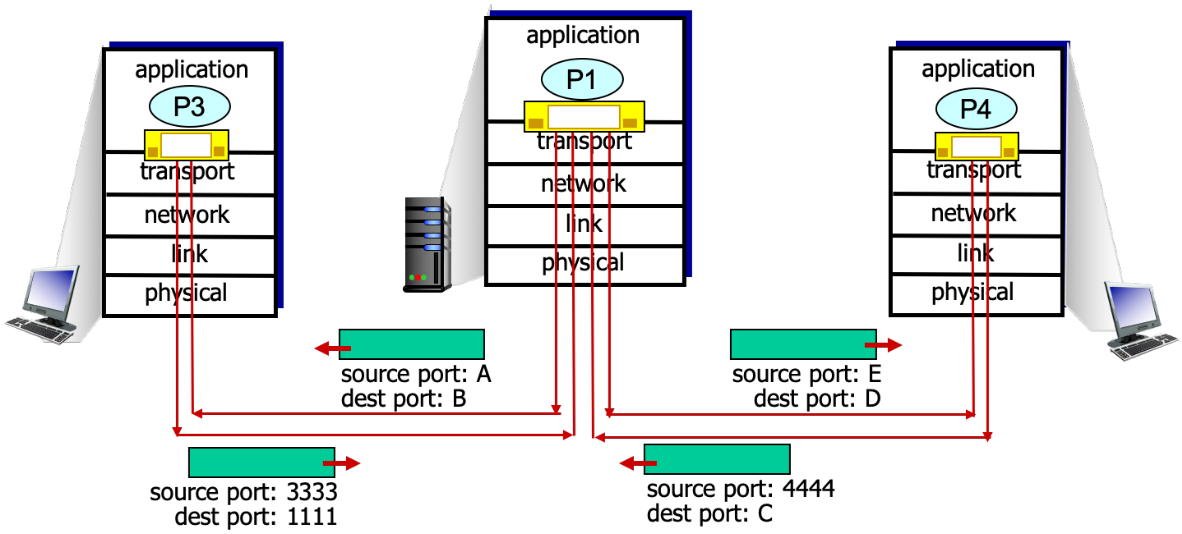

(a) connectionless demux

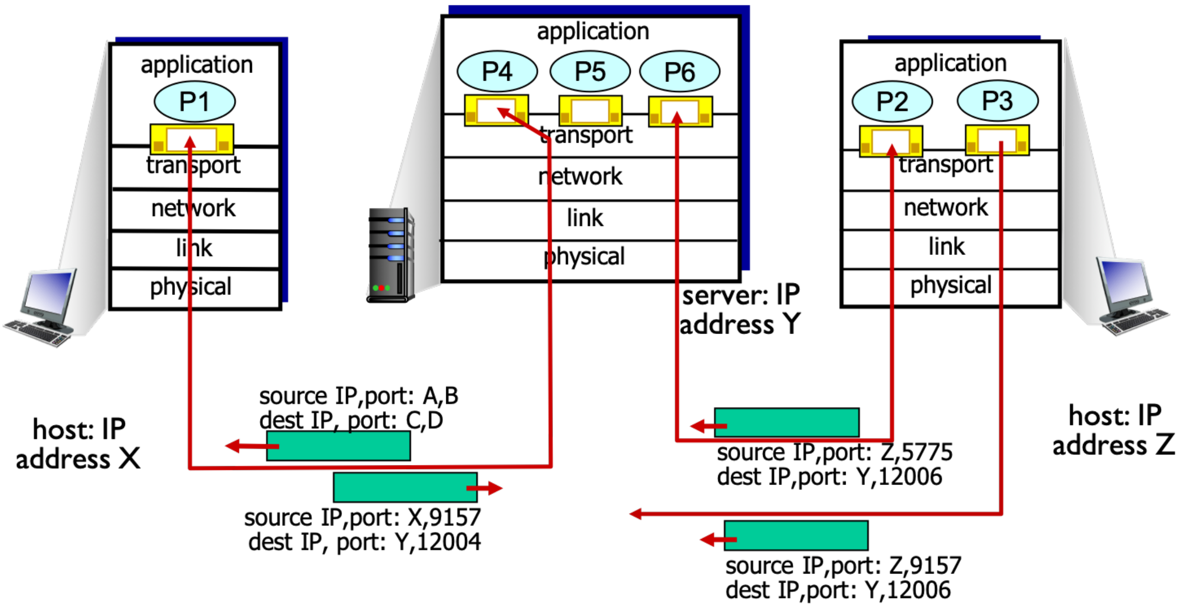

(b) connection-oriented demux

(1) What should be the value of A, B, C, D in Figure (a)? (2%)
(2) What should be the value of A, B, C, D in Figure (b)? (2%)
(3) Which process would the packet coming from P3 go to? And why? (4%)

Sample Solution:

(1) 1111, 3333, 1111, 4444
(2) Y, 12004, X, 9157
(3) None. The packet with (dest IP, port) : (Y,12006) is supposed to go to P6 as the packet coming from P2 carrying the same (dest IP, port) goes to P6. However, P6, running the connection-oriented demux, accepts packets only when the (src IP, port) also match. Given that P6's (src IP, port) are already (Z,5775) due to the connection to P2, packets coming from (Z:9157) won't match (Z:5775), and will therefore be discarded.

    Alternatively, P6 for it being the 2nd type of connection-oriented demux.

8. (Transport) Provided below are the FSMs of rdt 3.0 sender and receiver. Indicate the order of the transitions (in terms of t1, t2, …, t14) taking place for the following scenarios.

<u>rdt 3.0 sender:</u>

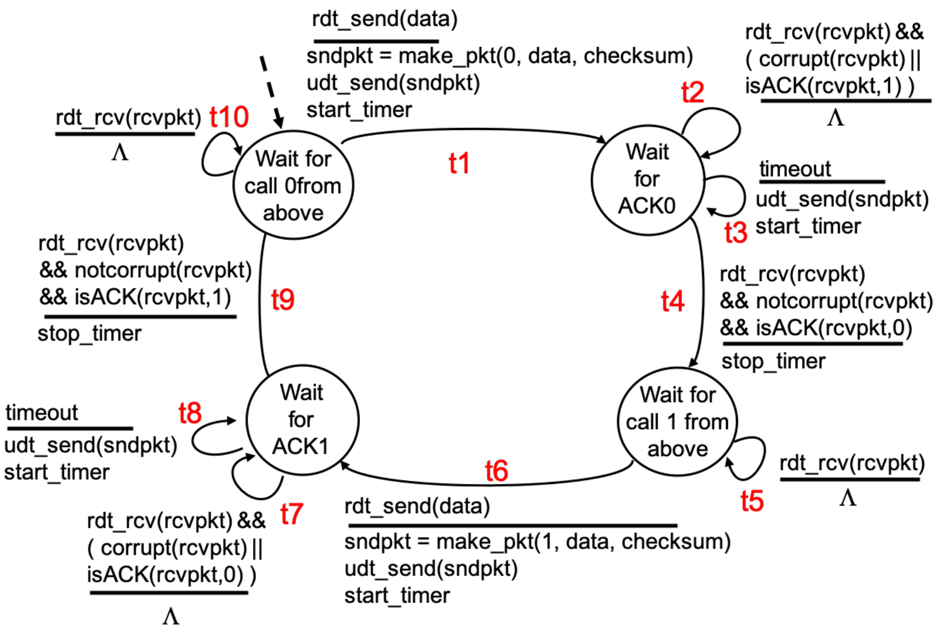

<u>rdt 3.0 receiver:</u>

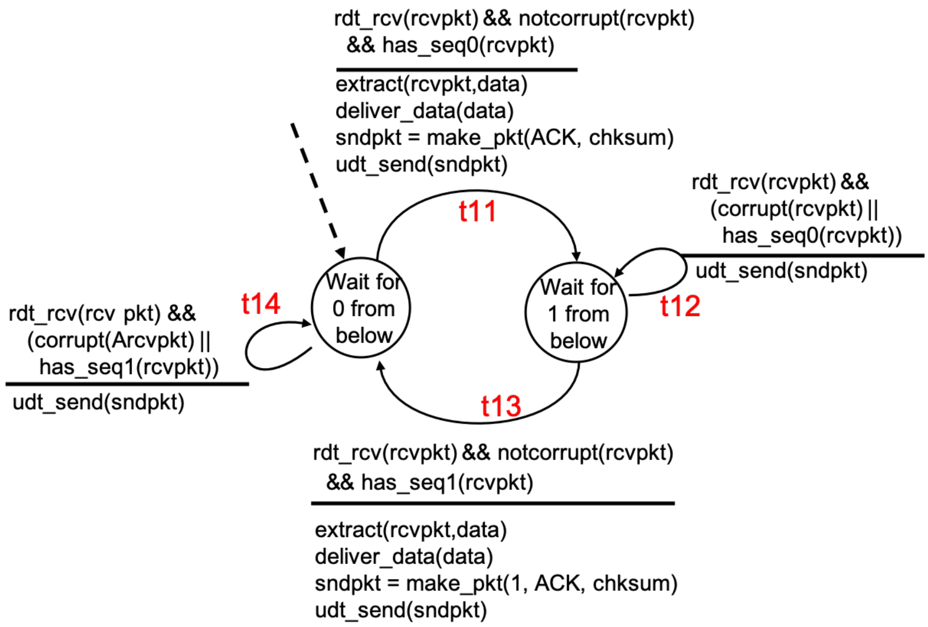

(1) Both sender and receiver start from the initial state. The sender gets one call from above and there is no bit error. (2%)
(2) Both sender and receiver start from the initial state. The sender gets one call from above and there is a bit error in the data packet. (4%).
(3) Both sender and receiver start from the initial state. The sender gets one call from above and there is a bit error in the ACK packet coming back. (4%)
(4) Both sender and receiver start from the initial state. The sender gets a call from above and the ACK packet does not come back at all. (4%)
(5) Extend rdt 3.0 such that the retransmission delay can be lowered and explain why. (4%)

Sample Solution:

(1) t1, t11, t4
(2) t1, t14, t2, t3, t11, t4
(3) t1, t11, t2, t3, t12, t4
(4) t1, t11, t3, t12, t4 (alternatively t1, t3…; alternatively t1, t14…)
(5) Suggest one and justify your design

9. `(Transport) Compare and contrast Go-Back-N (GBN) and Selective Repeat (SR).
    (1) Which method uses more network bandwidth to retransmit for losses? Why? (2%)
    (2) Which method requires more memory space at the receiver end? Why? (2%)
    (3) Which method requires more timers at the sender end? Why? (2%)
    (4) If the network bandwidth is abundant and memory is cheap which method would you choose to implement? Why? (4%)

Sample Solution:

(1) GBN. In GBN, if any of the packet in a flight is lost, the timer will eventually expire and trigger retransmission of a whole flight of packets in the worst case. In SR, however, if a packet is lost, only the packet is retransmitted.
(2) SR. The receiver needs to buffer packets arriving out of order. The worst case is N-1 packets (for the most unfortunate case where the 1^{st} packet in a flight is lost).
(3) SR. Each packet transmitted and not yet ack'd is tracked by a timer. For a window N, a fixed number of N timers will be required. GBN, on the other hand, needs only 1 timer.
(4) Pick and justify for your choice.