

Final

P. Prakash, W. An, P. Nirantar, R. Yu, X. Zhu (TAs) A. Chaintreau (instructor)

Please indicate your uni: _____

- This exam should be answered in 170mn. It is graded on a total of 350pt.

General Review (16pt × 9 questions = 144 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. It takes a single bit **no** longer to propagate over a 10Mb/s link than over a 100Mb/s link.

- (a) True
- (b) False
- (c) I need more information: is the link connected to a router or a switch?
- (d) May I still ask that on piazza? Or is 5mn after final begins too late? I still have left days!

Explanation: _____

2. The TCP protocol uses a sliding window protocol. The window size varies because:

- (a) Routers along the route advertise a varying window size to prevent congestion.
- (b) The destination advertises a reduced window size when its buffers are congested.
- (c) The destination advertises a reduced window size when packets take a long time to reach it.
- (d) The source reduces its window size when it detects that congestion is occurring.

Explanation: _____

3. Link layer. Which of the following are true?

- (a) An Ethernet switch can correctly operate if it has a 10Mb/s Ethernet network card and another 1Gb/s Ethernet card.
- (b) An Ethernet hub can correctly operate if it has 3 different 10Mb/s Ethernet network card.
- (c) An Ethernet switch cannot detect collisions until it has computed a checksum over the frame.
- (d) Ethernet Switches keep a forwarding table that they also use to answer future ARP requests.

Explanation: _____

4. An IP packet destined at a server in google.com domain is received on a network interface card in a switch, which of the following packet fields are modified by the switch?

- (a) IP header Source address
- (b) IP header Destination address
- (c) IP header TTL
- (d) IP header checksum

Explanation: _____

5. Consider a router inside one AS that receives a packet destined for a network inside another AS. To choose the next-hop router:

- (a) An intra-AS routing protocol is used
- (b) An inter-AS routing protocol (BGP) is used
- (c) Both of the above
- (d) None of the above

Explanation: _____

6. A datagram containing a UDP segment is forwarded through a NAT box going out from a private network, which of the following packet fields are updated?

- (a) IP header Source address
- (b) IP header Destination address
- (c) IP header checksum
- (d) Source UDP port number

Explanation: _____

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

- (a) Ethernet switches learn addresses from the source address of packets they forward.
- (b) Collisions occur less on a switched Ethernet network because links run faster.
- (c) In a network containing at least one switch, and after initialization, a packet is never received by another destination than the correct one.
- (d) Ethernet hubs and repeaters learn addresses of hosts from the origin address of packets they receive.

Explanation: _____

8. BGP:

- (a) is a Path Vector protocol
- (b) always eventually converges to a path that traverses a minimum number of ASes.
- (c) can handle some amount of security using the import policy.
- (d) none of the above

Explanation: _____

9. To improve the data-rate to your home, you connect your computer using separate links to ten different service providers (multi-homing)! To balance the load of traffic you have a device that assigns each packet transmitted to a provider, where the choice is determined using the source port number.
- (a) It is impossible to use traceroute from your computer to learn the different path from different ISPs.
 - (b) Assuming no NAT, and a single TCP connection, the download rate from a particular server to your computer may increase thanks to multi-homing.
 - (c) Assuming no NAT, and using multiple TCP connections, the download rate from a particular server to your computer may increase thanks to multi-homing.
 - (d) This increases the chance of packet fragmentation, even if all the links have the same MTU.

Explanation: _____

Exercise 1: Collision Detection (16pt) In CSMA with collision detection, whenever a collision is detected at the source, it always indicate that the packet was corrupted at the destination. Provide a proof or a counter example.

Exercise 2: Fairness (12pt) Which one(s) of the following are true? If a statement is not correct, provide a counterexample.

1. When all routers are working 100% of the time, hence at full capacity, the network overall is transporting the maximum total amount of information per unit of time.
2. When the network overall is transporting the maximum amount of information per unit of time, all routers are working at full capacity.
3. When all routers have the same capacity, max-min fairness always assigns the same throughput to all flows.
4. Max-min fairness assigns the same throughput to all flows in a path only if every router has the same capacity along that path.

Exercise 3: Inter-domain routing (20pt) An AS sometimes receives a path to a IP-prefix from a neighboring AS but this does not necessarily lead to this path being advertised to all other neighboring ASes. In fact, it may sometimes lead to no path being advertised at all.

1. Provide two situations in which an AS that is a provider for some customer ASes may not advertise any new path as a result of receiving a new IP path to a prefix.
2. Provide two other situations that can occur if you assume in addition that the AS is connected to two or more other ASs but none of these is a customer.

Exercise 4: Can you make high frequency trading with cheap equipment? (30pt)

Imagine you operate a network for a corporate customer, and you are able to design your own link layer protocol in terms of packet size, protocol etc. The customer is going to use this network for fast response: a job is to be transmitted, it is critical that the network communicates it to the destination with as short a delay as possible (think for instance of high frequency trading, but this could be also being first to response to a large data job). On the other hand, such job occurs randomly and may be frequent or relatively rare.


You have the choice between connecting destinations (think of offices submitting jobs, and a data-center computing the model) using either a series of 10 expensive network switches operating at 1Gb/s or 10 cheaper hubs operating at 200Mb/s. The propagation speed between two destinations on that network on any of these technology is 10ms.

1. Assume first that the customer requires that each job is sent in a single frame. Are there particular conditions (large or small job size? high or low loads?) under which the network of hubs make a better job?
2. In the context of the last question, are there other risks that should be considered and could justify one or the other technology?
3. Assume now you can set your own standard for the datagram and frames exchanged on this network, describe the conditions (on job size and load) under which using switch can be more advantageous?

NB: To answer these questions, it's important to give as much detail as possible on how job sizes and different network relate to the most important metric (delay), but your answer do not have to be quantitative when it comes to the effect of the load (as this is a more complex phenomenon).

Exercise 5: Dijkstra’s algorithm (25 pt)

We consider the following network, and, as a reminder, a table showing how to present the steps of Dijkstra’s algorithm (on a complete different network):



Step	N'	D(v) p(v)	D(w) p(w)	D(x) p(x)	D(y) p(y)	D(z) p(z)
0	u	7,u	3,u	5,u	∞	∞
1	uw	6,w	5,u	11,w	∞	∞
2	uwvx	6,w		11,w	14,x	
3	uwxv			10,v	14,x	
4	uwxyv				12,y	
5	uwxyz					

- Using Dijkstra’s algorithm, compute the shortest-path tree from router F to every other router in the network. Present the result in a table where each line denotes a step of the algorithm, and the current estimation for all nodes that remain to be treated by the algorithm, as shown in the example.

Exercise 6: Fragmentation (20pt) At the transport layer, because transport protocol may have specified packet size, application message are divided into segments that may have to be reordered or duplicated, this is why sequence number are used to number them. At the network layer, because links have specified maximum packet sizes, datagram can also be fragmented, but instead of using a sequence number they assign initially a packet ID, then include the number of bytes that were preceding this one for that particular ID, and a 0/1 flag that indicates that the packet was “fragmented”.

- Provide a counter example why the use of transport level sequence number would not work for network level packet fragmentation.
- Inversely, provide an example of what can go wrong if the format used for network level fragmentation is applied to a reliable data protocol.

Exercise 7: Managing a hotline (38 pt)

Lucky you! You have been chosen to manage the hotline of your newly funded start-up. Of course, you are going to outsource it. After a few days you made the following observations:

- The hotline providers is paid per number of request served. But they monitor how many customers they answer per minute, and they charge a significant premium free for “peaks” that is a fixed amount you pay for every minute in which they have more than 4 requests that arrived.
- You expect to receive on average 150 requests during an hour.
- Customers typically ask questions at time independently of each other. Your product is now global and the chance to receive an email at anytime during the day is exactly the same.

You would like to know how much they are going to charge.


1. What is the probability that a minute incurs a “peak” premium charge? (Hint: It is recommended that you first derive the probability that exactly k requests arrive during this minute.)
2. On your website you introduce an automatic response that is received by each request before they actually contact the hotline. This automatic response answers a few common questions, and asks for confirmation to request additional support. You notice that for every request, there is a $1/3$ probability that it is handled through this automatic response and hence never actually contact the hotline. This event, that occurs with probability $1/3$, is independent from everything else. If the request is not handled this way, the customers immediately click on the confirmation and the request arrives in the queue as usual.

Does that significantly reduce your premium “peak” fee? By how much.

(Hint: There are at least two ways to answer this question; one of them is very short and the other is very long. It means that if you end up with very complicated formulae you should probably come back and consider a simpler way to derive the result).

Exercise 8: Collision with partial deployment of carrier sensing (45 pt)

Consider 4 nodes pictured above, where A , B and C each wish to transmit to D . All solid edges are the same length, hence nodes separated by a solid edge are the same distance from one another, and the propagation delay between them is τ . The only exception is the pair of nodes A and C , who are separated by distance $c\tau$ where $c = \sqrt{3}$.



All nodes transmit frames whose transmission time L on the medium satisfies $L > \sqrt{3}\tau$. Prior to transmitting, each device runs a backoff timer that is exponentially distributed, with devices A and C using rate λ and device B using rate μ .

1. Suppose device B uses carrier sensing whereas devices A and C do not. When device B 's timer expires, if it senses another device transmitting, it resets its timer and backs off again. A and C always transmit when their timers expire. When device B transmits a frame, what is the probability that its transmission is successful (*i.e.*, the probability that A and C 's transmissions do not collide with B 's transmission attempt)?
2. In the above setting where B is still the only device utilizing carrier sensing, what is the probability that A 's transmission is successful (*i.e.*, the probability that B and C 's transmissions do not collide with A 's transmission attempt)?
3. Suppose in this part that all three devices use carrier sensing, $\lambda = \mu$, and C is transmitting a frame. Because of the different distance B will sense C terminating its transmission $(\sqrt{3} - 1)\tau$ time units before A . Does this give B an advantage (*i.e.*, does it mean that B has a larger probability to use the medium than A)? Explain your answer.