Computer Networks - Final Exam

January 19, 2021

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 50.
- This document contains 20 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	Points achieved	Out of
1		5
2		25
3		20
Total		50

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

Problem 1 (5 points)

For each question, please circle a single best answer.

1. Two Internet Service Providers (ISPs) have a "customer-provider" relationship when:
   (a) One pays the other for Internet connectivity. *(Correct)*
   (b) They exchange traffic for free.
   (c) They have a physical connection to each other.
   (d) They are connected to the same Internet eXchange Point (IXP).

2. End-system A is infinitely (forever) sending back-to-back same-size packets to end-system B over a link of transmission rate R and propagation delay d. The link is malfunctioning and droppping half the packets. The following is true:
   (a) The propagation delay from A to B is $\frac{d}{2}$.
   (b) The transmission delay from A to B is approximately $\frac{R}{2}$.
   (c) The average throughput from A to B is approximately $\frac{R}{2}$. *(Correct)*
   (d) All of the above.

3. The following is an example of a layering violation (a packet switch accesses headers that it should not, because they belong to a higher layer):
   (a) MAC learning.
   (b) IP forwarding.
   (c) Network Address Translation (NAT). *(Correct)*
   (d) Address Resolution Protocol (ARP).

4. Two users, one using end-system A, the other using end-system B, access the same URL at the same time, but get different web pages (different content) in response. What could be the explanation?
   (a) A and B have different network paths to the web server.
   (b) A uses a persistent TCP connection to communicate with the web server, whereas B does not.
   (c) The web server uses cookies. *(Correct)*
   (d) All of the above.

5. What does every end-system in the world need to know?
   (a) The DNS name of at least one DNS server.
   (b) The IP address of at least one DNS server. *(Correct)*
   (c) The IP address of at least one root DNS server.
   (d) The IP address of at least one authoritative DNS server.

6. According to what we said in class, how does a tracker compare to a distributed hash table (DHT)?

   (a) They are the same thing.

   (b) A tracker is more reliable than a DHT.

   (c) A tracker is more secure than a DHT.

   (d) They are different implementations of the same service. *(Correct)*

7. Process X running on end-system Y creates a socket and binds it to IP address `100.0.0.10` and port number 1000. The following is true:

   (a) If process X sends a packet through this socket, it has source IP address `100.0.0.10`.

   (b) If process X sends a packet through this socket, it has source port number 1000.

   (c) If a packet arrives at end-system Y with destination IP address `100.0.0.10` and destination port number 1000, it is sent to process X.

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

8. Why would an application use UDP instead of TCP as a transport-layer protocol?

   (a) To avoid the overhead of connection setup. *(Correct)*

   (b) To leverage the superior security propeties of UDP.

   (c) To leverage the superior reliability of UDP.

   (d) There is no good reason.

9. An Autonomous System (AS) changes from Dikjstra to Bellman-Ford for its intra-domain routing protocol. This may result in:

   (a) Shorter intra-domain routes (paths).

   (b) Longer intra-domain routes (paths).

   (c) Different convergence time. *(Correct)*

   (d) None of the above.

10. You successfully access a web server from your EPFL computer. You try to access the same web server from your home computer, and you cannot. What could be the reason?

    (a) There is a failure on the path from your home computer to the web server.

    (b) The web server is behind a firewall that blocks traffic from certain IP addresses/port numbers.

    (c) Any of (a) or (b). *(Correct)*

    (d) This cannot happen in today's Internet.

Problem 2 (25 points)

Consider the network in Figure 1, which includes:

- Workstations E_1, ... E_{500} (there are 500 of them).
- Workstations B_1, ... B_{1000} (there are 1000 of them).
- IP routers R_1 and R_2.
- Link-layer switch S_1.
- DNS server `dns.xxx.ch`.
- Web server `www.yyy.ch`.

Clarifications:

- The orange boxes represent network interfaces.
  For example, router R_1 has network interfaces x, y, z, and u.
- The "enterprise network" consists of everything on the left of network interface z of router R_1.
- End-systems E_2 to E_{500} and B_1 to B_{1000} also have one network interface each, we just don't show all these interfaces explicitly in the picture.

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.

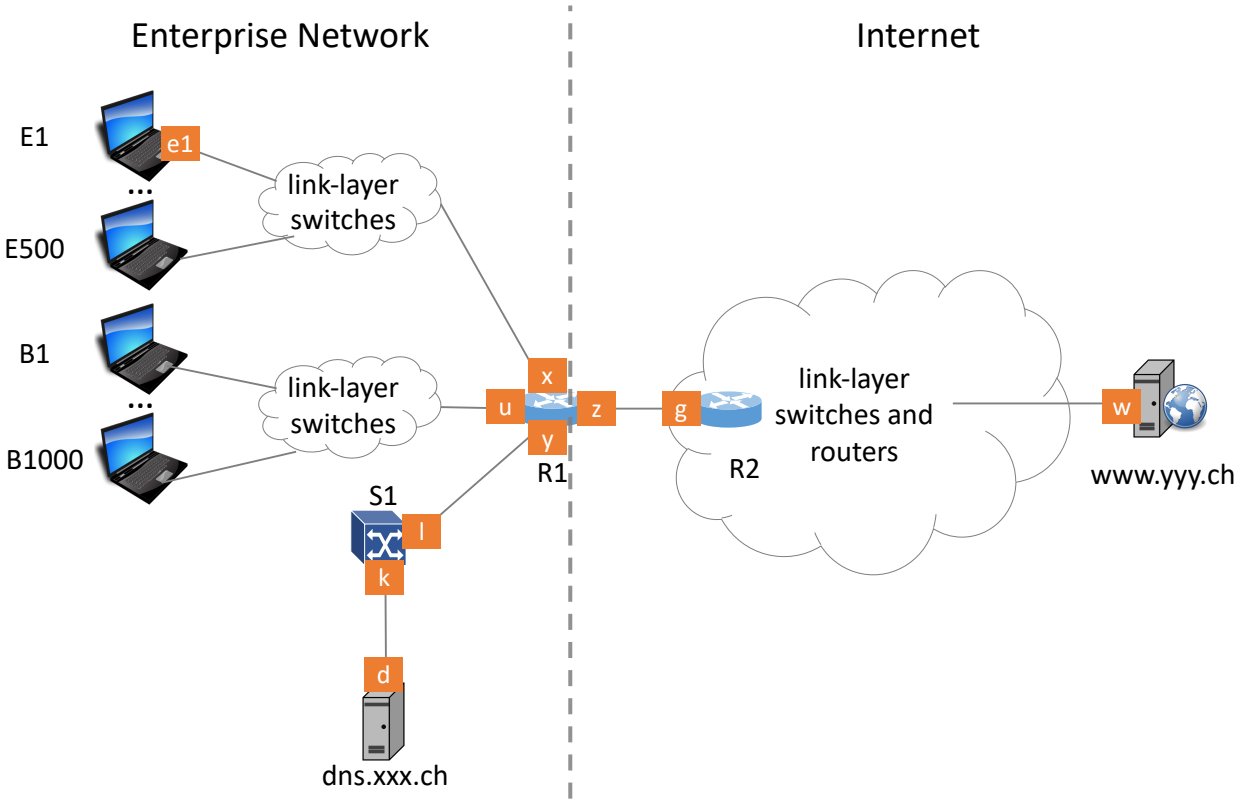

Figure 1: Network topology for Problem 2.

Question 1 (8 points):

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

- All IP prefixes and IP addresses must be allocated from $20.0.0.0/16$.
- 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 $20.0.0.0/16$ in its binary form:

$$00010100.00000000.xxxxxxxx.xxxxxxxx.$$

Subnet behind R_1, interface u:

- 1003 IP addresses (1000 for the end-systems, 1 for the router interface, 1 network IP address (optional), and 1 broadcast IP address). Hence, 10 bits.
- IP prefix $00010100.00000000.000000xx.xxxxxxxx = 20.0.0.0/22$.

Subnet behind R_1, interface x:

- 503 IP addresses, hence, 9 bits.
- IP prefix $00010100.00000000.0000010x.xxxxxxxx = 20.0.4.0/23$.

Subnet behind R_1, interface y:

- 4 IP addresses, hence, 2 bits.
- IP prefix $00010100.00000000.00000110.000000xx = 20.0.6.0/30$.

Subnet	IP prefix	Interfaces and IP addresses	Broadcast IP address
Behind R_1, interface u	$20.0.0.0/22$	u: 20.0.0.1 $B1$: 20.0.0.2 $B1000$: 20.0.3.232	1.0.3.255
Behind R_1, interface x	$20.0.4.0/23$	x: 20.0.4.1 $E1$: 20.0.4.2 $E500$: 20.0.5.246	20.0.5.255
Behind R_1, interface y	$20.0.6.0/30$	y: 20.0.6.1 d: 20.0.6.2	20.0.6.3

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

Question 2 (9 points):

All link-layer switches have just been rebooted, and all end-system caches are initially empty. Then, the user of workstation E_1 visits web page `www.yyy.ch/index.html`, which contains no embedded objects (e.g., no images).

State all the packets that are **received, forwarded, or transmitted by router R_1 until E_1's user can view the web page**. For example, if a packet follows the path $E_1 \rightarrow R_1 \rightarrow \ldots$ `www.yyy.ch`, then you should state it 2 times: when it is received by R_1, and when it is forwarded by R_1.

Answer by filling in Table 2. To denote the IP address or the MAC address of interface s, write "s". If a field is not applicable, write "–". To repeat a field from the above cell, write "." To illustrate the format, we have provided a hypothetical example entry.

#	Source MAC	Dest MAC	Source IP	Dst IP	Transp. prot.	Src Port	Dst Port	Application & Purpose
1	e_1	broadcast	-	-	-	-	-	ARP request for x's MAC
2	x	e_1	-	-	-	-	-	ARP reply
3	e_1	x	e_1	d	UDP	2000	53	DNS request for w's IP
4	y	broadcast	-	-	-	-	-	ARP request for d's MAC
5	d	y	-	-	-	-	-	ARP reply
6	y	d	e_1	d	UDP	2000	53	DNS request for w's IP
7	d	y	d	root	UDP	2500	53	DNS request for w's IP
8	z	broadcast	-	-	-	-	-	ARP request for g's MAC
9	g	z	-	-	-	-	-	ARP response
10	z	g	d	root	UDP	2500	53	DNS request for w's IP
11	g	z	root	d	UDP	53	2500	DNS response
12	y	d	root	d	UDP	53	2500	DNS response
13	d	y	d	e_1	UDP	53	2000	DNS response
14	x	e_1	d	e_1	UDP	53	2000	DNS response
15	e_1	x	e_1	w	TCP	3000	80	TCP SYN
16	z	g	e_1	w	TCP	3000	80	TCP SYN
17	g	z	w	e_1	TCP	80	3000	TCP SYN ACK
18	x	e_1	w	e_1	TCP	80	3000	TCP SYN ACK
19	e_1	x	e_1	w	TCP	3000	80	HTTP GET index
20	z	g	e_1	w	TCP	3000	80	HTTP GET index
21	g	z	w	e_1	TCP	80	3000	HTTP OK
22	x	e_1	w	e_1	TCP	80	3000	HTTP OK

Table 2: Packets received, forwarded, or transmitted by router R_1 in Question 2.

The above solution assumes that all devices use default gateway when they ARP. It also assumes that `dns.xxx.ch`'s DNS request for `www.yyy.ch`'s IP address is answered recursively by some root DNS server.

Question 3 (4 points):

Answer again Question 2, but assuming that router R_1 acts as a Network Address Translation (NAT) gateway between the enterprise network and the rest of the Internet on the right hand side.

Answer by filling in Table 3. List only the entries and cells that change relative to your answer to Question 2. For example, if the only change relative to your answer to Question 2 was that, in the 2nd entry, the transport protocol became TCP, then you would need to provide the hypothetical answer shown below.

#	Source MAC	Dest MAC	Source IP	Dst IP	Transp. prot.	Src Port	Dst Port	Application & Purpose
12	z	g	z	w	TCP	3100	80	TCP SYN
13	g	z	w	z	TCP	80	3100	TCP SYN ACK
16	z	g	z	w	TCP	3100	80	HTTP GET index
17	g	z	w	z	TCP	80	3100	HTTP OK

Table 3: Packets received, forwarded, or transmitted by router R_1 in Question 3. Only changes relative to Question 2 answer.

(Lab related) Question 4 (2 points):

The user of workstation E_1 types `ifconfig -a` into a terminal and receives the following answer:

```
e1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet x.x.x.x  netmask 255.255.255.0 broadcast 128.178.158.255
        ether a4:bb:2d:2f:f5:6b  txqueuelen 1000  (Ethernet)
        RX packets 1093229  bytes 580963781 (580.9 MB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 56084  bytes 12509105 (12.5 MB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
        device interrupt 20  memory 0xe4400000-e4420000
lo: flags=73<UP,LOOPBACK,RUNNING>  mtu 65536
        inet 127.0.0.1  netmask 255.0.0.0
        inet6 ::1  prefixlen 128  scopeid 0x10<host>
        loop  txqueuelen 1000  (Local Loopback)
        RX packets 64858  bytes 6028272 (6.0 MB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 64858  bytes 6028272 (6.0 MB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

How many network interfaces does E_1 have? What is the IP address and the MAC address of each network interface? What is `lo` and why does it not have a MAC address? What could such an interface be used for?

Two interfaces:

- e1, with IP address `x.x.x.x` and MAC address `a4:bb:2d:2f:f5:6b`.
- lo, with IP address 127.0.0.1.

lo is a loopback interface, which is not associated with an actual physical link, this is why it has no MAC address. It is typically used for communication between processes running on the local computer and for testing.

Question 5 (2 points):

The DNS name of workstation E_1 is `e1.xxx.ch`. Assume that E_1 knows its own DNS name and—of course—its own IP addresses.

The user of E_1 starts running a web-server process on E_1. Then, she visits web page `e1.xxx.ch/index.html`.

Which of E_1's network interfaces do you think will be involved in the resulting communication? State all the packets that are **transmitted or received by this network interface until the user can view the web page**. Answer by filling in Table 4.

You have not seen this exact scenario before, and it is normal not to be certain of the answer. We are asking you to make your best guess based on your understanding of network layers.

#	Source MAC	Dest MAC	Source IP	Dst IP	Transp. prot.	Src Port	Dst Port	Application & Purpose
1	–	–	e_1	e_1	TCP	3000	80	TCP SYN
2	–	–	e_1	e_1	TCP	80	3000	TCP SYN ACK
3	–	–	e_1	e_1	TCP	3000	80	HTTP GET index
4	–	–	e_1	e_1	TCP	80	3000	HTTP OK

Table 4: Packets transmitted or received by E_1's network interface in Question 5.

Problem 3 (20 points)

Assume the following for all the questions in this problem:

- Fast Retransmit/Fast Recovery are DISABLED.
- The maximum segment size is `MSS` $= 1$ byte.
- The TCP timeout is 1.5 RTT, where RTT is the round trip time from the sender to the receiver.
- A TCP receiver sends an ACK every time it receives a data segment.
- When computing delays, transmission delay is negligible.

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 Online-store.
- The acknowledgment numbers of all ACKs sent from Online-store 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):

Alice establishes a TCP connection with Online-store and (Alice) sends 12 bytes of data to Online-store. Online-store does not send any data to Alice.

The 8th segment sent by Online-store (counting the SYN-ACK as the 1st segment sent by Online-store) is dropped. No other segment, sent by Alice or Online-store, is dropped or corrupted.

Show all the segments sent by Alice and Online-store, including connection setup (not connection tear-down), by completing the diagram in Figure 2 on the next page.

cwnd [bytes]	ssthresh [bytes]	State of the congestion control algorithm for Alice
1	∞	-
1	∞	slow start
2	∞	slow start
4	∞	slow start
7	∞	slow start
1	3	slow start
4	3	congestion avoidance

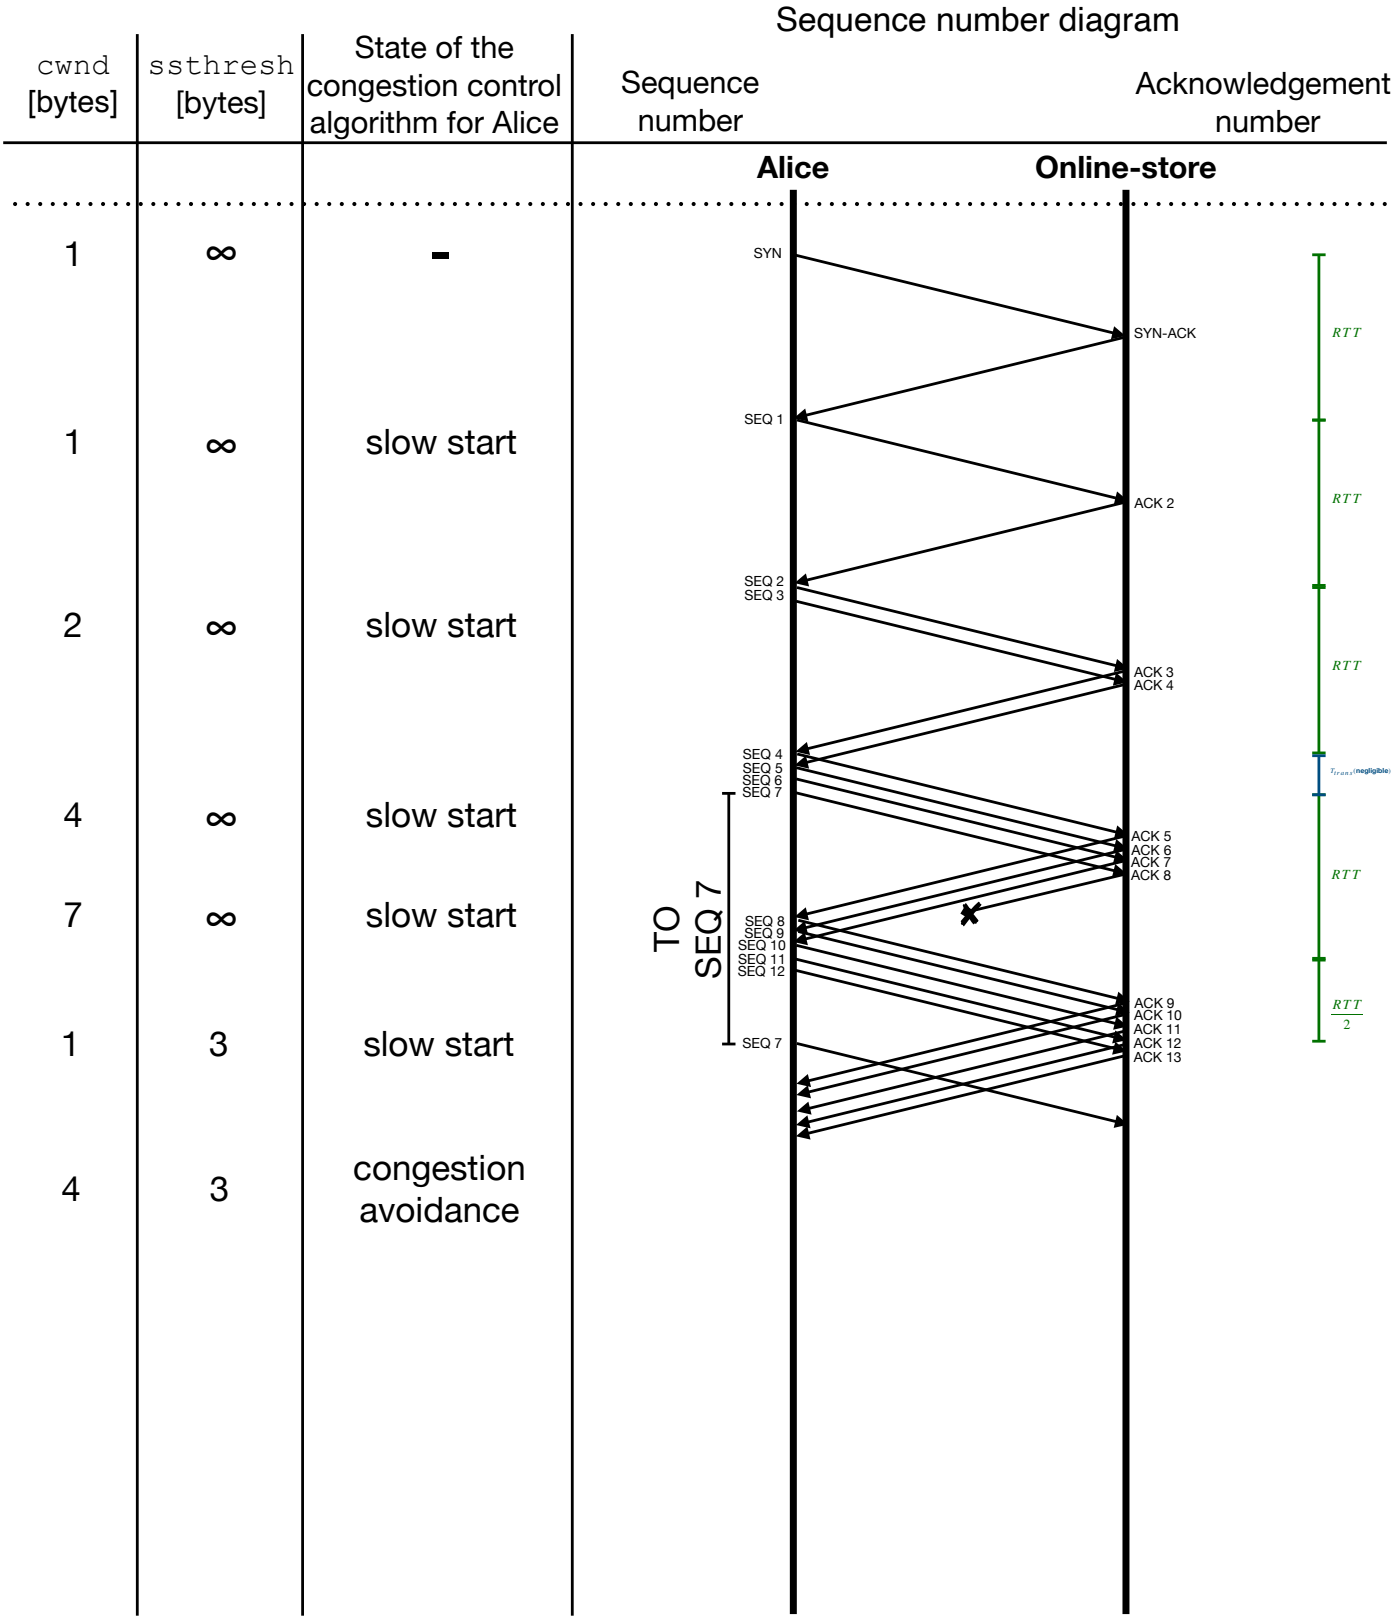

Figure 2: Sequence diagram to be completed for Question 1.

Question 2 (5 points):

Suppose Fig. 3 shows the network topology between Alice and Online-store:

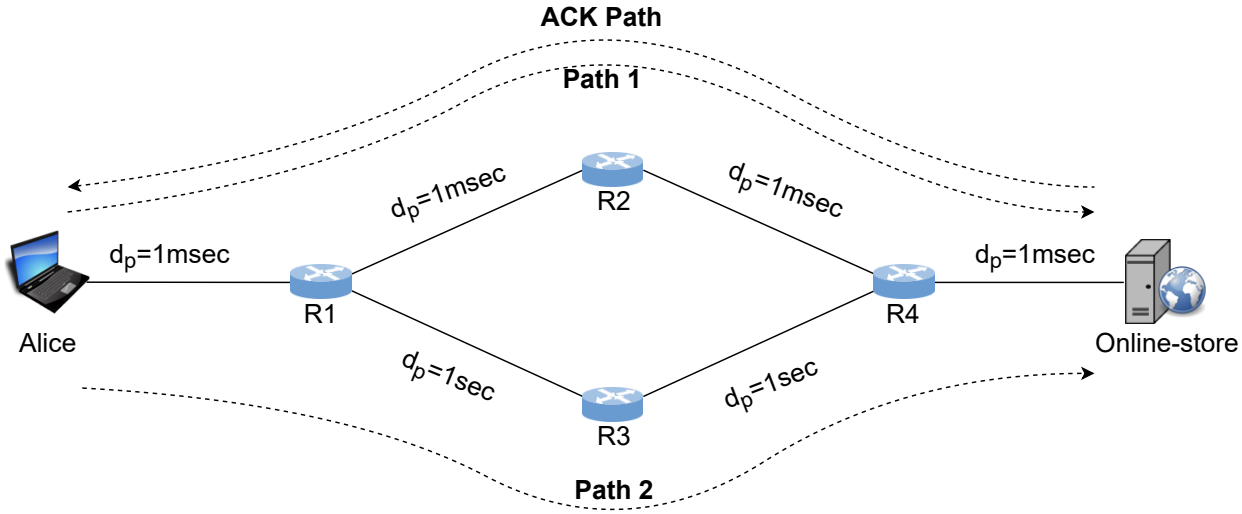

Figure 3: Network topology for Question 2.

The propagation delays of the links are shown on the figure:

- Two links in Path 2 have propagation delay 1 second.
- All the other links have propagation delay 1 millisecond.

There is no network traffic other than the one exchanged by Alice and Online-store.

(a) Suppose the packets from Alice to Online-store in Question 1 follow Path 1, while the packets from Online-store to Alice follow the ACK Path. What is the average throughput achieved in Question 1 from Alice's process to the Online-store process? In this type of question, "throughput" is the rate at which the Online-store's process receives data from Alice's process. Justify your answer.

Average throughput = $\frac{12\times 8\text{bits}}{\text{transfer time}}$.

With the connection-setup time included, transfer time is 4.5RTT (Fig. 2).

Since the packets from Alice to Online-store follow Path 1, RTT$=$ 8msec.

Hence, transfer time is 36msec or 0.036sec, and average throughput is $\frac{96}{0.036} = 2'666.66$bps.

(b) Now suppose that the packets from Alice to Online-store in Question 1 follow Path 2, while the packets from Online-store to Alice follow the ACK Path. What is the average throughput achieved in Question 1 from Alice's process to the Online-store process? Justify your answer.

Since the packets from Alice to Online-store follow Path 2, RTT$= 2.006$sec.

Hence, transfer time is = 9.027sec, and average throughput is $\frac{96}{9.027}$ = 10.63bps.

Question 3 (5 points):

At some point (after the events of Question 1), Alice continues to send data to Online-store.

At the same time, one of the following two scenarios occurs:

- Scenario A: The route from Alice to Online-store starts oscillating between Path 1 and Path 2: one packet follows Path 1, the next one Path 2, the next one Path 1, and so on.
- Scenario B: All packets from Alice to Online-store follow Path 3 (not shown in Fig. 3), which has propagation delay equal to the average between Path1's and Path2's propagation delays.

The packets from Online-store to Alice always follow the ACK Path in both scenarios. If there is packet loss, assume that the same packets are lost in both scenarios.

(a) In Scenario A, how will the route oscillation affect TCP behavior? Is there any particular aspect of TCP behavior that will be confused because of the oscillation? Answer in a few sentences.

RTT estimation will be confused.

RTT estimation is based on sample RTTs measured during the connection. The estimated RTT will be somewhere between the RTT of Path 1+ACK Path and the RTT of Path 2+ACK path.

As a result, for the data segments following Path 2, the estimated RTT may be too low, and Alice may timeout prematurely (timeout even if the segment was successfully received and ACK-ed). This will cause Alice's congestion window to be reset to 1 MSS, and will unnecessarily reduce throughput.

Moreover, for the data segments following Path 1, the estimated RTT may be too high. So, if one of these data segments is lost or corrupted, Alice will wait longer than necessary to timeout, which will also unnecessarily reduce throughput.

(b) Does the average throughput from Alice's process to the Online-store process change between the two scenarios? Which scenario do you expect to achieve higher average throughput? Justify your answer in a few sentences. (It is OK to answer that which scenario achieves higher throughput depends on some condition, but then you need to explain what that condition is.)

Scenario B should achieve higher throughput than A.

As stated above, Scenario A will most likely suffer from premature timeouts, causing the congestion window to be unnecessarily reset to 1, which will significantly reduce throughput.

Scenario B will most likely not suffer from such an effect, as there is only one path, and RTT estimation should work as expected.

Question 4 (5 points):

Now ignore the previous three questions.

Suppose Alice and Online-store secure their communication using the Secure Sockets Layer (SSL) as we saw in class. Persa is an adversary, sitting on the communication channel between Alice and Online-store.

(a) Persa tries to launch an impersonation attack, i.e., send Alice her own public key in place of Online-store's public key:

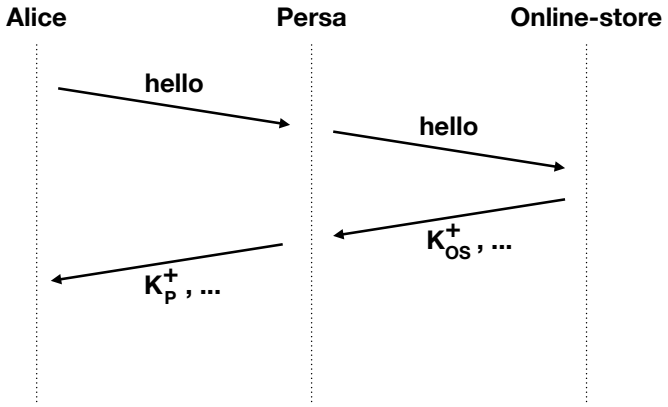

Figure 4: Impersonation attack in Question 4(a).

How does SSL protect Alice and Online-store against this attack? State in one or two sentences exactly which information will enable Alice and/or Online-store to detect the attack and how. Feel free to use a simple diagram or complete the missing information (indicated by the three dots) from Figure 4.

SSL prevents a man-in-the-middle attack using public-key certificates. The Online-store sends Alice a certificate along with its public key (K^+_{OS}). When Alice receives the message from Persa (containing $K^+_P, K^-_{CA}\{OS\ owns\ K^+_{OS}\}$), it will check the certificate (using CA's true public key) and realize that the public key sent does not belong to the Online-store.

The certificate is $K^-_{CA}\{OS\ owns\ K^+_{OS}\}$ (equivalently $K^-_{CA}\{K^+_{OS}|OS\}$).

Additional information, not required by the answer, but to address common mistakes:

- The certificate is a value initially computed by the CA during the certification of the Online-Store, the Online-Store does not and can not compute it as it does not know K^-_{CA}. The Online-Store only stores the certificate and transmits it every time it's required.
- The certificate is not a secret: the Online-Store will send the certificate to anyone who wants to initialize an SSL-based communication with the Online-Store (for example, even if Persa tries to order something from the Online-Store, the Online-Store may send Persa the same certificate).
- Persa can transmit to anyone the Online-Store's certificate, as it's a constant value that Persa can ask the Online-Store to send her, while Persa initializes communication with the Online-Store; although it would be useless for Man-in-the-middle attacks as it does not certify something relevant to Persa's public key.
- Persa may have a certificate, certifying Persa owns her public key: $K^-_{CA}\{Persa\ owns\ K^+_P\}$.
- The certificate does not certify that the Online-Store is a good actor, but only that the public key (K^+_{OS}) belongs to the physical entity "Online-Store".

(b) Alice sends two orders to Online-store, and Persa tries to reorder them:

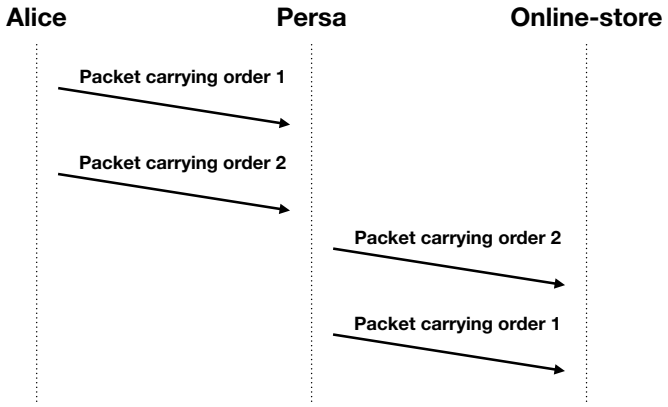

Figure 5: Reordering attack in Question 4(b).

How does SSL protect Alice and Online-store against this attack? State in one or two sentences exactly which information will enable Alice and/or Online-store to detect the attack and how. Feel free to use a simple diagram (a version of Figure 5 that provides more detail on what the packets contain).

SSL prevents re-ordering attacks using the "Message Authentication Code" (MAC). In particular, when Alice sends an order to the Online-store, Alice includes a sequence number (in addition to the rest of the message and a shared key) in the MAC computation. This way when the Online-store receives the reordered orders, the MAC validation will fail and the Online-store will detect the attack.

Additional information, not required by the answer, but to address common mistakes:

- The sequence number have a specific increment pattern; an order number (random binary string) or any random number wouldn't suffice.
- The starting point of the sequence should be known to both Alice and the Online-Store.

Scratch Paper

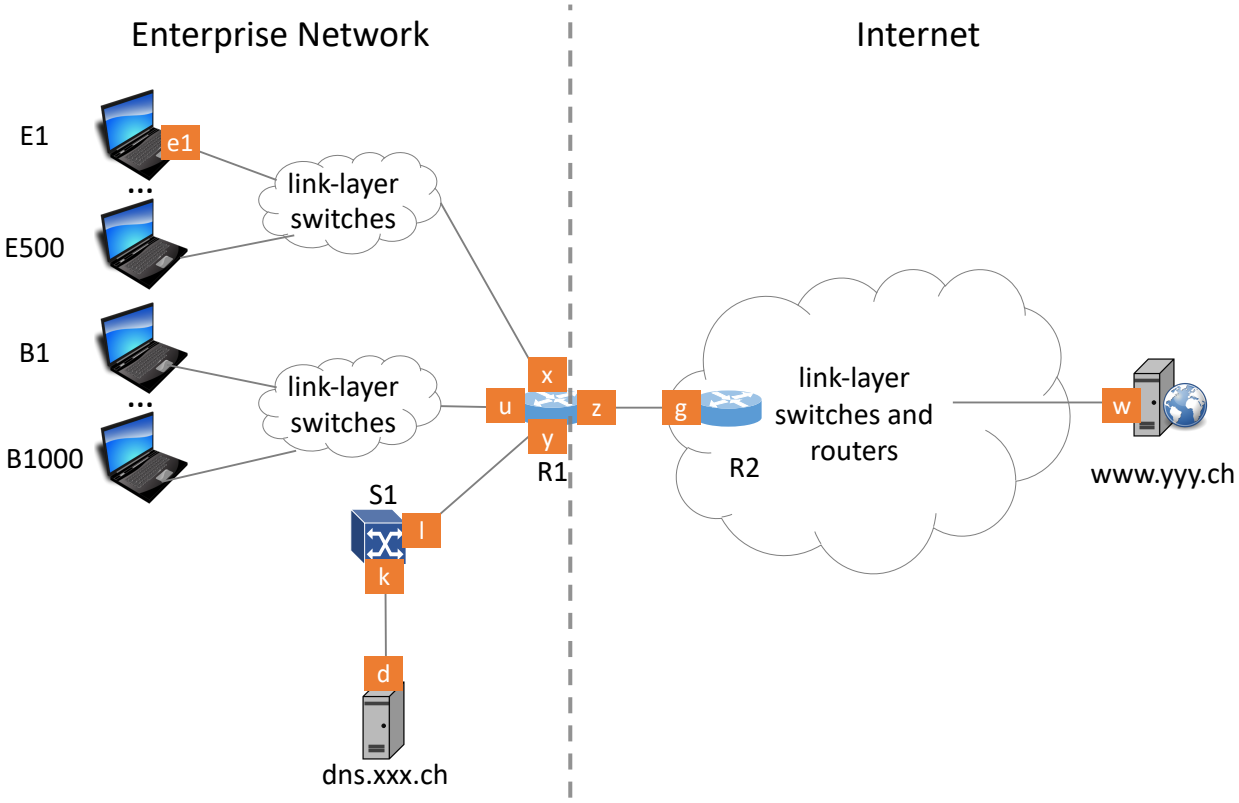

Figure 6: The Network Topology used in Problem 2