Computer Networks - Final Exam

January 21, 2020

Duration: 2:15 hours, closed book.

- This is a closed-book exam.
- Please write your answers on these sheets in a readable way, in English or in French.
- Please do **not** use a red pen.
- You can use extra sheets if necessary (don't forget to put your name on them).
- The total number of points is 60.
- This document contains 24 pages.
- Good luck!

Last Name (Nom):

First Name (Prénom):

SCIPER No:

Division: ☐ Communication Systems ☐ Computer Science
☐ Other (mention it): . . . . . .

Year: ☐ Bachelor Year 2 ☐ Bachelor Year 3
☐ Other (mention it): . . . . . .

Problem	Total Points	Points Achieved
1	10	
2	25	
3	10	
4	15	

(answers to the questions are shown in italic and blue)

Problem 1 (10 points)

For each question, please circle a single best answer.

1. As long as all packets from end-system A to end-system B follow the same sequence of links, they experience the same:

   (a) Propagation delay. *(Correct)*

   (b) Queuing delay.

   (c) Loss rate.

   (d) All of the above.

2. You and your friend visit the same URL but get different webpages in response. Which of the following are plausible reasons?

   (a) A DNS impersonation attack.

   (b) A malfunctioning proxy web server.

   (c) Cookies.

   (d) All of the above. *(Correct)*

3. A DNS client sends a DNS request to its local DNS server D, asking for the IP address of `www.epfl.ch`.
   D responds. In which scenario is D's response guaranteed to be correct? Assume that all DNS servers operate correctly, and none are under any type of attack.

   (a) Never, because D may have cached a stale mapping for `www.epfl.ch`.

   (b) D is a root DNS server.

   (c) D is an authoritative server for domain `epfl.ch`. *(Correct)*

   (d) Always.

4. End-system A wants to distribute a file to a large number of other end-systems. Suppose that A has infinite upload capacity. Based on the bounds we computed in class, would client-server or peer-to-peer file distribution be faster?

   (a) Client-server.

   (b) Peer-to-peer.

   (c) About the same. *(Correct)*

   (d) It depends on the size of the file and the exact network topology.

5. The only active socket on your computer is a TCP listening socket, bound to port number 80.
   Which of the following packets will your computer receive successfully?

   (a) Any packet with destination port number 80.

   (b) Any TCP segment with destination port number 80.

   (c) Any TCP SYN segment with source port number 80.

   (d) Any TCP SYN segment with destination port number 80. *(Correct)*

6. End-system A is sending data non-stop to end-system B using TCP at the transport layer. There is 0 packet loss between A and B. Which of the following factors will determine A's throughput?

   (a) TCP flow control. *(Correct)*

   (b) TCP congestion control.

   (c) The round-trip-time (RTT) from A to B.

(d) All of the above.

7. A distributed application running on end-systems A and B sends data from A to B using UDP at the transport layer. Some packets from A to B get lost due to network congestion. How does A react?

   (a) It reduces its transmission rate to avoid creating further congestion.

   (b) It increases its transmission rate to compensate for the data loss.

   (c) It does not react at all because it has no way of knowing that packets were lost.

   (d) It depends on the application. *(Correct)*

8. Packet P_1 has destination IP address $5.0.0.1$. Packet P_2 has destination IP address $5.0.0.101$. IP router R receives P_1 and P_2. R's forwarding table does not change. Will P_1 and P_2 match the same or different entries in R's forwarding table?

   (a) Different, because they have different IP addresses.

   (b) The same, because their destination IP addresses belong to the same IP prefix.

   (c) It depends on the contents of R's forwarding table. *(Correct)*

   (d) It depends on the contents of the packets.

9. Alice and Bob exchange messages over TCP without any extra security mechanism. Persa is an on-path adversary who can intercept/copy/manipulate their packets. Can Persa launch a successful message-reordering attack against Alice and Bob?

   (a) Yes. *(Correct)*

   (b) No, because TCP will detect that the messages were reordered.

   (c) No, because TCP will put the messages back in order.

   (d) I don't have enough information to answer this question.

10. All IP routers in the same Autonomous System (AS) participate in the same:

    (a) MAC learning protocol.

    (b) Intra-domain routing protocol. *(Correct)*

    (c) Inter-domain routing protocol.

    (d) All of the above.

Problem 2 (25 points)

Consider the network in Figure 1, which includes:

- Web server `www.as1.ch` and DNS server `dns.as1.ch`.
- Personal computers C_0, C_1 ... C_{499} (there are 500 of them).
- IP routers R_1, R_2, and R_3.
- Link-layer switches S_1 and S_2.

You can find a copy of this network topology at the end of the exam. You can detach it so that you can look at the topology while solving the problem, without having to turn the pages back and forth.

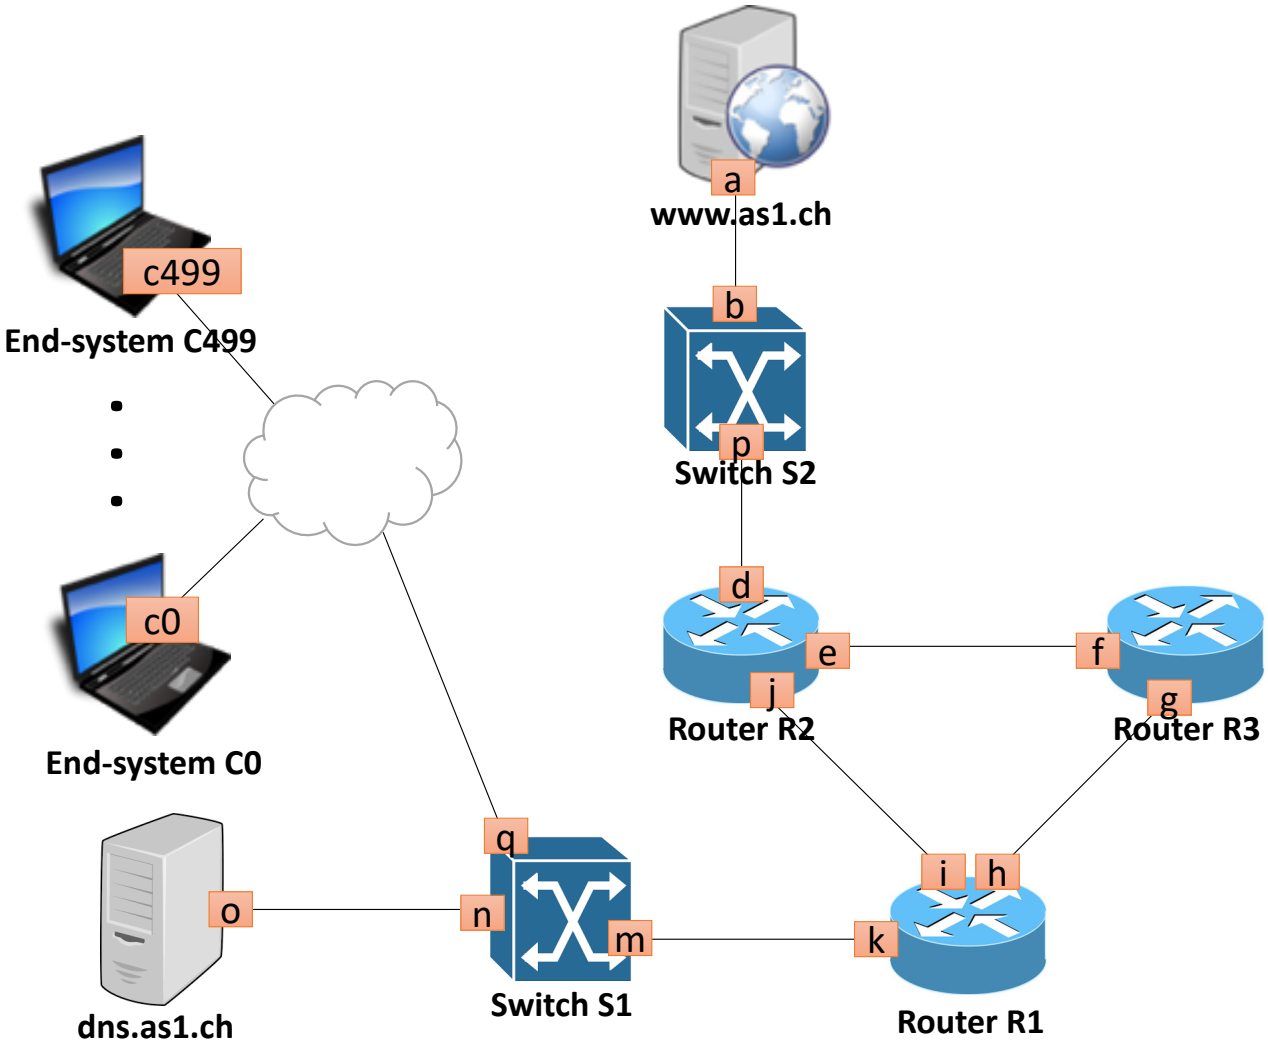

Figure 1: The Network Topology used in Problem 2

Question 1 (8 points):

Allocate an IP prefix to each IP subnet. Then allocate an IP address to each end-system network interface and to each IP-router (but not link-layer switch) network interface. Follow these rules:

- All IP prefixes and IP addresses must be allocated from $1.0.0.0/16$.
- Each IP subnet must be allocated the smallest possible IP prefix and must have one broadcast IP address.

Explain in one or two sentences how you compute each IP prefix and fill in Table 1 on the next page.

First, let's write $1.0.0.0/16$ in its binary form:

$$00000001.00000000.xxxxxxxx.xxxxxxxx.$$

Subnet behind R_1, interface k:

- 503 IP addresses (501 for the end-systems, 1 for the router interface, and 1 broadcast IP address). Hence, 9 bits.
- IP prefix $00000001.00000000.0000000x.xxxxxxxx = 1.0.0.0/23$.

Subnet behind R_2, interface d:

- 3 IP addresses, hence, 2 bits.
- IP prefix $00000001.00000000.00000010.000000xx = 1.0.2.0/30$.

Subnet behind R_1, interface i:

- 3 IP addresses, hence, 2 bits.
- IP prefix $00000001.00000000.00000010.000001xx = 1.0.2.4/30$.

Subnet behind R_1, interface h:

- 3 IP addresses, hence, 2 bits.
- IP prefix $00000001.00000000.00000010.000010xx = 1.0.2.8/30$.

Subnet behind R_2, interface e:

- 3 IP addresses, hence, 2 bits.
- IP prefix $00000001.00000000.00000010.000011xx = 1.0.2.12/30$.

Subnet	IP prefix	Interfaces and IP addresses	Broadcast IP address
Behind R_1, interface k	$1.0.0.0/23$	k: $1.0.0.1$ o: $1.0.0.2$ $c0$: $1.0.0.3$ $c499$: $1.0.1.246$	$1.0.1.255$
Behind R_2, interface d	$1.0.2.0/30$	d: $1.0.2.1$ a: $1.0.2.2$	$1.0.2.3$
Behind R_1, interface i	$1.0.2.4/30$	i: $1.0.2.5$ j: $1.0.2.6$	$1.0.2.7$
Behind R_1, interface h,	$1.0.2.8/30$	h: $1.0.2.9$ g: $1.0.2.10$	$1.0.2.11$
Behind R_2, interface e	$1.0.2.12/30$	e: $1.0.2.13$ f: $1.0.2.14$	$1.0.2.15$

Table 1: Allocation of IP prefixes and IP addresses for the network in Figure 1

Question 2 (2 points):

Routers R_1, R_2, and R_3 participate in a least-cost path routing protocol that has converged. All links have the same cost in both directions.

Fill in R_1's forwarding table right below:

Destination IP prefix	Output link
$1.0.2.0/30$	i
$1.0.2.4/30$	i
$1.0.2.8/30$	h
$1.0.2.12/30$	i (or h)
$1.0.0.0/23$	k

Question 3 (1 point):

All the computers $C_i, i = 0..499$, are configured with a default gateway.

Which is the IP address of their default gateway?

The IP address of interface k.

Question 4 (9 points):

All link-layer switches have just been rebooted, and all end-system caches are initially empty. Then, a user of computer C_0 visits web page `www.as1.ch/index.html`, which contains no images.

State all the packets that are **received, forwarded, or transmitted by router R_2 until C_0's user can view the web page**. For example, if a packet follows the path $C_0 \rightarrow R_1 \rightarrow R_2 \rightarrow$ `www.as1.ch`, then you should state it 2 times: when it is received by R_2, and when it is forwarded by R_2.

Answer by filling in Table 2. When you want to refer to the IP address of interface x, write "x". When you want to refer to the MAC address of interface x, write "x". If a field is not applicable, indicate that with a "–". To repeat a field from the above cell, write "."

#	Source MAC	Dest MAC	Source IP	Dst IP	Transp. prot.	Src Port	Dst Port	Application & Purpose
1	i	broadcast	-	-	-	-	-	ARP request for j's IP
2	j	i	-	-	-	-	-	ARP reply
3	i	j	c0	a	TCP	3000	80	TCP SYN
4	d	broadcast	-	-	-	-	-	ARP request for a's IP
5	a	d	-	-	-	-	-	ARP reply
6	d	a	c0	a	TCP	3000	80	TCP SYN
7	a	d	a	c0	TCP	80	3000	TCP SYN ACK
8	j	i	a	c0	TCP	80	3000	TCP SYN ACK
9	i	j	c0	a	TCP	3000	80	HTTP GET index
10	d	a	c0	a	TCP	3000	80	HTTP GET index
11	a	d	a	c0	TCP	80	3000	HTTP OK
12	j	i	a	c0	TCP	80	3000	HTTP OK

Table 2: Packets received, forwarded, or transmitted by router R_2 in Question 4

Question 5 (2 points):

Show the forwarding table of link-layer switch S_1 right after the last packet you stated above has arrived at its destination. Assume that no other traffic was exchanged.

Destination MAC address	Output link
*c*0	q
o	n
k	m

Question 6 (3 points):

Suppose the topology in Figure 1 is part of Autonomous System (AS) AS1, which is connected to another AS, AS2, as shown in Figure 2.

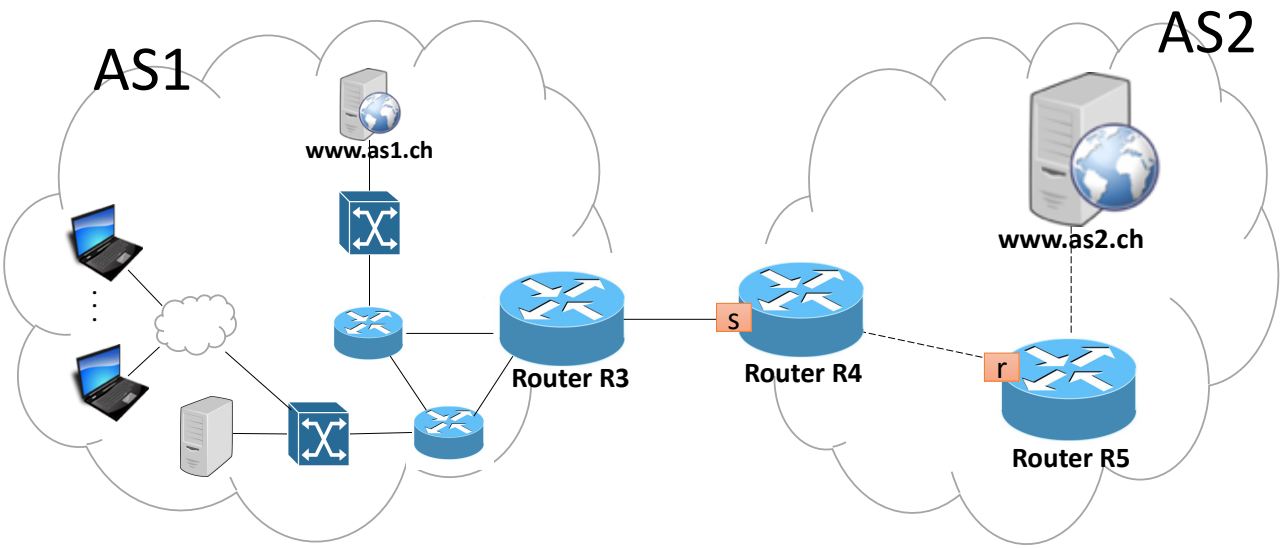

Figure 2: The Network Topology used in Problem 2, Question 6

Web server `www.as2.ch`, which is located in AS2, sends a packet to `www.as1.ch`. The packet traverses router R_5, then router R_4, before exiting AS2.

(a) Show the entry of R_5's forwarding table that matches the packet by filling in the one-row table right below. Which routing protocol created this entry? Justify your answer in one sentence.

Destination IP prefix	Output link
1.0.0.0/22	r

The intra-domain routing protocol run inside AS2. Because R_5 is not an AS border router, so it learns its routes through the local intra-domain routing protocol.

(b) Show the entry of R_4's forwarding table that matches the packet by filling in the one-row table right below. Which routing protocol created this entry? Justify your answer in one sentence.

Destination IP prefix	Output link
$1.0.0.0/22$	s

BGP. Because R_4 is an AS border router, so it learns its routes through BGP.

Problem 3 (10 points)

Every time Alice has messages to send to Bob, she uses the following protocol to send her messages to Bob:

- Alice $\rightarrow$ Bob: [Hello, K_A^+]
- Bob $\rightarrow$ Alice: [Hello, $Nonce$]
- Alice $\rightarrow$ Bob: [m_1, $K_A^-\{(Nonce + 1)|m_1\}$]
- Alice $\rightarrow$ Bob: [m_2, $K_A^-\{(Nonce + 2)|m_2\}$]
- Alice $\rightarrow$ Bob: ...
- Alice $\rightarrow$ Bob: [m_N, $K_A^-\{(Nonce + N)|m_N\}$]
- Alice $\rightarrow$ Bob: [Bye]

where:

- $[xxx]$ denotes a packet with payload (content) xxx.
- K_A^+ is Alice's public key.
- K_A^- is Alice's private key.
- $Nonce$ is an integer.
- $m_1, m_2, ...m_N$ are messages sent by Alice to Bob.
- $|$ denotes concatenation, e.g., Alice produces "$K_A^-\{(Nonce + 1)|m_1\}$" by putting $(Nonce + 1)$ and m_1 one after the other and encrypting the outcome with her private key.

Bob does not know Alice's public key K_A^+ before their communication begins.

Manuel is a bad guy who has taken control of a network device on the path between Alice and Bob. So, Manuel can copy, intercept, and manipulate any packet sent by Alice to Bob and vice versa.

Question 1 (3 points):

Describe in two-to-four sentences how Manuel can launch a successful reordering attack against Alice and Bob, i.e., make Bob believe that Alice sent her messages in a different order than she actually did. As part of your answer, show the messages sent by Alice and Manuel.

A reordering attack is possible due to the absence of a trusted third party vouching for a particular key belonging to Alice. Specifically, Manuel drops any messages Alice sends to Bob, and initiates a conversation with Bob by sending to him Manuel's own public key. Bob is not able to distinguish between Alice's and Bob's public key, and therefore, Manuel uses his own private key to send to Bob any kind of messages in whichever order Manuel decides.

The exchange of messages that respect the protocol and result in a successful reordering attack for Manuel can be seen in Figure 3.

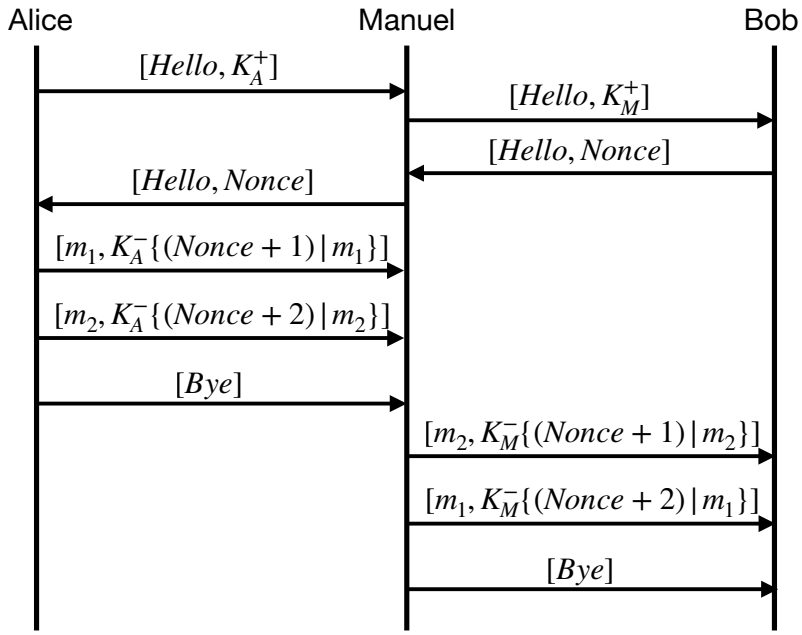

Figure 3: Message exchanges in the reordering attack.

Question 2 (2 points):

Make the minimum change(s) to Alice and Bob's protocol to prevent Manuel from launching a successful reordering attack. State the corrected protocol.

Have Alice providing Bob with a certificate, along with her public key. The certificate is vouching for Alice's key belonging to Alice and is issued by a Certificate Authority (CA) that both Alice and Bob trust.

The corrected protocol is same as the initial one except for the first message which now becomes:

- Alice $\rightarrow$ Bob: [Hello, K_A^+, $K_{CA}^-\{H(K_A^+|Alice)\}$]
- ...

where K_{CA}^- is CA's private key, and $H()$ is a globally known cryptographic hash function.

Question 3 (2 points):

Does your protocol still need the $Nonce$? Why (not)? Justify your answer in one or two sentences.

Yes, the protocol needs the Nonce to prevent both replay and reordering attacks (swapping messages between two different sessions).

Question 4 (1 point):

Does your protocol provide confidentiality? Justify your answer in one sentence.

No, as messages are sent in the clear.

Question 5 (2 points):

How can you reduce the amount of computation required by your protocol (without changing the security properties that it provides)? State the corrected protocol.

In the beginning of the protocol, use asymmetric crypto to bootstrap trust to a symmetric key K shared only between Alice and Bob. The rest of the conversation remains the same, the only difference being the replacement of KA-/KA+ by K. The aforementioned changes reduce the amount of computation required by the protocol as, for most of the part of the conversation, asymmetric crypto is replaced by the less computation-heavy symmetric crypto.

Problem 4 (15 points)

Assume the following for all the questions in this problem:

- The maximum segment size is `MSS` $= 1$ byte.
- Each TCP receiver sends an ACK every time it receives a data segment.

When you complete the diagram in Question 1, the following information should be visible:

- All the segments (including the ACKs) exchanged between the communicating end-systems.
- The sequence numbers of all data segments sent from Alice to Bob.
- The acknowledgment numbers of all ACKs sent from Bob to Alice.
- The state of Alice's congestion-control algorithm.
- The size of Alice's congestion window (`cwnd`) in bytes.
- The value of Alice's congestion threshold (`ssthresh`) in bytes.
- Any dropped segments.
- If your answer includes any timeouts, mark them clearly and indicate the sequence number of the data segment that timed out.

Question 1 (5 points):

In this question, Fast Retransmit/Fast Recovery are DISABLED.

Alice establishes a TCP connection with Bob and sends 5 bytes of data.
The 3rd segment sent by Alice (so, not the SYN, not the 1st data segment, but the 2nd data segment) is dropped. No other segment, sent by Alice or Bob, is dropped or corrupted.

Show all the segments sent by Alice and Bob, including connection setup (not connection teardown), by completing the diagram in Figure 4 on the next page.

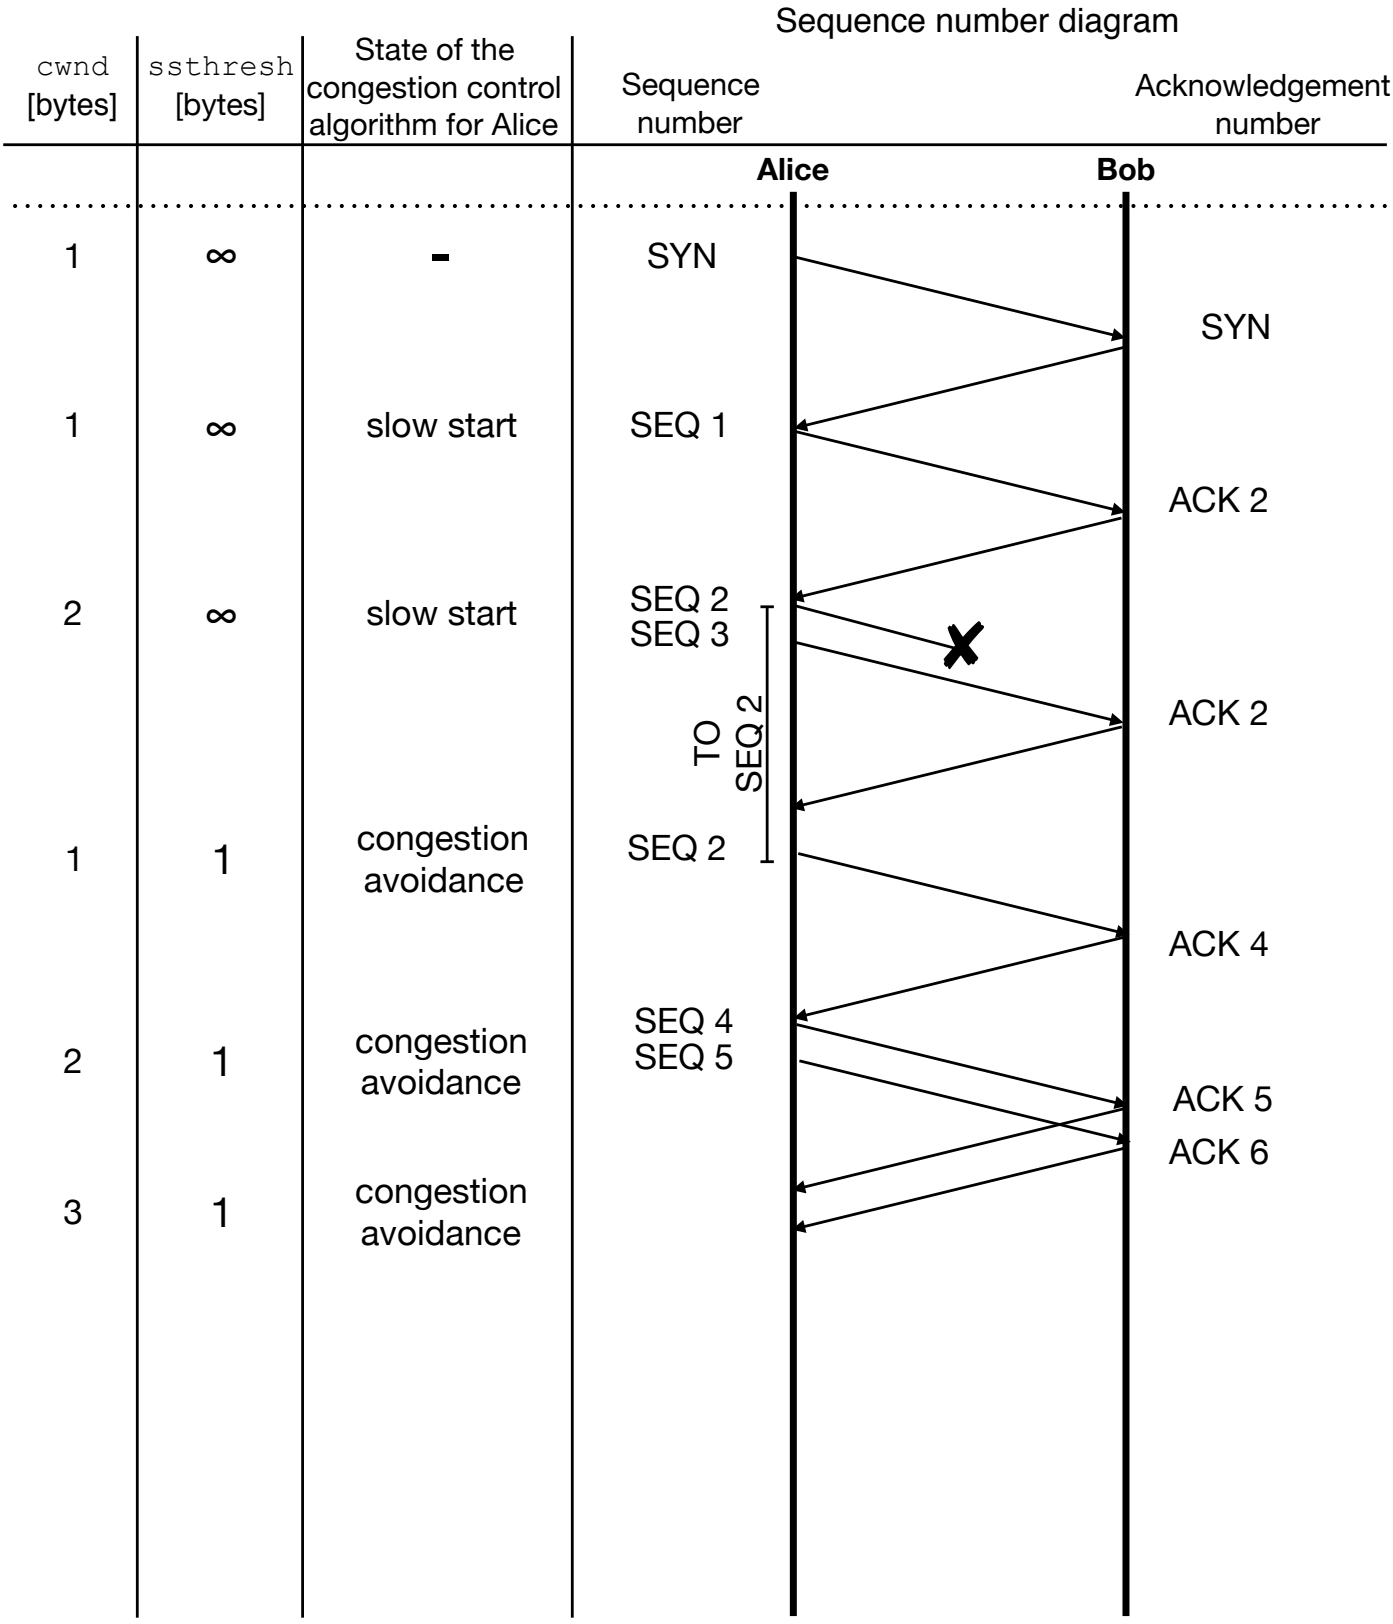

Figure 4: Sequence diagram to be completed for Question 1.

(Lab) Question 2 (5 points):

In this question, Fast Retransmit/Fast Recovery **may** be ENABLED (you will be asked to guess).

Alice has been sending data to Bob over TCP. At some point in time t, Alice's congestion-window size takes value 1 MSS and starts evolving as shown in Figure 5 (darker line).

At the same time, another sender, Céline, starts sending data to another receiver, Dabir. The Céline-Dabir traffic crosses the same bottleneck link as the Alice-Bob traffic. Figure 5 also shows the evolution of Céline's congestion-window size (lighter-colored line, with x's).

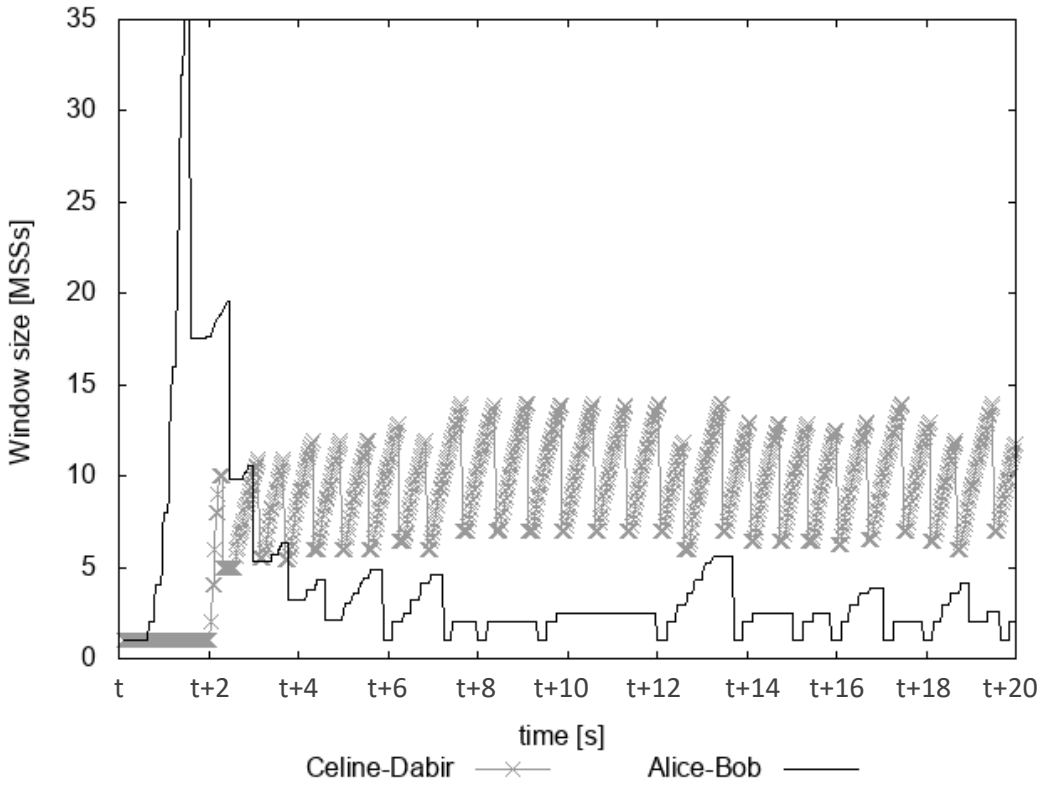

Figure 5: Congestion-window sizes for Question 3.

(a) Is Fast Retransmit/Fast Recovery enabled or disabled in Céline's congestion-control algorithm? Justify your answer in one sentence.

It is enabled. We can tell, because Céline's window size often decreases by half. If Fast Retransmit/Fast Recovery were disabled, Céline's window size could only be reset to 1 MSS after a packet loss.

(b) How do the round-trip times (RTTs) experienced by the two flows (Alice-Bob and Céline-Dabir) compare to each other? Are they approximately the same? Is one longer than the other? Justify your answer in two-to-four sentences.

We know that, in the congestion-avoidance state, the window size increases by 1 MSS (in our case 1 byte) every RTT. Hence, we can estimate the RTT by dividing the increase in window size by the time it took for this increase to occur. Alice's window increases roughly by 4 bytes in 1 sec. Céline's window increases roughly by 4 bytes in 0.5sec. Hence, the Alice-Bob RTT is rougly twice the Céline-Dabir RTT.

(c) Estimate very approximately how the average-throughput values achieved by the two flows compare to each other. Are they approximately the same? Is one approximately twice (or four times? or ten times?) as much as the other? Justify your approximation in two-to-four sentences.

The average throughput is approximately the average window size divided by the average RTT. The Alice-Bob average window size is about 2 bytes, while the Céline-Dabir average window size is about 10 bytes. The Alice-Bob RTT is about twice the Céline-Dabir RTT. So, Céline's average throughput should be about 10 times higher than Alice's average throughput.

(d) Is TCP fair to the two flows (the Alice-Bob flow, and the Céline-Dabir flow)? Depending on your answer, explain how TCP achieves fairness or why it fails to achieve fairness in this scenario. Your answer should not be longer than four sentences.

No, TCP is not fair: the Alice-Bob flow achieves significantly lower average throughput than the Céline-Dabir flow. The reason for this unfairness is that the Alice-Bob flow has a longer RTT: Given that the window increases every RTT, it increases more slowly for flows with longer RTTs. In the presence of congestion, where the window keeps being reset or halved due to packet loss, this causes flows with longer RTTs to never consume their fair share of the bottleneck link's bandwidth.

Question 3 (5 points):

Imagine an alternative Internet that uses connection switching (virtual circuits) with resource reservation, as we saw in class. So, whenever an end-system sends data to another end-system, that happens over a network-layer connection (virtual circuit). Do you think that, in this alternative Internet, end-systems should use TCP as their transport-layer protocol? Justify your answer in two-to-four sentences.

No, they should not. In a virtual-circuit network, there is no network congestion, as resources are reserved in advance for all generated traffic. Hence, there is no need for TCP's congestion control.

End-systems might still benefit from flow control and in-order delivery (which TCP also offers), but they could use a simpler/more lightweight protocol than TCP for that.

Scratch Paper

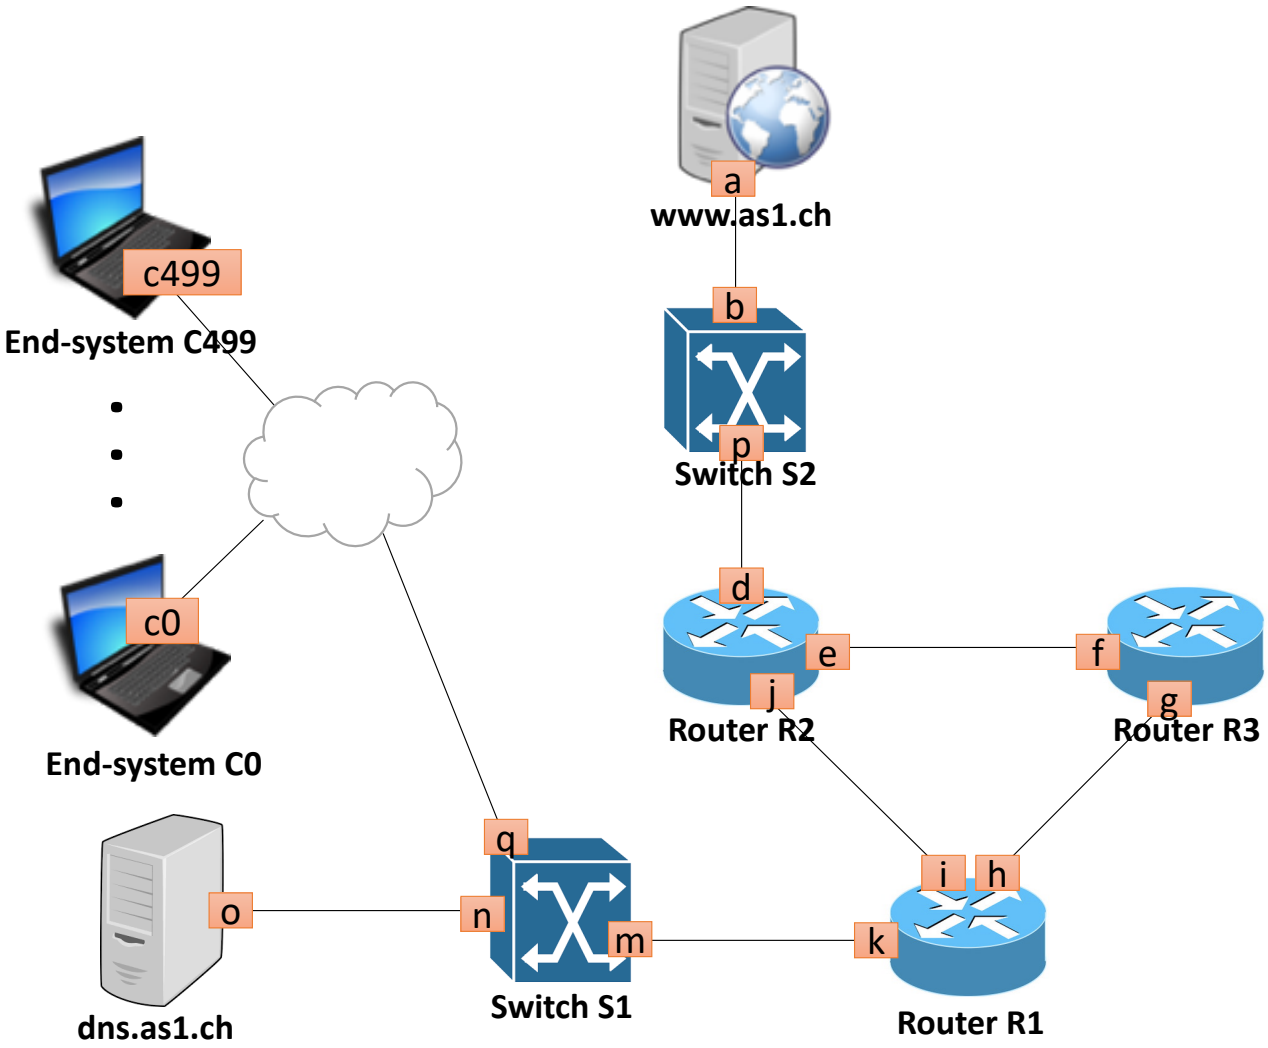

Figure 6: The Network Topology used in Problem 2