

Midterm

Srirakumar Balasubramanian, Zheng Cang, Li Yan (TAs), A. Chaintreau (instructor)

Please indicate your uni: _____

- This exam should be answered in 75mn. It is graded on a total of 150pt.

General Review ($12 \times 10 = 120$ pt)

For each question, you have to circle all the correct answer(s) (there may be one, several or even zero) and cross all the incorrect answer(s). If you believe these answers are obvious you may not need to justify it, but for any question that may require justification, you can clarify your answer **on the answer sheet!** (Be brief, no more than a couple of sentences).

1. What happens in CSEE 4119 when you ask a question?
 - (a) you receive a fruit (most often, an apple) if that's your first question asked today.
 - (b) you receive a free iPad.
 - (c) you receive a fruit only if you ask the question in French.
 - (d) you receive a fruit only if you are an undergraduate student.

Explanation: _____

2. The following are advantages of using a packet switched architecture.
 - (a) All bits are grouped into packets, so no information gets lost.
 - (b) The transmission delay depends on the packet size but not on the number of routers that relay the information to the destination.
 - (c) When users send information sporadically, a larger population can be served by the same link.
 - (d) Links and routers with various speeds can relay information.

Explanation: _____

3. The following are direct immediate advantages of organizing network into layer
- (a) Because layers work independently, it's easier to make sure that all operations are secure.
 - (b) It is guaranteed that the minimum overhead information is added to the actual useful data being sent.
 - (c) The network edge (laptop, cellphones) implements no physical and no link layer protocol.
 - (d) A wireless links can divide its bandwidth equally among potential communication pairs.

Explanation: _____

4. Assume $N \geq 50$ end-hosts are connected via a 20Mbps shared Ethernet link (whenever one uses it, the medium can't be used by any other host in the LAN), which is connected to the rest of the Internet by a 1Mbps link. Under the most optimistic scenario, what is the maximum gain of installing a web-cache inside the LAN.
- (a) Instead of all hosts simultaneously receive information at 1Mbps, they all receive it at 20Mbps.
 - (b) The end hosts can simultaneously receive at a rate 20 times larger than without cache.
 - (c) The end hosts can simultaneously receive at a rate 19 times larger than without cache.
 - (d) The end hosts can simultaneously receive at a rate $N - 1$ times larger than without cache.

Explanation: _____

5. Under the most optimistic scenario, what is the maximum gain that installing a web cache on the LAN can provide, if you now assume $N = 10$.
- (a) Instead of all hosts simultaneously receive information at 1Mbps, they all receive it at 20Mbps.
 - (b) The end hosts can simultaneously receive at a rate 20 times larger than without cache.
 - (c) The end hosts can simultaneously receive at a rate 19 times larger than without cache.
 - (d) The end hosts can simultaneously receive at a rate $N - 1$ times larger than without cache.

Explanation: _____

6. In which of the following situations introducing shortcuts within a DHT helps?
- (a) You have a very large group and you want to make sure that every query is answered quickly.
 - (b) Nodes frequently leave and you need to continue to operate without reconfiguring the whole DHT.
 - (c) You have a lot of information to store and you want to ensure the all peers have an equal number of objects to store.

(d) You run out of key values and you want to include new peers in the groups.

Explanation: _____

7. Why should a server bind to a port number whereas it is not required for a client?

Explanation: _____

8. Detail how many sockets are simultaneously open in every machine when 10 clients connect to a webserver using persistent HTTP to request 5 objects each, and the webserver answers those queries by sending web-objects.

Explanation: _____

9. Same question for 10 clients connect through 2 different applications to a local DNS server.

Explanation: _____

10. An adversary attempts at conducting a Denial of Service attack on a server by generating a lot of requests. Knowing that an adversary is trying to infiltrate, the server maintains a white list of IP addresses. Unfortunately, the adversary can spoof the origin IP address and hence poses as a legitimate user. As a consequence, an additional secure initialization step is ensured at the beginning of the communication but the adversary could in theory produce additional packets in the middle of the communication to create more work on the server.


Assuming the adversary can tell that a particular source is currently active, but cannot intercept the packet itself, would it be harder for an adversary to pose as this legitimate active user if the server uses UDP or if it uses TCP? Why?

Explanation: _____

11. In TCP, if we assume $\text{cwnd} = W$ packets, what is the value of cwnd after the next ACK packet arrives?
- (a) W
 - (b) $W + \text{mss}$
 - (c) $W + 1/\text{mss}$
 - (d) I need more information

Explanation: _____

12. Consider the network above containing 6 flows (f_1 through f_6) passing through links with bandwidth capacities 10, 6, 6, and 2.



Which of the following allocation is max-min fair? If an allocation is not max-min fair, explain which flow can be increased by a small amount without reducing a flow that receives a rate less or equal.

NB: Each allocation is written $(a_1, a_2, a_3, a_4, a_5, a_6)$, where a_i is the bandwidth allocated to flow f_i .

- (a) $(1, 1, 1, 1, 1, 1)$
- (b) $(4, 3, 3, 0, 1, 1)$
- (c) $(6, 2, 2, 2, 2, 0)$
- (d) $(6, 2, 2, 2, 1, 1)$
- (e) $(8, 1, 1, 4, 1, 1)$

Explanation: _____

Are sequence numbers needed? (30 pt, including 1pt for question 4)

Motivation: Imagine you give your credit card number on the phone but your cell phone is noisy. It's annoying to give a sequence number to every digit (something like "the first digit is 0, the second digit is 7"). That's why many of us like to give the information as is and say "I'll repeat" in case something seemed wrong and we repeated a number (to avoid it is written twice). In this exercise, you will see if the same idea works between computers or whether it is in general confusing.

In order to detect duplicates while avoiding using sequence numbers, one may think of adding a different field in the header of any packet. Two alternatives are:

retransmission flagging: A unique bit is added to the header of the packet. This bit is a "0" if this is the first time this particular data is sent, or a "1" if this is not the first time it is sent.

retransmission counting: An integer value is added to the header of the packet, that indicates the number of previous transmissions of this data. This number will be "0" when this is the first time. Later, in future retransmissions it will be changed to "1", "2" etc.

You may conceptually think of it as a "sequence number" associated with all packets transmitting the same data.

Note that we assume that the receiver, having no sequence number on the packet it acknowledges, simply answers either with a "ACK" or a "NACK" at the reception of a packet.

1. (\rightarrow) Prove using a simple counter example that none of these schemes can correctly operate if packets can get reordered.

For the following three questions, we assume that packets are not reordered. However, they may be delayed or loss. The sender implements a time-out to avoid deadlock.

2. (\curvearrowright) Prove that **retransmission counting** does not correctly operate with packet losses. To do that, you will provide a counter-example that uses the minimum number of packets and packet loss.
