Computer Networks - Final Exam

January 27, 2022

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. Which of the following is always true for a Tier-1 Internet service provider (ISP)?
   (a) It does not have a provider. *(Correct)*
   (b) It does not have customers.
   (c) It does not have peers.
   (d) None of the above.

2. Two packets traverse the same path (sequence of devices and links). Could they experience different transmission delays?
   (a) Yes, if they have different sizes. *(Correct)*
   (b) Yes, if they carry different content.
   (c) Yes, if they encounter different levels of congestion.
   (d) Any of the above.

3. Two packets traverse the same path (sequence of devices and links). Could they experience different propagation delays?
   (a) Yes, if they have different sizes.
   (b) Yes, if they carry different headers.
   (c) Yes, if they encounter different levels of congestion.
   (d) No, unless the properties of a link change between packet traversals. *(Correct)*

4. Process P_1 running on end-system E_1 sends a continuous stream of data to process P_2 running on end-system E_2. The throughput from P_1 to P_2 is T. What does this say about the network links between E_1 and E_2?
   (a) Each link has transmission rate at least T *(Correct)*.
   (b) Each link has transmission rate at most T.
   (c) The sum of the transmission rates of the links is T.
   (d) None of the above.

5. You type in your web browser www.epfl.ch. Your browser times out (fails to show you the requested page). What could be the reason?
   (a) Your computer is not connected to the Internet.
   (b) Your local DNS server is unreachable.
   (c) The web server that hosts www.epfl.ch is broken.
   (d) Any of the above. *(Correct)*

6. You visit a given URL. Then you change laptops and you visit again the same URL, but you get a different web page (different colors, some differences in content). What is the likeliest explanation?

   (a) Your new laptop hasa a different IP address.

   (b) Your new laptop uses a different proxy web server(*).

   (c) The web server that hosts the given URL uses cookies. *(Correct)*

   (d) None of the above makes sense.

7. Alice communicates with Amazon over the secure sockets layer (SSL), as we saw in class: Alice establishes a TCP connection to Amazon and sends an "SSL Hello" message. Amazon responds with its public key and certificate. As a result of this exchange:

   (a) Alice is certain(**) that she is communicating with Amazon. *(Correct)*

   (b) Amazon is certain(**) that it is communicating with someone named Alice.

   (c) Both of the above.

   (d) None of the above.

8. TCP provides reliable data delivery. Some link layers also provide reliable data delivery. Why provide reliable data delivery at two different layers?

   (a) Just in case one of the two layers malfunctions.

   (b) Each layer recovers from different loss/corruption events.

   (c) As a performance optimization. *(Correct)*

   (d) No good reason—it's a bad idea.

9. A link-layer switch normally never does the following:

   (a) Generates or responds to an ARP request. *(Correct)*

   (b) Updates its forwarding table.

   (c) Broadcasts a packet.

   (d) Participates in the spanning tree protocol.

10. Which of the following is true?

   (a) The Internet contains multiple Autonomous Systems (ASes).

   (b) An Autonomous System (AS) typically contains multiple IP routers and IP subnets.

   (c) An IP subnet contains at least one network link and potentially link-layer switches.

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

(*) Assume that all proxy web servers work correctly.
(**) Certain, assuming the certificate authority and Amazon's keys are not compromised.

Problem 2 (25 points)

Consider the Autonomous System AS0 shown in Figure 1, which includes:

- End-systems E_1, ... E_{10000} (there are 10000 of them).
- End-systems B_1, ... B_{99} (there are 99 of them).
- DNS server `dns.xxx.com`.
- IP routers R_2, R_3, and R_4.
- Link-layer switches S_1, S_2, and others (not explicitly shown).

The orange boxes represent network interfaces.
For example, router R_3 has network interfaces f, g, and h.

All the links shown in the figure have cost 1 in each direction.

All end-systems in AS0 use `dns.xxx.com` as their local DNS server.

The intra-domain routing protocol of AS0 is the Bellman-Ford protocol with poisoned reverse that we saw in class.

Web server `www.yyy.com` and IP router R_1 belong to other ASes.

The time-to-live (TTL) of DNS records, L2-forwarding-table entries, and ARP-table entries is an hour.

You can find a copy of this network topology at the end of the exam (next to last page). 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.

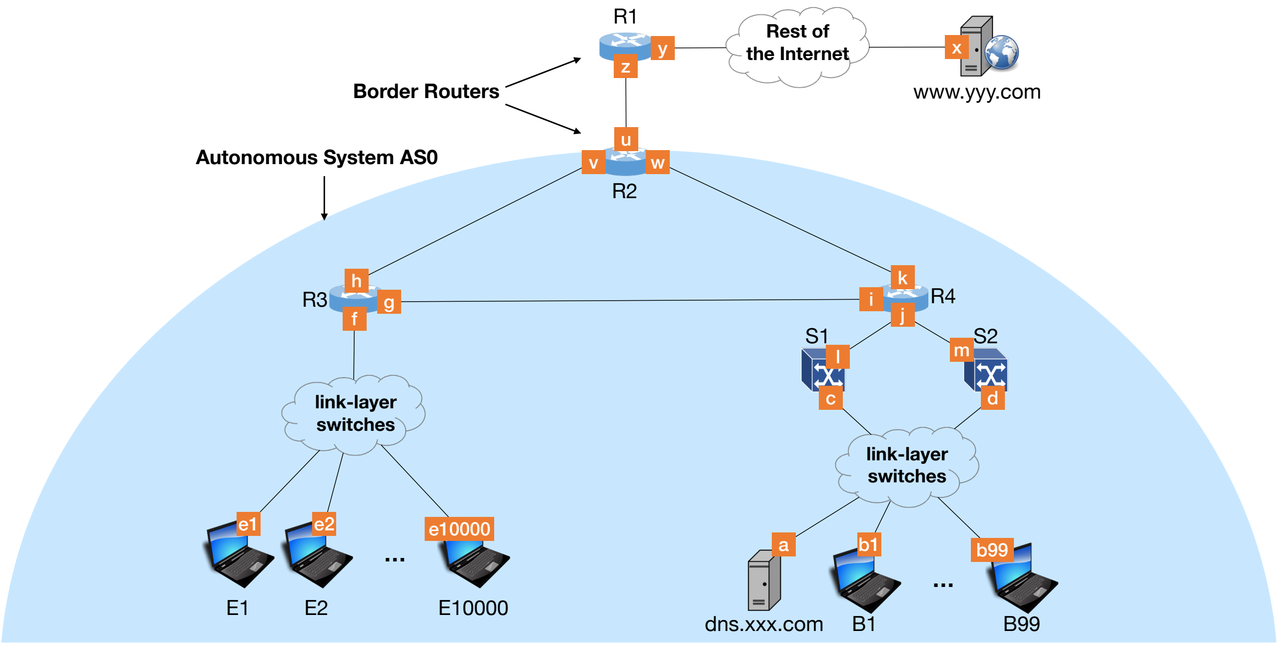

Figure 1: Network topology for Problem 2.

Question 1 (5 points):

Allocate an IP prefix to each IP subnet of AS0 (the blue bubble only) following these rules:

- All IP prefixes must be allocated from $4.0.0.0/8$.
- Each IP subnet must be allocated the smallest possible IP prefix.
- Assume one IP address per end-system and per-IP-router interface (but not for link-layer switches).
- Assume one broadcast IP address per IP subnet.
- You do not need to assume a network address per IP subnet (but it's not a mistake if you do).

Explain in one or two sentences how you compute each IP prefix.

In decreasing order of size, we have subnets: E (contains end-systems Ex), B, R23 (subnet between routers 2 and 3), R24, and R34 - five in total. We will allocate IP prefixes in this order but there are many possible solutions.

- Subnet E needs 10000 addresses for end-hosts, one for interface f, and one broadcast address. To assign 10002 addresses we need 14 bits ($2^{14} = 16384$) and thus the mask size will be $32 - 14 = 18 bits$. We can therefore assign the following address range to subnet E:

  0000 0100.0000 0000.00xx xxxx.xxxx xxxx (where the non-x digits are subnet part of each address)

  which is equivalent to:

  4.0.0.0/18

- Subnet B needs 99+1(DNS)+1(j)+1(broadcast) = 102 addresses. Therefore we need 7 bits for the host part and the mask size is $32 - 7 = 25 bits$.

  Continuing from where the previous range ends we have:

  0000 0100.0000 0000.0100 0000.0xxx xxxx or

  4.0.64.0/25

- Subnet R_{23} needs 2 (routers) + 1 (broadcast) = 3 addresses (2 bits). The mask size is $32 - 2 = 30 bits$.

  0000 0100.0000 0000.0100 0000.1000 00xx or

  4.0.64.128/30

- similarly for subnet R_{24}:

  0000 0100.0000 0000.0100 0000.1000 01xx or

  4.0.64.132/30

- similarly for subnet R_{34}:

  0000 0100.0000 0000.0100 0000.1000 10xx or

  4.0.64.136/30

Question 2 (8 points):

All link-layer switches have just been rebooted, and all end-system caches/ARP tables are initially empty.

The user of end-system E_1 visits web page `www.yyy.com`, which contains no embedded objects (e.g., no images). Immediately after E_1's user views `www.yyy.com`, the user of end-system E_2 visits the same web page.

State all the packets that are **received, forwarded, or transmitted by router R_3 as a result of $\underline{E_2}$'s actions and until E_2's user can view the web page**. For example, if a packet follows path $E_2 \rightarrow R_3 \rightarrow$ ...`www.yyy.com`, then you should state it 2 times: when it is received by R_3, and when it is forwarded by R_3.

Answer by filling in Table 1. 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_2	broadcast	-	-	-	-	-	ARP request for f's MAC
2	f	e_2	-	-	-	-	-	ARP reply
3	e_2	f	e_2	a	UDP	2000	53	DNS request for x's IP
4	g	i	e_2	a	UDP	2000	53	DNS request for x's IP
5	i	g	a	e_2	UDP	53	2000	DNS response
6	f	e_2	a	e_2	UDP	53	2000	DNS response
7	e_2	f	e_2	x	TCP	3000	80	TCP SYN
8	h	v	e_2	x	TCP	3000	80	TCP SYN
9	v	h	x	e_2	TCP	80	3000	TCP SYN ACK
10	f	e_2	x	e_2	TCP	80	3000	TCP SYN ACK
11	e_2	f	e_2	x	TCP	3000	80	HTTP GET index
12	h	v	e_2	x	TCP	3000	80	HTTP GET index
13	v	h	x	e_2	TCP	80	3000	HTTP OK
14	f	e_2	x	e_2	TCP	80	3000	HTTP OK

Table 1: Packets received, forwarded, or transmitted by router R_3 in Question 2.

Question 3 (5 points):

(a) In which routing protocol do routers R_1 and R_2 participate together? Justify your answer.

R_1 and R_2 are border routers of two different ASes and therefore participate in BGP—specifically external BGP (eBGP).

(b) Show the entry or entries of R_1's forwarding table that match/matches packets addressed to AS0. Keep in mind that only the part of $4.0.0.0/8$ that you allocated to AS0's IP subnets belongs to AS0; the rest of $4.0.0.0/8$ may be allocated/belong to other ASes.

Destination IP prefix	Output interface
4.0.0.0/18	z
4.0.64.0/25	z
4.0.64.128/29*	z
4.0.64.136/30	z

Table 2: R_1's forwarding table for Question 3(b).

* BGP router R1 merged subnets 4.0.64.128/30 and 4.0.64.132/30 since their first 29 bits are the same and the address range is continuous. 4.0.0.0/17 was also considered a correct response despite the fact that it includes address ranges that are not allocated to AS0.

Question 4 (5 points):

Assume the link between IP routers R_2 and R_4 breaks, and the intra-domain routing protocol of AS0 converges again.

(a) Which IP routers send/receive new routing announcements as a result of the failure? Justify your answer.

Routers R_2, R_3, and R_4 exchange routing announcements because they participate in the same intra-domain routing protocol.

(b) Show the forwarding table of IP router R_3 after the failure and subsequent convergence of the intra-domain routing protocol by filling in Table 3.

Destination IP prefix	Output interface
4.0.0.0/18	f
4.0.64.0/25	g
4.0.64.128/30	h
4.0.64.132/30	h/g*
4.0.64.136/30	g
other	h

Table 3: R_3's forwarding table for Question 4(b).

*Not listing 4.0.64.132/30 was also considered to be correct.

(c) Will IP router R_3 have any poisoned reverses in its Bellman-Ford table? If yes, how many? Justify your answer.

R_3's Bellman-Ford table will have two poisoned reverses. R_4 routes traffic to R_2 through R_3. Therefore, R_4 will inform R_3 that R_3 shouldn't route its traffic through R_4. What applies to R_4 also applies to R_2.

Question 5 (2 points):

Assume the link between IP router R_4 and link-layer switch S_2 breaks.

Which IP routers send/receive new routing announcements as a result of this failure? Justify your answer.

None of the IP routers send or receive routing announcements as a result of this failure. This is because this failure happens inside an IP subnet: forwarding traffic inside an IP subnet is the responsibility/concern of the L2 switches, not the IP routers.

Problem 3 (20 points)

Assume the following for all the questions in this problem:

- Alice and Bob communicate using TCP at the transport layer.
- Fast Retransmit/Fast Recovery are DISABLED.
- The maximum segment size is $\texttt{MSS} = 1$ byte.
- The TCP timeout is 1.5 RTT, where RTT is the sender's estimate of the round trip time from sender to receiver. Both RTT and the TCP timeout stay fixed throughout the TCP connection.
- Transmission delays are negligible.
- A TCP receiver sends an ACK every time it receives a data segment.
- Alice and Bob do not use any security mechanism (like SSL), i.e., any network device that sees their packets could modify them without Alice and Bob realizing it.

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):

A process running on Alice's computer establishes a TCP connection with a process running on Bob's computer and sends to it some large number of bytes. Bob's process does not send any data to Alice's process.

No segment, sent by Alice or Bob, is corrupted or dropped. The 4th segment sent by Alice (not the TCP SYN, not the 1st or the 2nd segment carrying data, the 3rd segment carrying data) takes 2RTT to be acknowledged by Bob. All other segments sent by Alice take RTT to be acknowledged by Bob.

Show all the segments sent by Alice and Bob, including connection setup, until the point where Alice has transmitted 7 distinct bytes of data. Use the diagram in Figure 2 on the next page. You can find a copy of the same diagram at the end of the exam, to first do a draft, if wou wish.

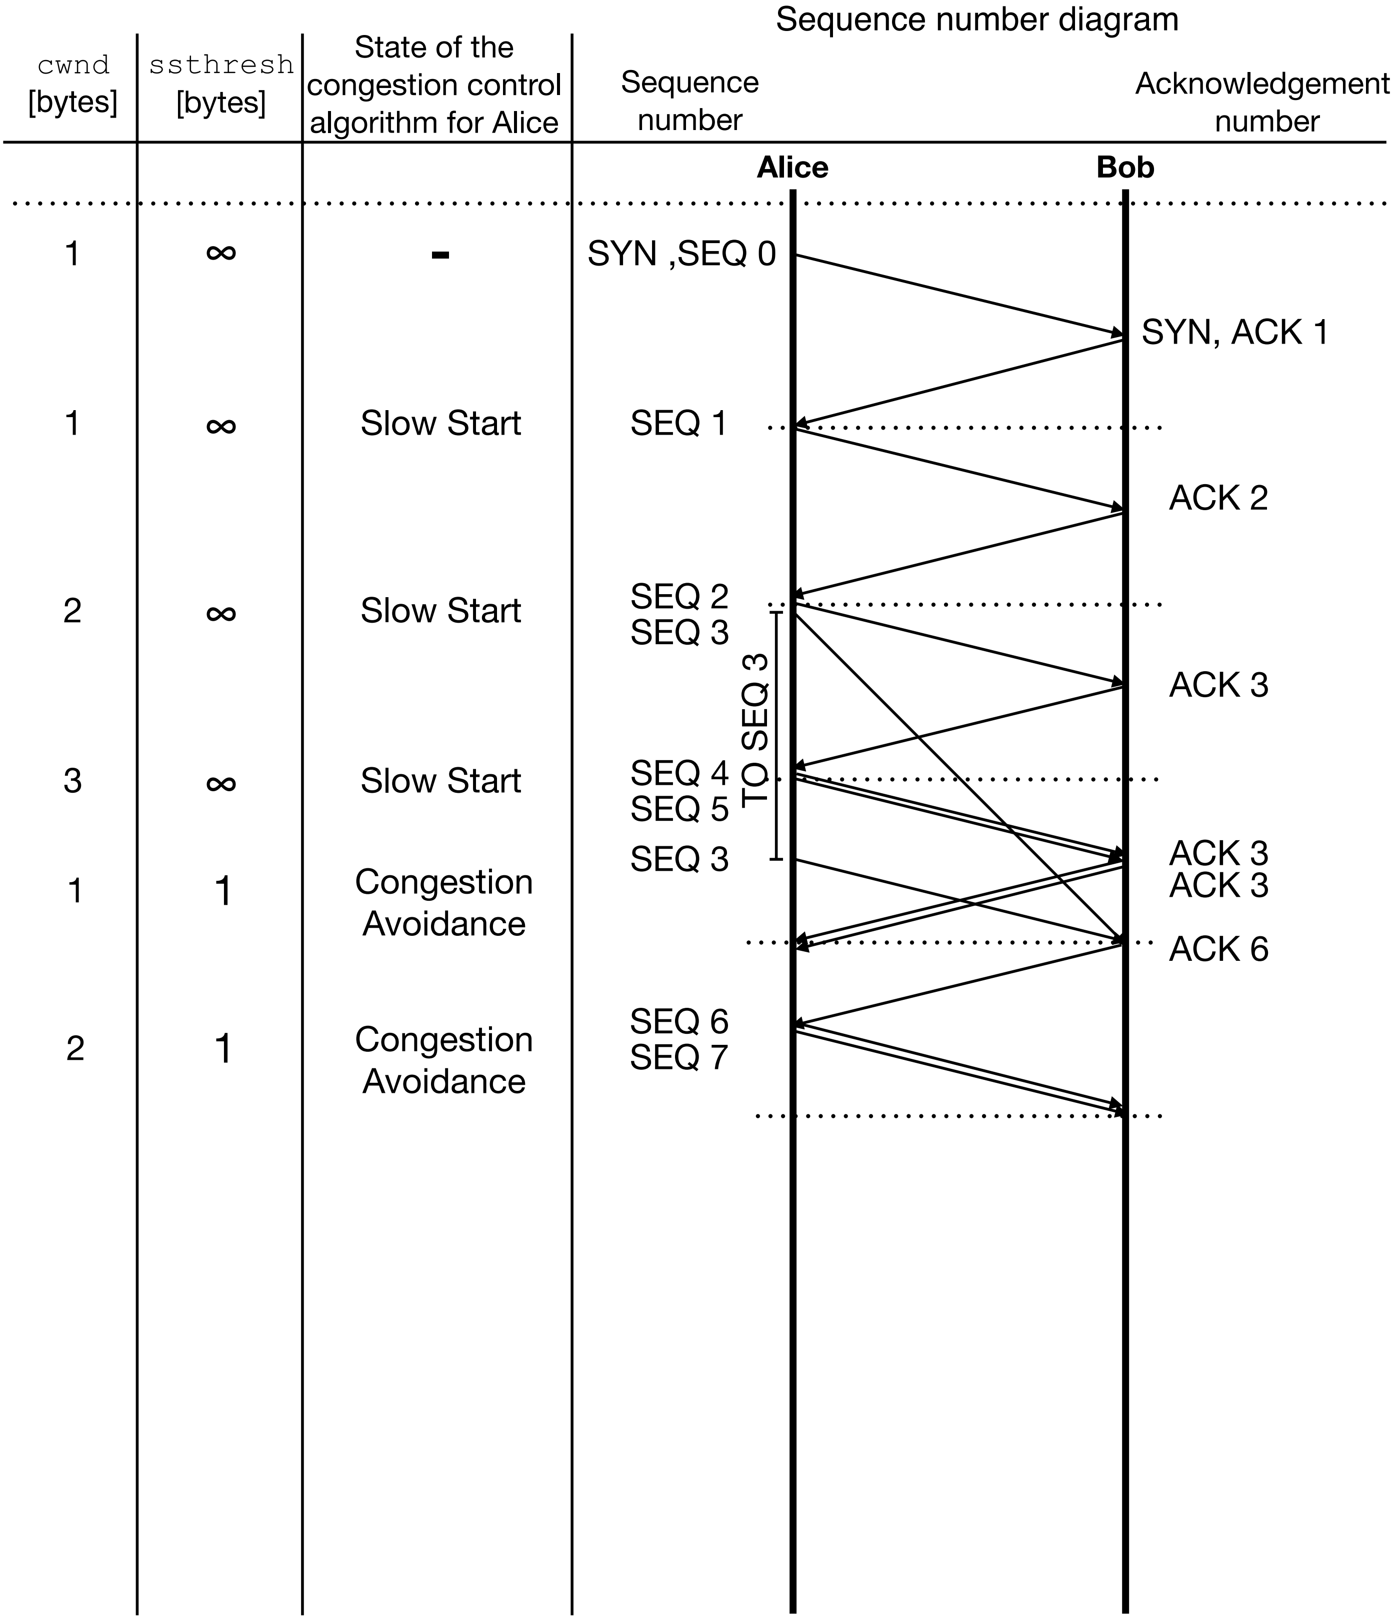

cwnd [bytes]	ssthresh [bytes]	State of the congestion control algorithm for Alice
1	∞	-
1	∞	Slow Start
2	∞	Slow Start
3	∞	Slow Start
1	1	Congestion Avoidance
2	1	Congestion Avoidance

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

Question 2 (5 points):

Fig. 3 shows the network topology between Alice and Bob:

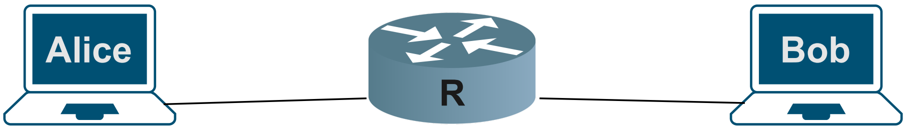

Figure 3: Network topology for Question 2.

The operator of IP router R wants to force Alice to send up to 100 bytes every RTT. Ideally, the operator would like the throughput from Alice to Bob to be as stable as possible at 100 bytes every RTT (e.g., ideally, the throughput should not fluctuate between 1 to 100 bytes per RTT). How should router R behave to achieve this goal? Describe clearly what R will do and how that will achieve the operator's goal.

The sending rate of a TCP source is determined by: (1) flow control and (2) congestion control.

One option is for R to manipulate Alice's flow control: modify the "receiver window" field in Bob's ACKs and set it such that Alice sends 100 bytes every RTT.

Another option is for R to manipulate Alice's congestion control: whenever Alice sends close to 100 bytes in an RTT, R drops one of Alice's packets, hence causes a timeout and a window reset to 1 MSS. However, this leads the congestion window to fluctuate between 1 and 100 bytes, so it is not an ideal solution.

Question 3 (5 points):

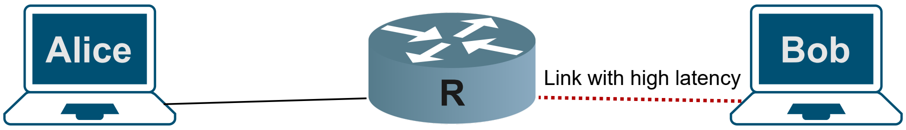

Figure 4: Network topology for Question 3.

The operator of IP router R wants to hide from Alice the fact that the link between itself and Bob has unusually high latency. How should router R behave to achieve this goal, while affecting Alice's throughput to Bob as little as possible? Describe clearly what R will do and how that will achieve the operator's goal.

R can pretend to be Bob and ACK Alice's segments as soon as it receives them (and drop the real ACKs coming from Bob). Alice thinks the ACKs are coming from Bob. This way, Alice does not "see" the high latency of the link between R and Bob.

Question 4 (5 points):

(a) Fig. 5 shows part of the server code that was given as a solution to Lab 5.

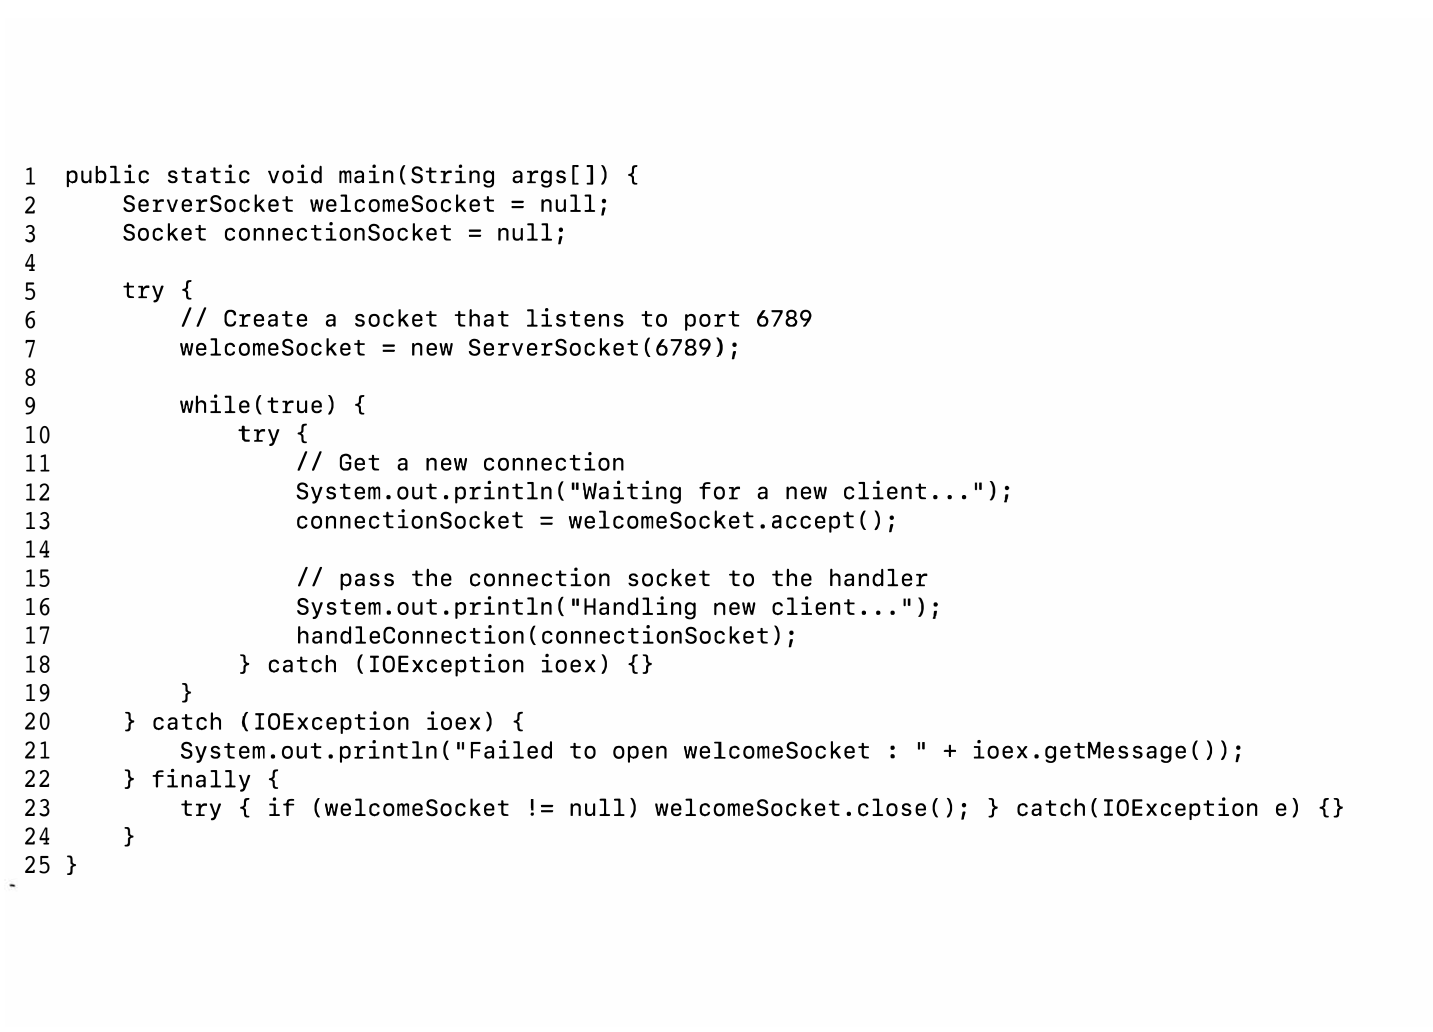

```
1  public static void main(String args[]) {
2      ServerSocket welcomeSocket = null;
3      Socket connectionSocket = null;
4
5      try {
6          // Create a socket that listens to port 6789
7          welcomeSocket = new ServerSocket(6789);
8
9          while(true) {
10             try {
11                 // Get a new connection
12                 System.out.println("Waiting for a new client...");
13                 connectionSocket = welcomeSocket.accept();
14
15                 // pass the connection socket to the handler
16                 System.out.println("Handling new client...");
17                 handleConnection(connectionSocket);
18             } catch (IOException ioex) {}
19         }
20     } catch (IOException ioex) {
21         System.out.println("Failed to open welcomeSocket : " + ioex.getMessage());
22     } finally {
23         try { if (welcomeSocket != null) welcomeSocket.close(); } catch(IOException e) {}
24     }
25 }
```

Figure 5: Server code for Question 4(a).

Which line is invoked (executed) when the server receives a TCP SYN segment? Which line causes a TCP SYN ACK segment to be sent in response?

Line 13 for both questions.

(b) Google likes TCP's reliable data delivery(*) but not TCP's performance. So, it wants Google Chrome browsers and Google's web servers to communicate reliably but without using TCP. Is this possible? I.e., is it possible for Google Chrome browsers and Google's web servers to achieve the same kind of reliable data delivery that they would achieve if they communicated over TCP but without actually communicating over TCP? Describe how or explain why it is not possible.

(*) Packets that are corrupted or dropped by the network are retransmitted by the sender until received or until some timeout occurs.

Yes, it is possible. Google can modify the software of Chrome web browsers and Google web servers, such that they use UDP instead of TCP at the transport layer, and such that they implement themselves reliable data delivery (i.e., reliable data delivery is implemented at the application layer as opposed to the transport layer).

This is similar to what you did at the last sockets lab, where you wrote client and server code that implement reliable data delivery themselves.

Scratch Paper

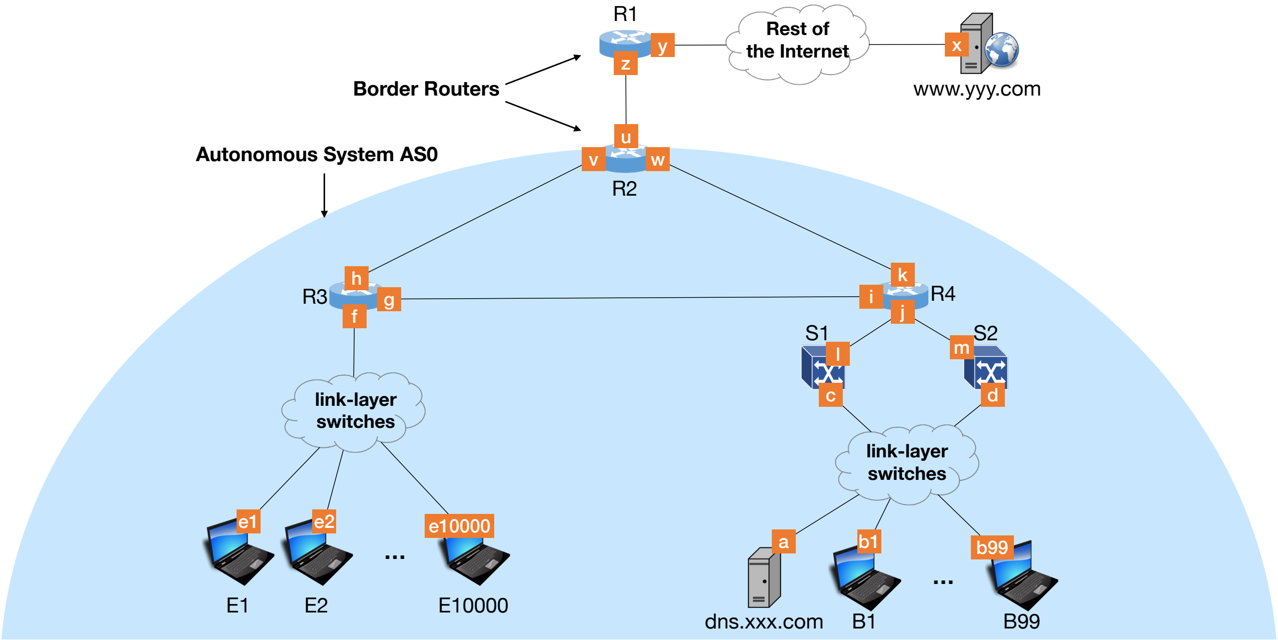

Figure 6: The Network Topology used in Problem 2.

cwnd [bytes]	ssthresh [bytes]	State of the congestion control algorithm for Alice
1	∞	-
1	∞	Slow Start
2	∞	Slow Start
3	∞	Slow Start
1	1	Congestion Avoidance
2	1	Congestion Avoidance

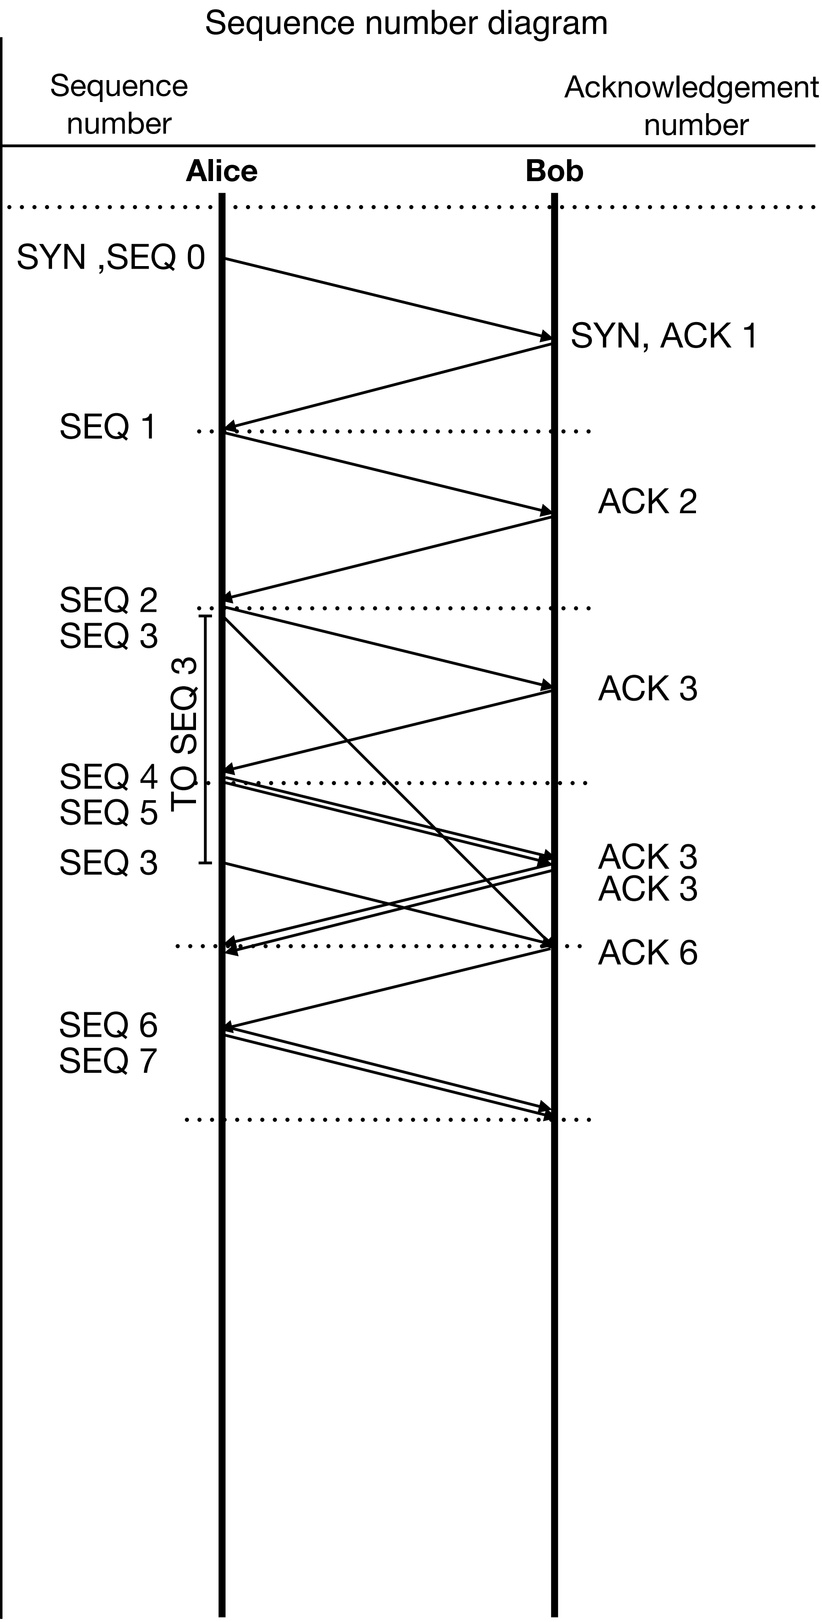

Figure 7: Copy of the sequence diagram to be completed for Problem 3, Question 1.