

#5 – Link Layer

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

How to read this assignment : Exercise levels are indicated as follows

(→) “elementary”: the answer is not strictly speaking obvious, but it fits in a single sentence, and it is an immediate application of results covered in the lectures.

Use them as a checkpoint: it is strongly advised to go back to your notes if the answer to one of these questions does not come to you in a few minutes.

(↪) “intermediary”: The answer to this question is not an immediate translation of results covered in class, it can be deduced from them with a reasonable effort.

Use them as a practice: how far are you from the answer? Do you still feel uncomfortable with some of the notions? which part could you complete quickly?

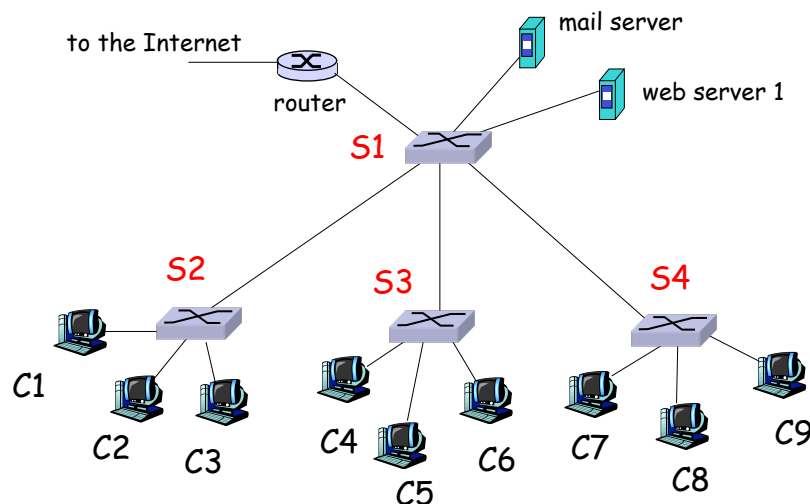
(↻) “tortuous”: this question either requires an advanced notion, a proof that is long or inventive, or it is still open.

Use them as an inspiration: can you answer any of them? does it bring you to another problem that you can answer or study further? It is recommended to work on this question only AFTER you are done with the rest!

Exercise 1: Getting to know 802.11 MAC and association protocol (15 pt) Complete the Wireshark lab for 802.11 that is provided on the website.

Exercise 2: Hubs and Switches (5 pt)

We consider the following Local Area Network containing 3 local Ethernet Switches (S2,S3,S4) and 1 Central Ethernet Switch (S1), as well as various end-hosts (9 clients and 2 servers). We assume that all links that appear on the Figure are 50Mbps full duplex (they can be used at this rate in both directions).



1. (\curvearrowright) Assuming that you can choose arbitrarily to whom each host is sending (either to a server or a client on this LAN or even to another machine on the Internet), what is the maximum amount of data that can simultaneously be sent from all the hosts (clients and servers)?
2. (\curvearrowright) Same question if the local switch S2, S3 and S4 are replaced by a hub.

Exercise 3: Answering Email request to SSOL (10 pt)

You are operating a support service for SSOL through email. Through experiences you made the following observations:

- Students typically ask questions totally independently of each other. In particular the chance to receive an email at anytime during the day is the same.
- You receive on average 120 request during an hour.

You would like to know what is the chance that you get a burst of request arriving in the queue, which would create delay and stress in your staff.

1. (\curvearrowright) Through some observations, you notice that receiving more than 3 requests in a minute is a psychological threshold that essentially makes things worse.


What is the probability that, in a given minute, there is at least 3 request arriving in the queue to request support? (Hint: It is recommended that you first derive the probability that exactly k requests arrive during this minute.)

2. (\curvearrowright) You introduce an automatic response that is received by all request before they arrive to the queue. 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 $2/3$ probability that it is handled through this automatic response and hence never actually arrives in the queue. This event, that occurs with probability $2/3$, is independent from everything else. If the request is not handled this way, the student immediately click on the confirmation and the request arrives in the queue.

What is the probability that, in a given minute, there is at least 3 requests that are confirmed and arrive in the queue? (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 4: Collision with partial deployment of carrier sensing (10 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. (\curvearrowright) 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. (\curvearrowright) 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. (\curvearrowright) 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.