Computer Networks - Final Exam

January 23, 2024

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		29
3		16
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. The transmission delay of a packet on a link depends on:
   (a) The packet's size and the transmission rate of the link. *(Correct)*
   (b) The link's length and propagation speed.
   (c) The queuing delay experienced by the packet.
   (d) All of the above.

2. The queuing delay experienced by a packet may be affected by:
   (a) The sizes of the other packets with which it shares the network.
   (b) The transmission rates of the links that it crosses (which affect the transmission delays of the aforementioned other packets).
   (c) The processing delays of the packet switches that it crosses.
   (d) All of the above. *(Correct)*

3. Link L_1 has twice the transmission rate of link L_2. What can you say about L_1's lower propagation delay?
   (a) It is lower than L_2's.
   (b) It is half as much as L_2's.
   (c) It is twice as much as L_2's.
   (d) Nothing. Propagation delay is independent from transmission rate. *(Correct)*

4. Two end-systems can communicate over path P_1 or path P_2. When they communicate over P_1, they achieve twice the throughput they achieve when they communicate over P_2. What can you say about the end-to-end delay experienced by packets on P_1?
   (a) It is lower than P_2's.
   (b) It is half as much as P_2's.
   (c) It is twice as much as P_2's.
   (d) I don't have enough information to say anything. *(Correct)*

5. Which of the following scenarios may benefit from statistical multiplexing of network resources?
   (a) One end-system connected to one server over a dedicated link (not used by anyone else).
   (b) Many end-systems connected to the same server, each one over a separate, dedicated link.
   (c) Many end-systems connected to the same server over the Internet. *(Correct)*
   (d) All of the above.

6. Consider an IP packet that has encapsulated a TCP segment. Which of the following is true?
   (a) The IP packet's header contains a TCP header.
   (b) The IP packet's data (the part that comes after the header) contains a TCP header. *(Correct)*
   (c) The TCP segment's data contains an IP header.
   (d) Any of the above could be true.
7. What does it mean that peer-to-peer (P2P) "scales better" than client-server?
   (a) In P2P, the time it takes for all downloaders to download a file grows more slowly with the number of downloaders. *(Correct)*
   (b) In P2P, the time it takes for the fastest downloader to download a file grows more slowly with the number of downloaders.
   (c) In P2P, it always takes less time for all downloaders to download a file.
   (d) All of the above.
8. Alice sends a message to Bob. In which of the following scenarios can they achieve confidentiality?
   (a) They share a secret key that nobody else knows.
   (b) Alice knows Bob's true public key.
   (c) Any of the above. *(Correct)*
   (d) None of the above.
9. Alice sends a message to Bob. In which of the following scenarios can they achieve authenticity and data integrity?
   (a) They share a secret key that nobody else knows. *(Correct)*
   (b) Alice knows Bob's public key.
   (c) Any of the above.
   (d) None of the above.
10. Which of the following attacks could be prevented by using a nonce?
   (a) A malicious end-system steals Alice's private key and impersonates Alice to Bob.
   (b) A malicious link-layer switch suppresses Alice's ARP requests.
   (c) A malicious router on the path from Alice to Bob keeps copies of Alice's messages to Bob and re-sends them to Bob at some later point. *(Correct)*
   (d) All of the above.

Problem 2 (29 points)

Suppose the Internet consisted only of the topology shown in Figure 1, which includes:

- Two Autonomous Systems (ASes), AS1 and AS2.
- End-systems A_1, ... A_{1200} (there are 1200 of them).
- End-systems B_1, ... B_5 (there are 5 of them).
- End-systems C_1, ... C_{700} (there are 700 of them).
- End-systems D_1, ... D_{10} (there are 10 of them).
- IP routers R_1, R_2, R_3, R_4, R_5, R_6 and R_7.
- Various link-layer switches (not explicitly shown).

The orange boxes represent network interfaces. For example, IP router R_1 has network interfaces e and f.

Each link between IP routers shown in the figure has routing cost 1, in each direction.

All end-systems in AS1 use B_1 as their local DNS server. D_1 is a web server with DNS name d1.ethz.ch.

The intra-domain routing protocol used within AS1 relies on the Dijkstra algorithm that we saw in class, while the intra-domain routing protocol used within AS2 relies on the Bellman-Ford algorithm that we saw in class. The inter-domain routing protocol is the same as the one used in the Internet.

The time-to-live (TTL) of DNS records and ARP-table entries is 24 hours.

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.

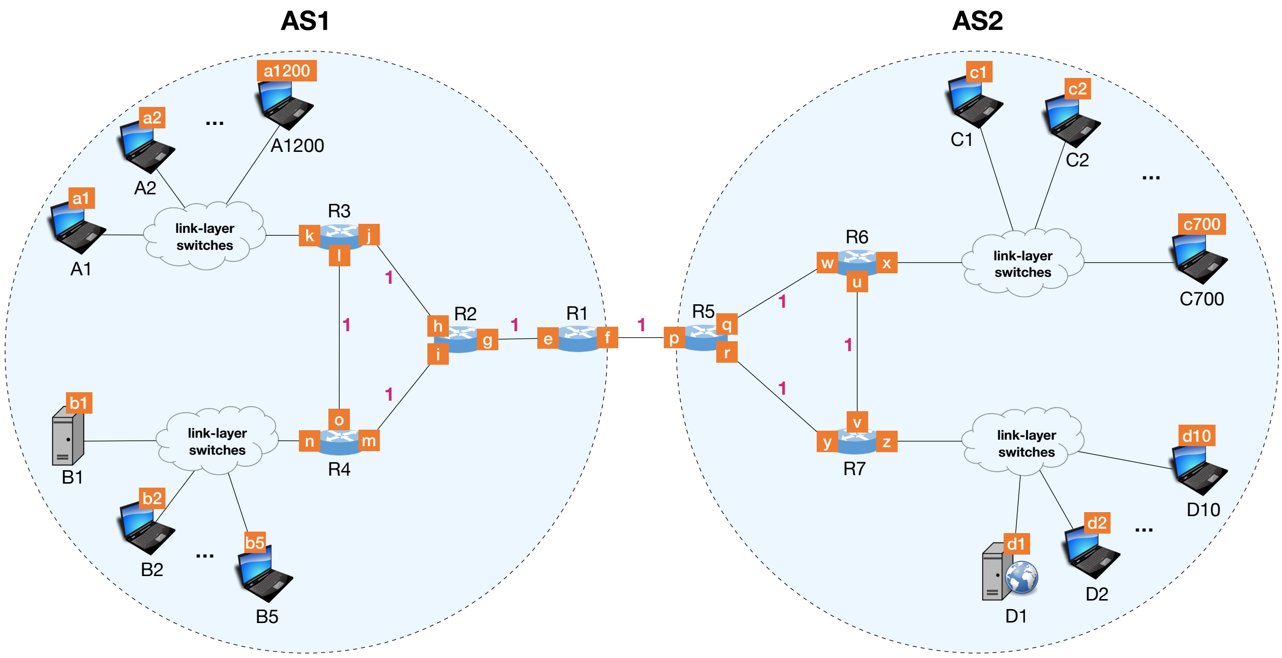

Figure 1: Network topology for Problem 2.

Question 1 (6 points):

Allocate an IP prefix to each IP subnet that contains end-systems, following these rules:

- All IP prefixes in AS1 must be allocated from 5.0.0.0/16.
- All IP prefixes in AS2 must be allocated from 8.0.0.0/16.
- 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.

There are 4 IP subnets that contain end-systems: A (contains end-systems A_x), B, C, and D. There are many possible solutions. Here is one:

- IP subnet A needs 1200 addresses for end-systems, one for interface k, and one broadcast address. To assign 1202 addresses we need 11 bits ($2^{11} = 2048$) and thus the mask size is $32 - 11 = 21$bits. We can therefore assign the following address range:

  0000 0101.0000 0000.0000 0xxx.xxxx xxxx

  which is equivalent to:

  5.0.0.0/21

- IP subnet B needs 5 addresses for end-systems, one for interface m, and one broadcast address. To assign 7 addresses we need 3 bits and thus the mask size is $32 - 3 = 29$bits. Continuing from where the previous range ends, we have:

  0000 0101.0000 0000.0000 1000.0000 0xxx

  which is equivalent to:

  5.0.8.0/29

- IP subnet C needs 700 addresses for end-systems, one for interface x, and one broadcast address. To assign 702 addresses we need 10 bits and thus the mask size is $32 - 10 = 22$bits. We can therefore assign the following address range:

  0000 1000.0000 0000.0000 00xx.xxxx xxxx

  which is equivalent to:

  8.0.0.0/22

- IP subnet D needs 10 addresses for end-systems, one for interface z, and one broadcast address. To assign 12 addresses we need 4 bits and thus the mask size is $32 - 4 = 28$bits. Continuing from where the previous range ends, we have:

  0000 1000.0000 0000.0000 0100.0000 xxxx

  which is equivalent to:

  8.0.4.0/28

Question 2 (8 points):

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

The user of end-system A_1 visits `http://d1.ethz.ch/hello.html`. Immediately after, the user of end-system A_2 visits the same web page.

State all the packets that are **sent or received by** A_2 **as a result of the action of** A_2**'s user and until the user can view the web page**.

Answer by filling in Table 1. To denote the IP address or the MAC address of interface x, write "x". 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 (the first entry in the table).

#	Source MAC	Dest MAC	Source IP	Dst IP	Transp. prot.	Src Port	Dst Port	Application & Purpose
1	a_2	broadcast	-	-	-	-	-	ARP request for k's MAC
2	k	a_2	-	-	-	-	-	ARP reply
3	a_2	k	a_2	b_1	UDP	2000	53	DNS request for d_1's IP
4	k	a_2	b_1	a_2	UDP	53	2000	DNS reply
5	a_2	k	a_2	d_1	TCP	4000	80	TCP SYN
6	k	a_2	d_1	a_2	TCP	80	4000	TCP SYN ACK
7	a_2	k	a_2	d_1	TCP	4000	80	HTTP GET index
8	k	a_2	d_1	a_2	TCP	80	4000	HTTP OK

Table 1: Packets sent or received by end-system A_2 in Question 2.

Question 3 (3 points):

Which of the packets listed in Table 1 caused R_3 to perform an IP lookup operation? For each of these packets, show the IP prefix that matched the packet's destination IP address in R_3's forwarding table, and state which routing protocol(s) caused that entry to appear in R_3's forwarding table.

- Packet 3: 5.0.8.0/29, Dijkstra
- Packet 4: 5.0.0.0/21, Dijkstra
- Packet 5: 8.0.0.0/16, BGP
- Packet 6: 5.0.0.0/21, Dijkstra
- Packet 7: 8.0.0.0/16, BGP
- Packet 8: 5.0.0.0/21, Dijkstra

Question 4 (2 points):

Which of the packets listed in Table 1 caused R_5 to perform an IP lookup operation? For each of these packets, show the IP prefix that matched the packet's destination in R_5's forwarding table, and state which routing protocol(s) caused that entry to appear in R_5's forwarding table.

- Packet 5: 8.0.4.0/28, Bellman-Ford
- Packet 6: 5.0.0.0/16, BGP
- Packet 7: 8.0.4.0/28, Bellman-Ford
- Packet 8: 5.0.0.0/16, BGP

Question 5 (10 points):

Forget about AS1. Assume that only AS2 exists.

The Bellman-Ford intra-domain routing protocol used within AS2 has converged (i.e., all routers have computed the least-cost path to each other).

*(a) Show the routing state of the three routers by filling in the tables below (the same way we filled them in class). Remember: the table of router R_5 shows the cost of the least-cost path from each AS2 router to each AS2 router, from the point of view of R_5.

Router R_5:

from / to	R_5	R_6	R_7
R_5	0	1 (R6)	1 (R7)
R_6			
R_7			

Router R_6:

from / to	R_5	R_6	R_7
R_6	1 (R5)	0	1 (R7)
R_5			
R_7			

Router R_7:

from / to	R_5	R_6	R_7
R_7	1 (R5)	1 (R6)	0
R_5			
R_6			

**(b) Suppose the link between R_5 and R_7 breaks and cannot be repaired. Show the routing state of the three routers right after the failure and after each round of message exchange, until the intra-domain routing protocol reconverges.

State right after the link failure:

Router R_5:

from / to	R_5	R_6	R_7
R_5	0	1 (R6)	2 (R6)
R_6	1	0	1
R_7			

Router R_7:

from / to	R_5	R_6	R_7
R_7	2 (R6)	1 (R6)	0
R_5			
R_6	1	0	1

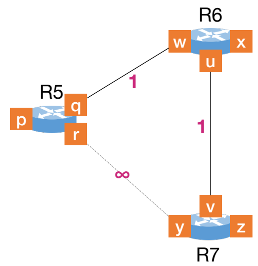

State after the first message exchange:

Router R_5:

from / to	R_5	R_6	R_7
R_5	0	1 (R6)	2 (R6)
R_6	1	0	1
R_7			

Router R_6:

from / to	R_5	R_6	R_7
R_6	1 (R5)	0	1 (R7)
R_5	0	1	∞ (Pr)
R_7	∞ (Pr)	1	0

Router R_7:

from / to	R_5	R_6	R_7
R_7	2 (R6)	1 (R6)	0
R_5			
R_6	1	0	1

State after the second message exchange:

Router R_5:

from / to	R_5	R_6	R_7
R_5	0	1 (R6)	2 (R6)
R_6	1	0	1
R_7			

Router R_6:

from / to	R_5	R_6	R_7
R_6	1 (R5)	0	1 (R7)
R_5	0	1	∞ (Pr)
R_7	∞ (Pr)	1	0

Router R_7:

from / to	R_5	R_6	R_7
R_7	2 (R6)	1 (R6)	0
R_5			
R_6	1	0	1

Problem 3 (16 points)

Question 1 (3 points):

A process running on Alice's computer has established (sometime in the past) a TCP connection with a process running on Bob's computer.

Suppose you want to compute the maximum possible rate at which Bob's process can receive data from Alice's process (when there is 0 congestion in the network that interconnects them). What information do you need in order to compute this? Justify your answer (explain why you need each piece of information, how you would use it).

- The lowest transmission rate B on the path from Alice to Bob, because this is the maximum rate at which Alice can send data to Bob (independently from the transport-layer protocol).
- The size W_r of Bob's TCP receive buffer, because it is the maximum receiver window that Bob can advertize.
- The propagation delay D from Alice to Bob and back, because this is the minimum round-trip time (RTT) between them.
- The maximum rate is $\min\{B, W_r \cdot D\}$.

Question 2 (7 points):

Assume the following:

- Alice's process always has an infinite amount of data to send to Bob's process.
- Bob's process never sends any data to Alice's process.
- The maximum TCP segment size is $\texttt{MSS} = 1$ byte.
- Bob sends an ACK every time he receives a data segment.
- Fast Retransmit and Fast Recovery are DISABLED.

At some point in time, Alice sends 8 new data segments, one after the other, to Bob, and they all get lost.

Figure 2 partially shows the 13 segments sent by Alice to Bob after the aforementioned loss event. However, it does not show which of these segments are successfully received by Bob, nor which segments (if any) are retransmissions. Also, it does not show Bob's ACKs.

Observe the information shown in Figure 2, guess what happened during the corresponding time period, and complete Figure 2, such that the following are visible:

- All the ACKs sent by Bob to Alice.
- The sequence numbers of all data segments sent by Alice to Bob.
- The acknowledgment numbers of all ACKs sent by 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 you think there were any timeouts, mark them clearly (on the side where the timeout occurs) and indicate the sequence number of the data segment that timed out.

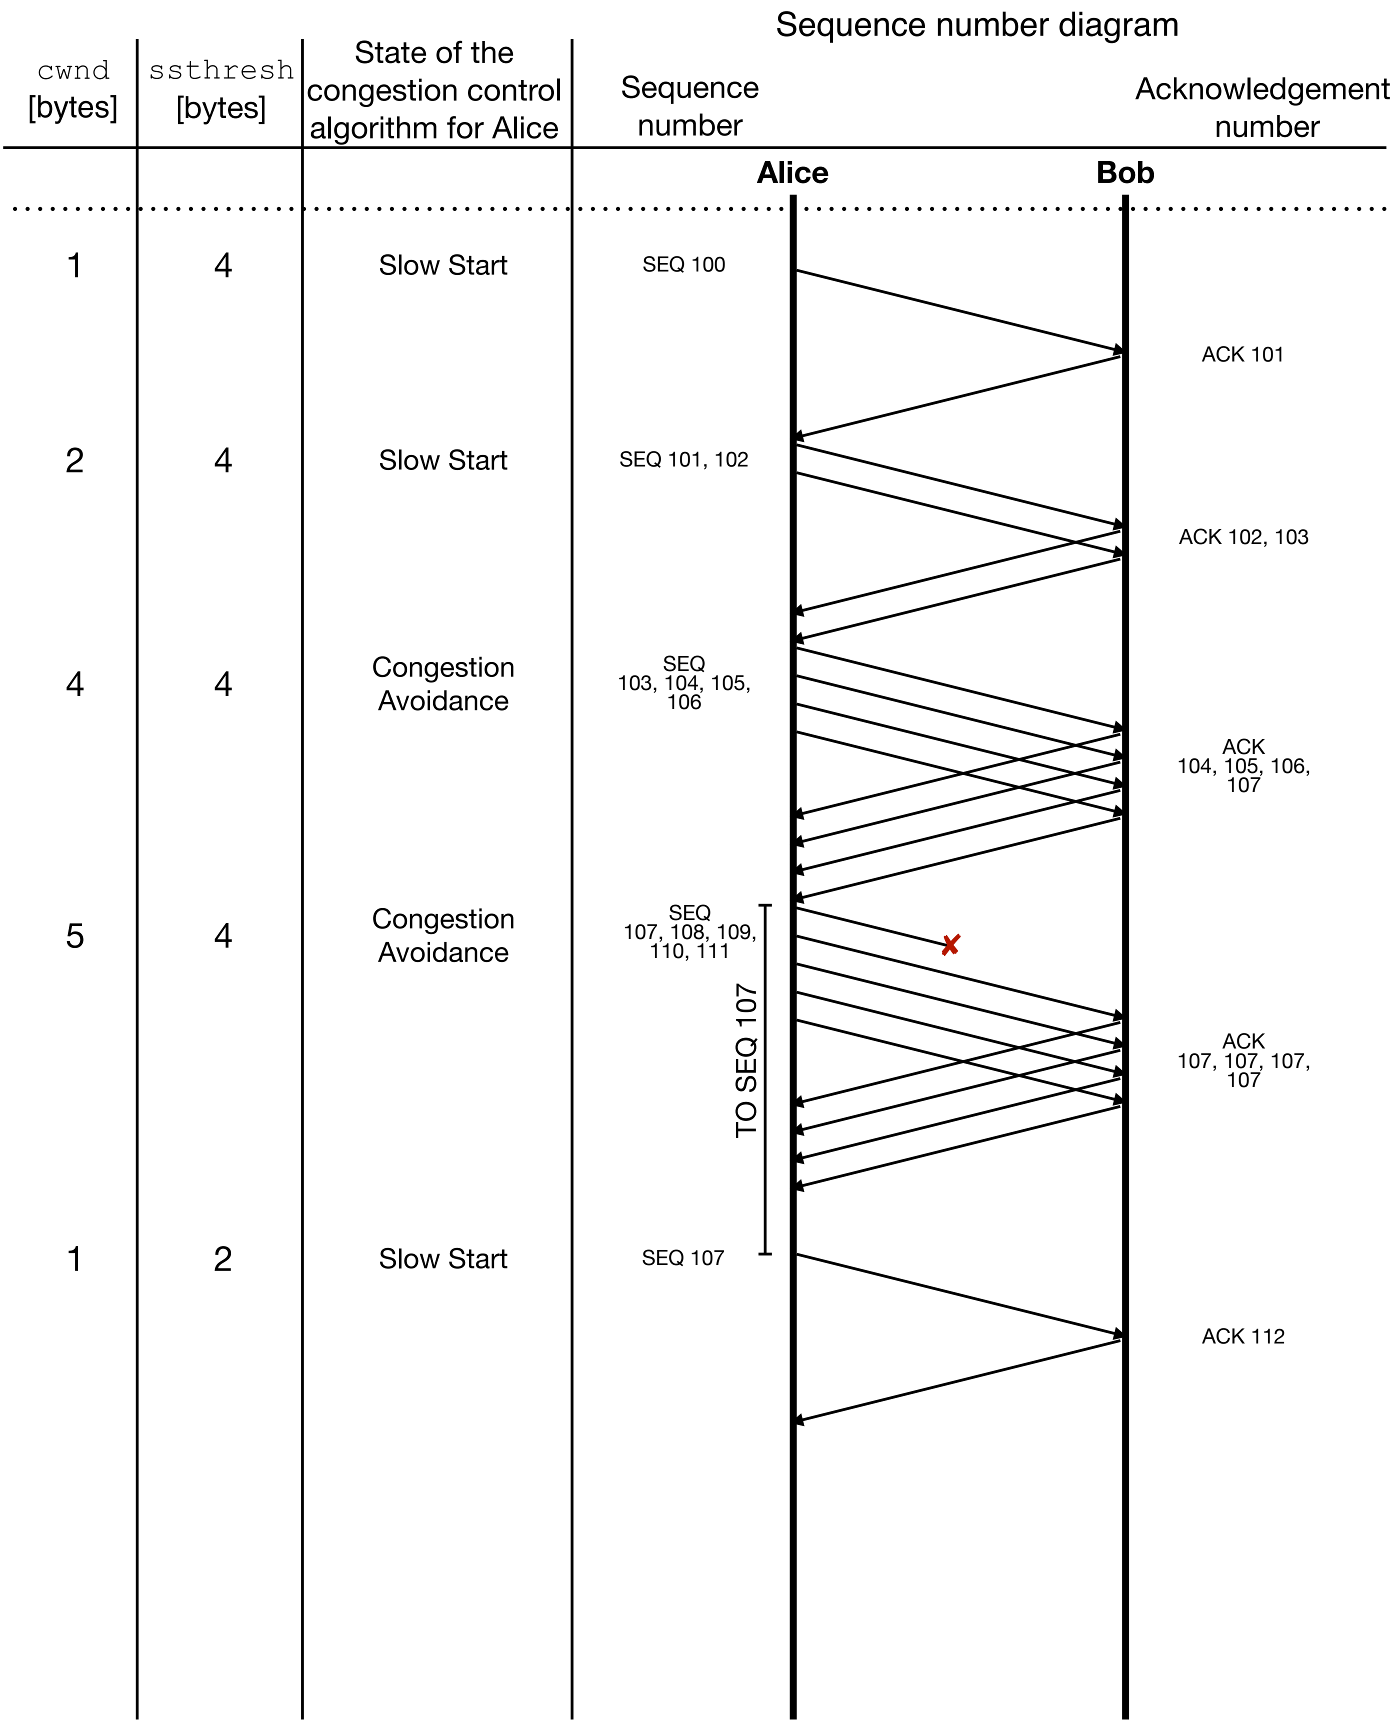

Figure 2: Sequence diagram to be completed for Question 2.

***Question 3 (6 points):**

When you answered Question 2, you guessed whether any segment(s) was(were) lost and, if yes, which one(s).

Assume the same loss scenario. E.g., if in your answer to Question 2, the x-th segment sent by Alice to Bob was lost, assume that, in this question, too, the x-th segment sent by Alice to Bob is lost.

Complete Figure 3 under the following changed condition: Fast Retransmit and Fast Recovery are ENABLED. The completed figure should show the same information as the completed figure in Question 2.

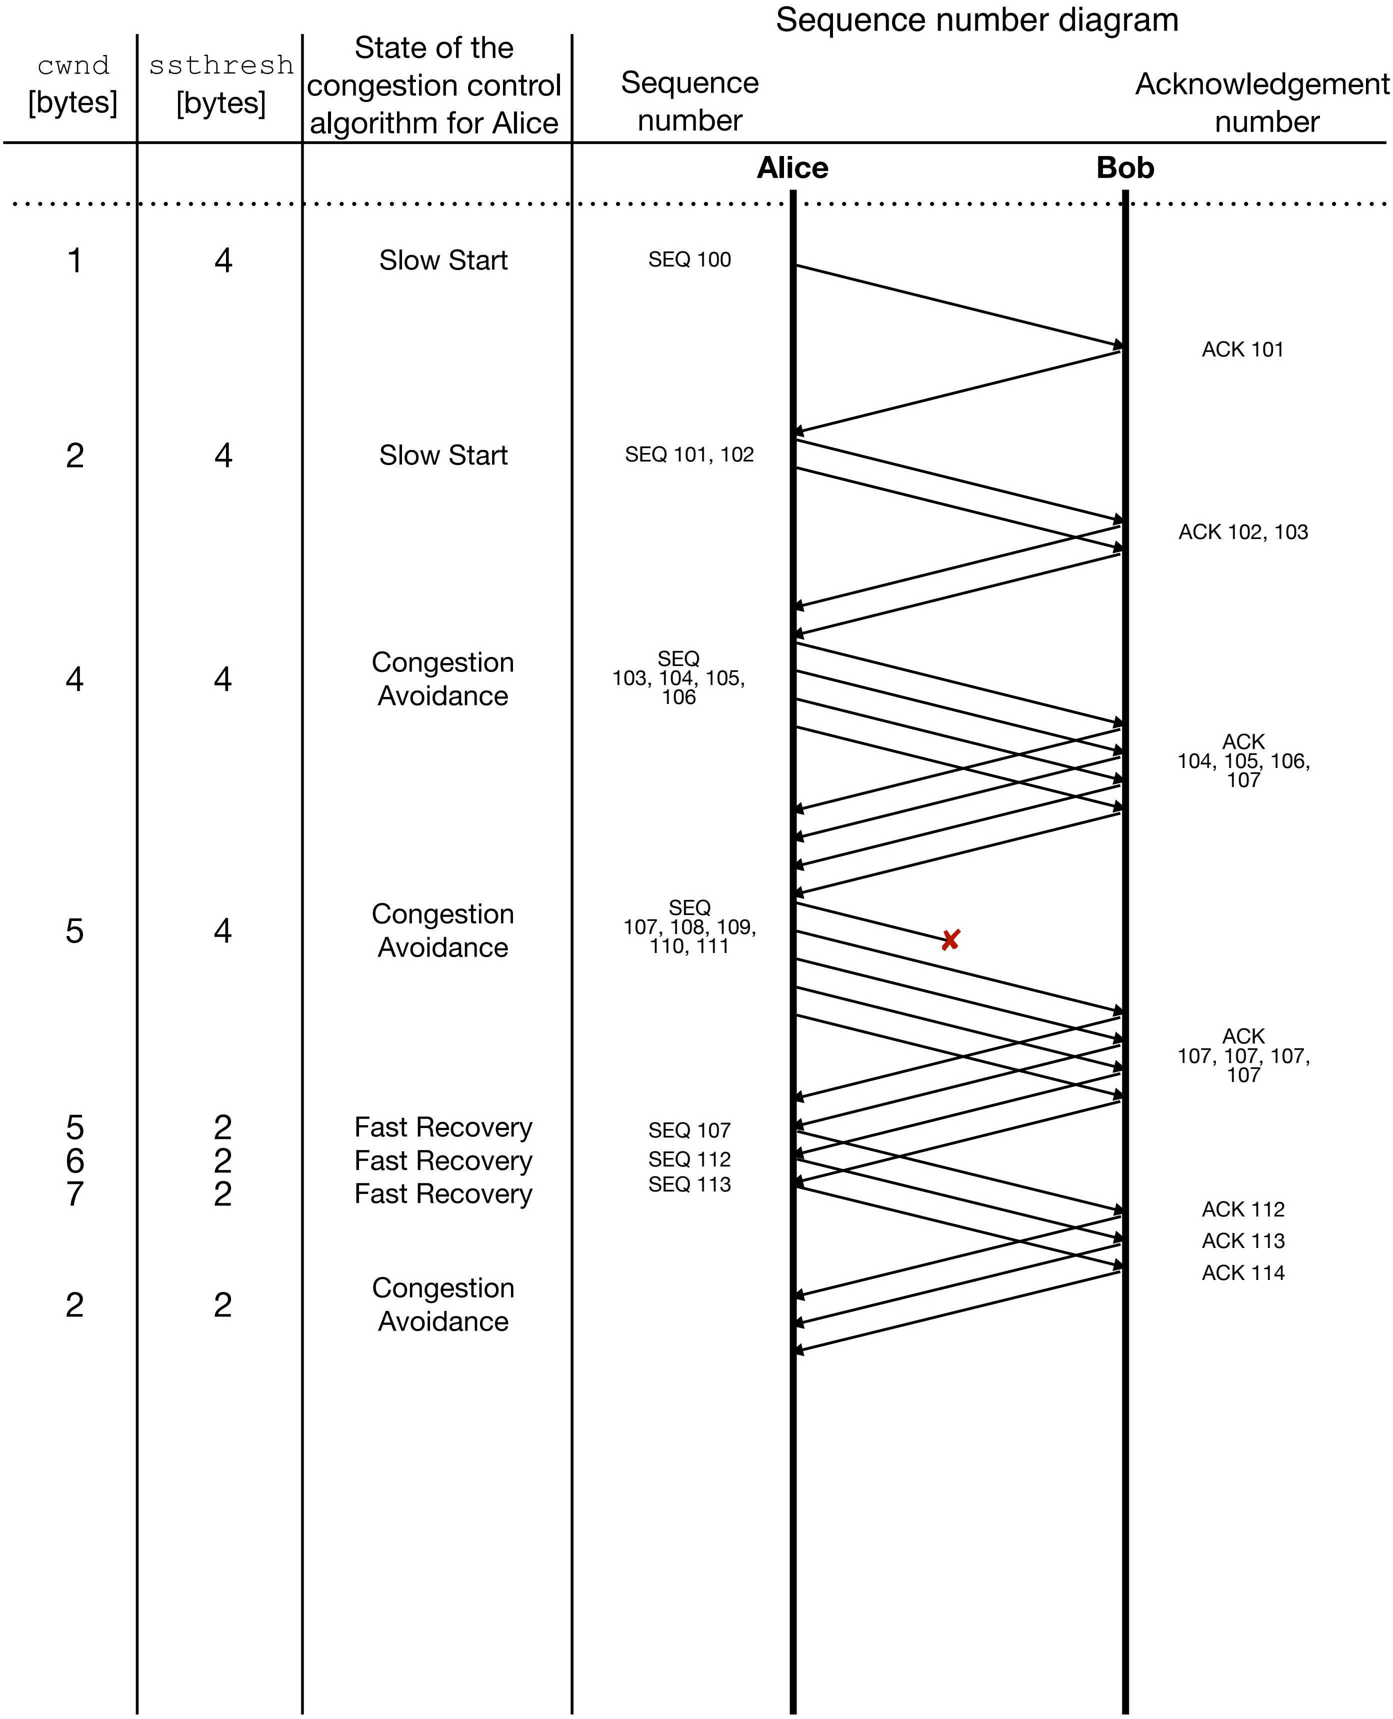

Figure 3: Sequence diagram to be completed for Question 3.

Scratch Paper

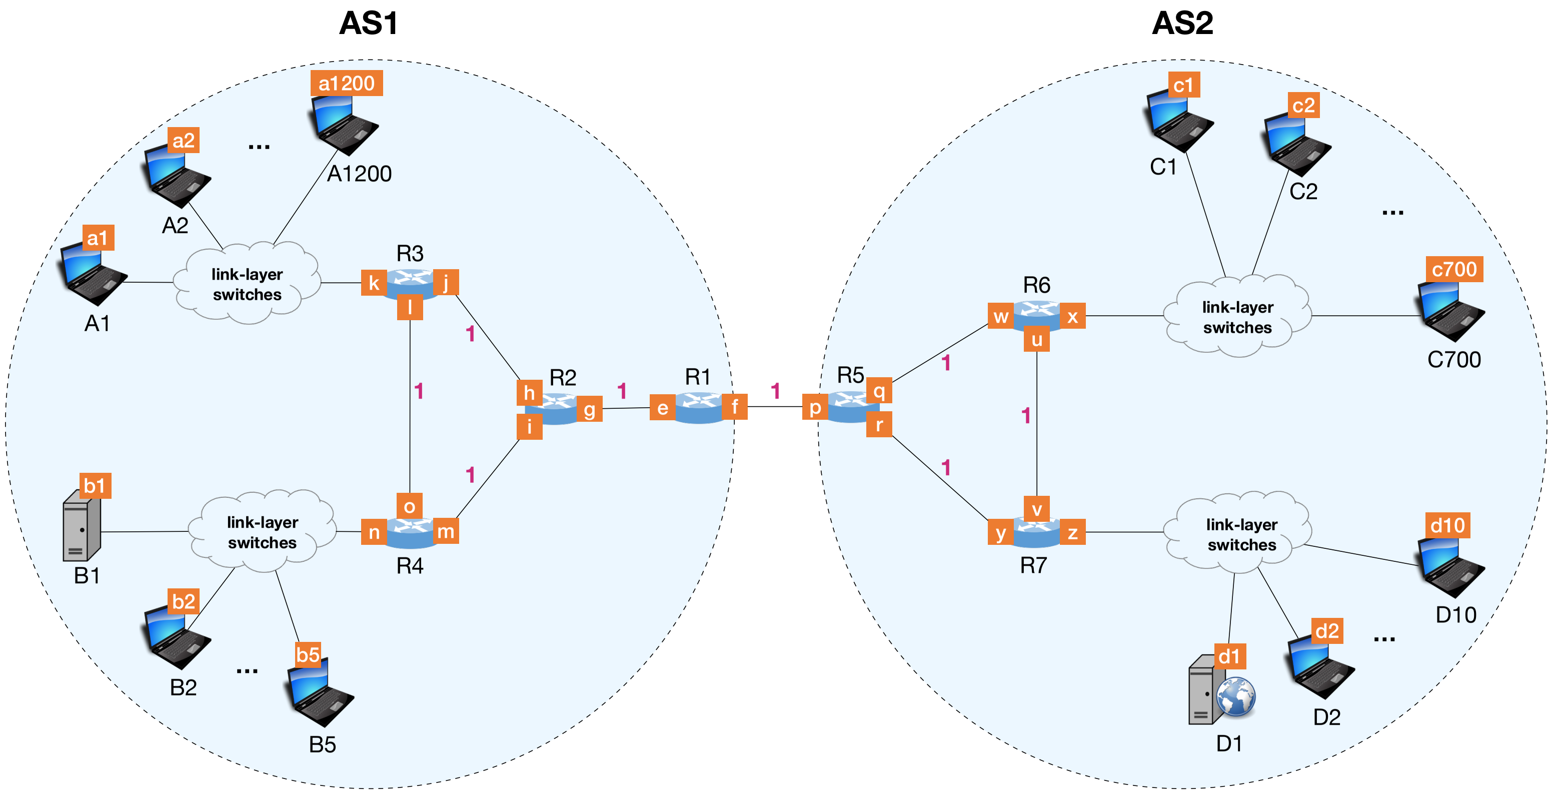

Figure 4: The Network Topology used in Problem 2.

State before the link failure:

Router R_5:

from \ to	R_5	R_6	R_7
R_5			
R_6			
R_7			

Router R_6:

from \ to	R_5	R_6	R_7
R_6			
R_5			
R_7			

Router R_7:

from \ to	R_5	R_6	R_7
R_7			
R_5			
R_6			

State after the link failure:

Router R_5:

from \ to	R_5	R_6	R_7
R_5			
R_6			
R_7			

Router R_6:

from \ to	R_5	R_6	R_7
R_6			
R_5			
R_7			

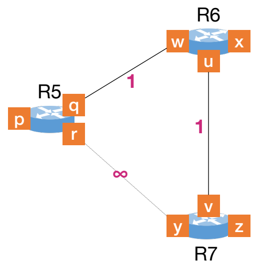

After the first exchange:

Router R_5:

from \ to	R_5	R_6	R_7
R_5			
R_6			
R_7			

Router R_6:

from \ to	R_5	R_6	R_7
R_6			
R_5			
R_7			

Router R_7:

from \ to	R_5	R_6	R_7
R_7			
R_5			
R_6			

After the second exchange:

Router R_5:

from \ to	R_5	R_6	R_7
R_5			
R_6			
R_7			

Router R_6:

from \ to	R_5	R_6	R_7
R_6			
R_5			
R_7			

Router R_7:

from \ to	R_5	R_6	R_7
R_7			
R_5			
R_6			

After the third exchange:

Router R_5:

from \ to	R_5	R_6	R_7
R_5			
R_6			
R_7			

Router R_6:

from \ to	R_5	R_6	R_7
R_6			
R_5			
R_7			

Router R_7:

from \ to	R_5	R_6	R_7
R_7			
R_5			
R_6			

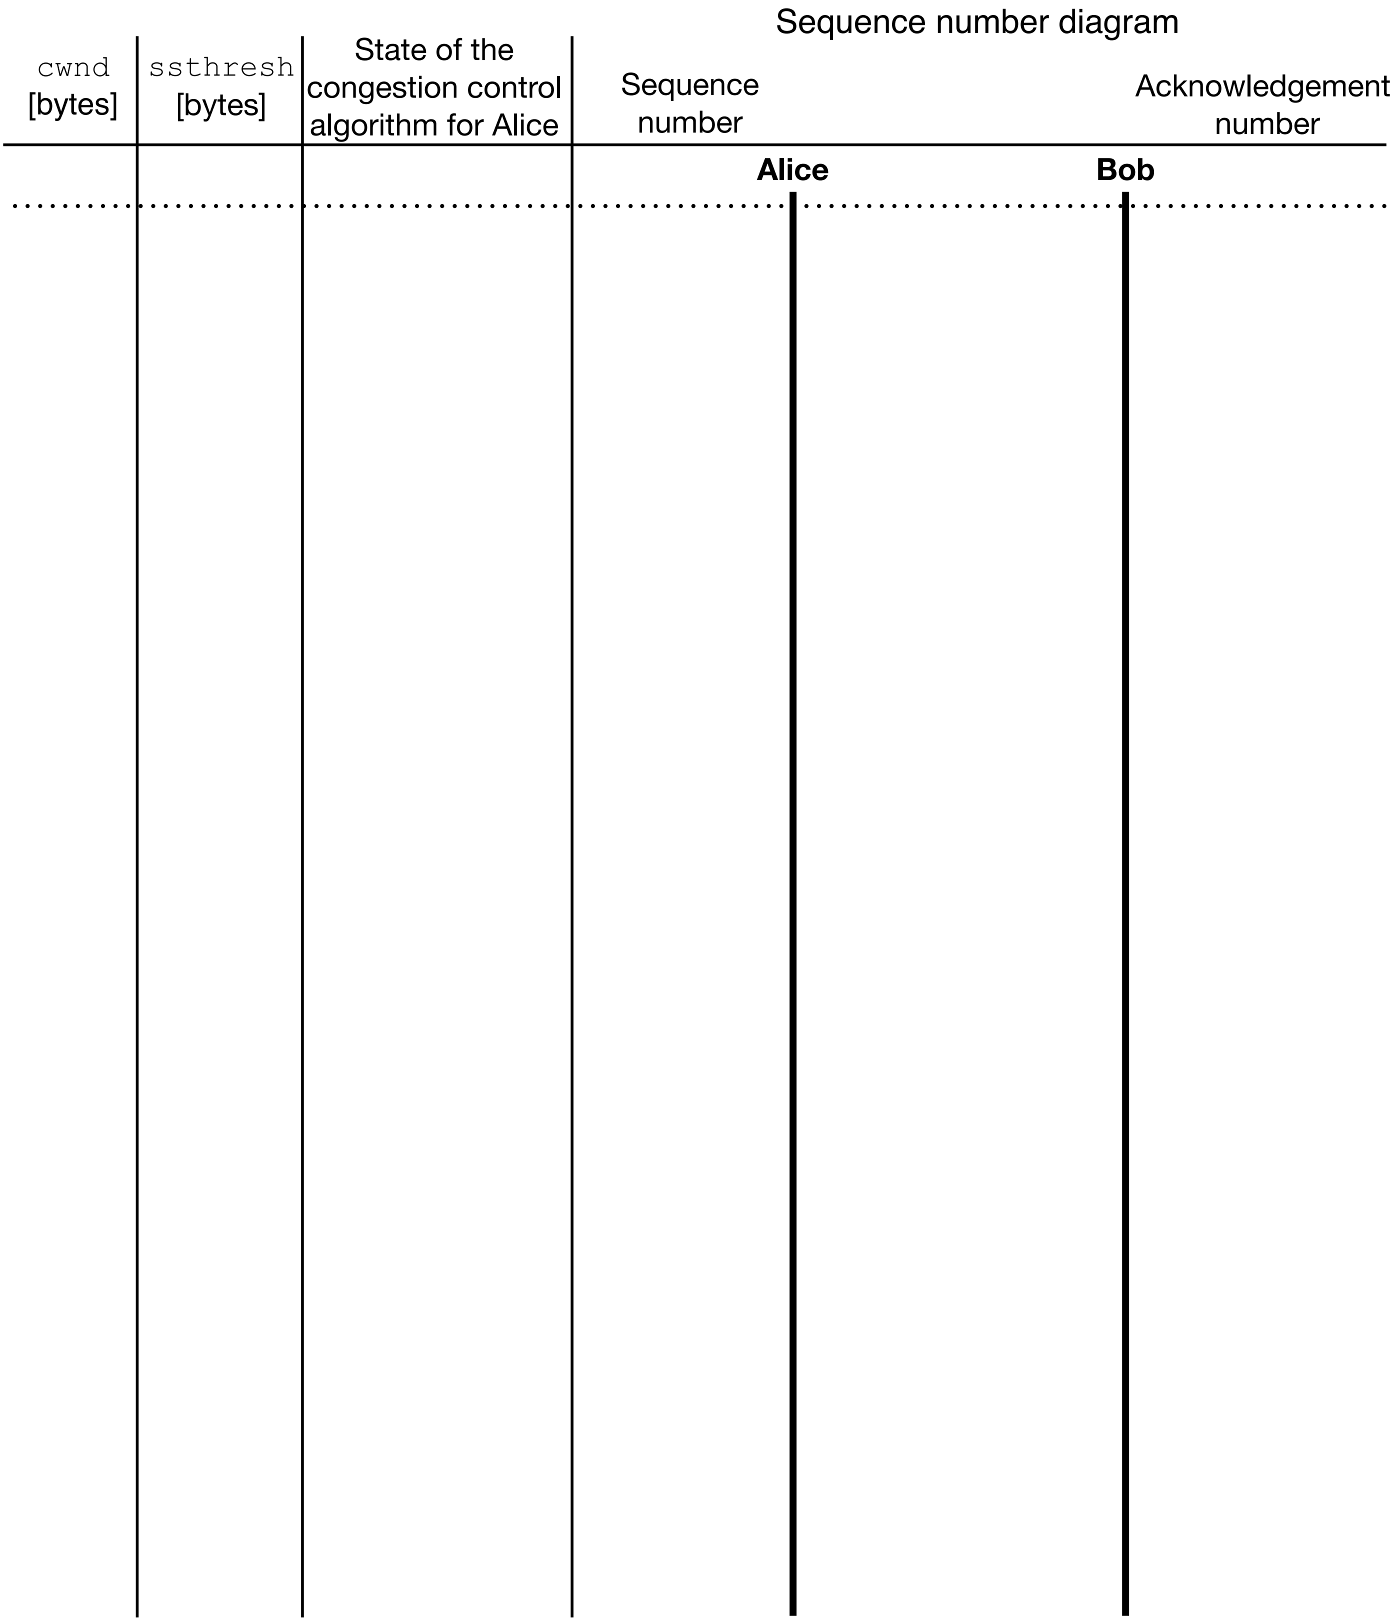

Figure 5: Copy of the sequence diagram to be completed for Problem 3, Question 2.

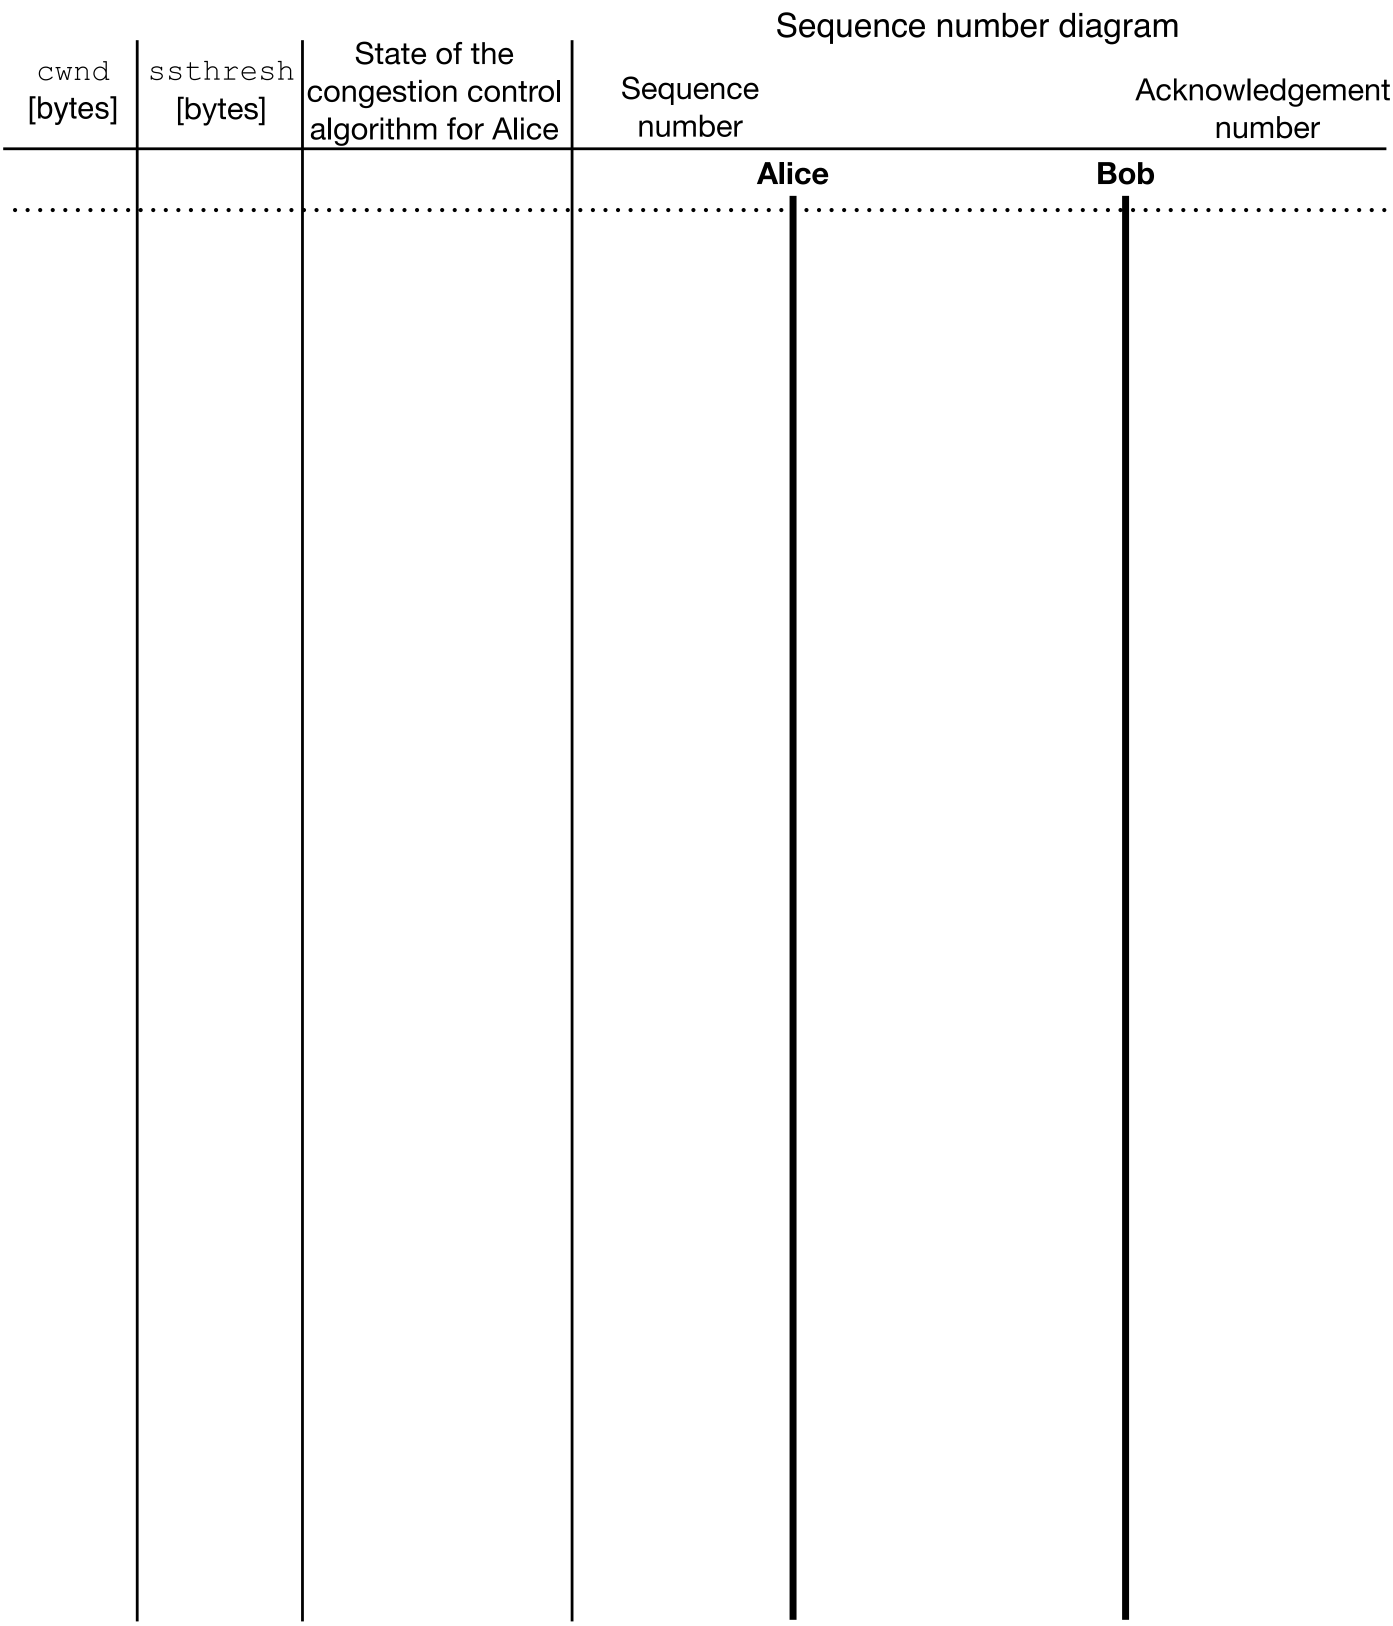

Figure 6: Copy of the sequence diagram to be completed for Problem 3, Question 3.