SCIPER: _____

First name: _____ Family name: _____

EXAM
TCP/IP NETWORKING
Duration: 3 hours

Jean-Yves Le Boudec

February 2022

INSTRUCTIONS

1. Verify that you have 4 problems + one figure sheet.
2. Write your solution into this document and return it to us (you do not need to return the figure sheet). You may use additional sheets if needed. Do not forget to write your name on **each of the four problem sheets** and **all** additional sheets of your solution.
3. All problems have the same weight.
4. Briefly justify your answer. For grading, the justification is as important as the solution itself.
5. If you find that you need to make additional assumptions in order to solve some of the questions, please describe such assumptions explicitly.
6. Figures are on a separate sheet, for your convenience.
7. No documents, no electronic equipments are allowed.

PROBLEM 1

Consider the network in the figure sheet. A, B and S are hosts. R1 is a router. B1, B2 and B3 are bridges. The cloud represents the public internet. O1, O2, , ... , O7 are observation points. All machines are dual-stack.

Unless otherwise specified, the network box N operates as a router for IPv6 and a NAT for IPv4.

All links are full duplex Ethernet. We assume that all machines are correctly configured (unless otherwise specified), proxy ARP is not used and there is no VLAN. A's default gateway is R1 and B's default gateway is R1. The bridges and routers have been running for some time and their different protocols are in steady-state. There is no other system or interface than shown on the figure.

The figure shows explicit IP addresses. The symbols A and B also represent the MAC addresses of the hosts A and B. The other MAC addresses are denoted with R1s, R1e, etc.

1. All interfaces in the private network have the same IPv4 network mask. The table below proposes values for the IPv4 network mask. Say which ones are valid by crossing the correct box.

   For this table and the next, the usual rules of quizzes apply: (i) if you cross the correct box in one row you obtain the full score for this row; (ii) if you cross zero or two boxes in one row, your score is 0 for this row; (iii) if you cross exactly one box in one row and it is incorrect your score for this row is negative, namely a penalty equal to half the score for this row.

Proposed net mask in private network	valid	invalid
255.255.255.0	☐	☐
255.255.0.0	☐	☐
255.255.240.0	☐	☐

2. All interfaces in the private network have the same IPv6 network mask. The table below proposes values for the length of the IPv6 subnet prefix. Say which ones are valid by crossing the correct box.

Proposed length of the IPv6 subnet prefix	valid	invalid
64	☐	☐
68	☐	☐
80	☐	☐

3. B is restarted; it is correctly configured but its caches are empty; B immediately sends one UDP packet to 10.2.4.3 (i.e to A), using IPv4 and with source TTL equal to 64. We assume that this packet is not lost. We observe all packets resulting from this activity at observations points O1 and O2. Write the values of the fields in the table below (if present). In each row, use as many lines as needed. The "type" field is the one contained in the MAC header (Ethertype).

at	srce MAC addr	dest MAC addr	type	srce IP addr	dest IP addr	TTL
O1						
O2						

If instead of observing the packets at O1 and O2 we also observe them at O3, O4 and O5, what do we see ?

4. A downloads a large web page from S using HTTPS over TLS over TCP over IPv4. The local port used by A is 5432. At about the same time, B also downloads a large web page from S using HTTPS over TLS over TCP over IPv4. By coincidence, the local port used by B is also 5432. The port used by HTTPS over TLS over TCP is 443. At O6 and O7 we observe the packets sent by S to A and B. Write the values of the fields in the table below.

From S to A					
At	srce IP addr	dest IP addr	protocol	srce port	dest port
O7					
O6					

From S to B					
At	srce IP addr	dest IP addr	protocol	srce port	dest port
O7					
O6					

Assume that a packet of this transfer is lost between N and R1. Which system, if any, will re-transmit the data that was lost in the packet ?

5. A downloads a large web page from S using HTTP over QUIC over IPv6. The local port used by A is 9876. At about the same time, B also downloads a large web page from S using HTTP over QUIC over IPv6. By coincidence, the local port used by B is also 9876. The port used by HTTP over QUIC at S is 443. At O6 and O7 we observe the packets sent by S to A and B. Write the values of the fields in the table below.

From S to A					
At	srce IP addr	dest IP addr	protocol	srce port	dest port
O7					
O6					

From S to B					
At	srce IP addr	dest IP addr	protocol	srce port	dest port
O7					
O6					

Assume that a packet of this transfer is lost between N and R1. Which system, if any, will re-transmit the data that was lost in the packet ?

6. In this question we assume that N acts as an HTTP tunnel. A downloads a page from S using HTTP over TLS over TCP over IPv6. A packet of this transfer is lost between N and R1. Which system, if any, will re-transmit the data that was lost in the packet ?

SCIPER: _____

First name: _____ Family name: _____

PROBLEM 2

Consider the network for problem 2 in the figure sheet. There are four ASs, A, B, C and D with routers A1, A2, B1, B2, C1, D1, R1 and R2. The physical links are shown with plain lines. Each AS uses OSPF with Equal Cost Multipath as IGP, and every router inside each AS uses OSPF. The cost of every link and every directly attached network is 1, except when otherwise specified.

The figure shows one stub network at router D1 with its IPv6 network prefix. The lower case symbols such as a1n, b1s, etc. represent IPv6 addresses.

Routers A1, A2, B1, B2, C1 and D1 use BGP with their external neighbours and as required with their internal neighbours. The routers R1 and R2 may or may not use BGP, depending on the question. No confederation or route reflector is used.

We assume that the BGP decision process use the following criteria in decreasing order of priority. BGP identifiers are router names such as A1, A2...

1. Highest LOCAL-PREF
2. Shortest AS-PATH
3. E-BGP is preferred over I-BGP
4. Shortest path to NEXT-HOP, according to IGP
5. Lowest BGP identifier of sender of route is preferred. Here comparison is lexicographic with A<B<C<D and $1 < 2$; for example A1 is preferred over A2, A2 is preferred over B1, etc...

Furthermore, we assume the following.

- Unless otherwise specified, when receiving an E-BGP announcement, every BGP routers tags it with LOCAL-PREF = 0. No other optional BGP attribute (such as MED, etc.) is used in BGP messages.
- No aggregation of route prefixes is performed by BGP.
- The policy in A, B, C and D is such that all available routes are accepted and propagated to neighbouring ASs, as long as the rules of BGP allow.
- Every router redistributes internal OSPF destinations into BGP. Unless otherwise specified, there is no other redistribution.
- Every router performs recursive forwarding-table lookup.
- Equal Cost Multi-Path routing is supported by all routers.
- When writing the AS-PATH attribute received by a router X, do not write the AS that this router is in. For example, if router X in AS A receives a route from router Y in AS B, who received it from AS C, the AS-PATH stored by X in its RIB is `AS-PATH = B C`. However, if X propagates this route over E-BGP to a router in AS D, the message sent by X will have `AS-PATH = A B C`.

1. In this question, we assume that R1 and R2 run BGP. At time t_1, BGP and OSPF have converged in all ASs.

   (a) At time t_1, what is the best BGP route to `2001:1:1::/48` selected by A2 ? From which BGP peer was it received ? What are its BGP NEXT-HOP and AS-PATH attributes ? Same question for A1, R1, and R2. Give your answers in the table below, with a short justification.

At	From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH
A1		2001:1:1::/48		
A2		2001:1:1::/48		
R1		2001:1:1::/48		
R2		2001:1:1::/48		
Justification:				

   (b) Still at time t_1, what is the best BGP route to `2001:1:1::/48` selected by B1 ? From which BGP peer was it received ? What are its BGP NEXT-HOP and AS-PATH attributes ? Give your answer in the table below, with a short justification.

At	From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH
B1		2001:1:1::/48		
Justification:				

(c) Still at time t_1, what is the list of all BGP routes received and stored by B2 in its RIB, with destination = 2001:1:1::/48 ? For every route, indicate from which BGP peer it was received from and give the BGP NEXT-HOP and AS-PATH attributes. Give your answer in the table below, with a short justification (put as many rows as necessary).

At B2 :				
From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH	Best route ?
	2001:1:1::/48 2001:1:1::/48 ...			
Justification:				

(d) Still at time t_1, B1 and B2 each have one packet to forward with destination `2001:1:1:2:3:4::1`. Assuming these packets are not lost due to transmission errors or buffer overflows, will the packet reach the router D1 ? If so, over which path (given as a sequence of routers) ?

2. In this question we assume that the network is restarted, with the following changes in the configurations

   - All routers inside AS A tag all E-BGP announcements received from AS B with LOCAL-PREF = 100; other incoming E-BGP announcements are tagged with LOCAL-PREF = 0, same as before. BGP routers in other ASs continue to tag incoming E-BGP announcements with LOCAL-PREF = 0, same as before.

   We assume that R1 and R2 run BGP. At time t_2, BGP and OSPF have converged in all ASs.

   (a) At time t_2, what is the best BGP route to `2001:1:1::/48` selected by A2 ? From which BGP peer was it received ? What are its BGP NEXT-HOP and AS-PATH attributes ?
   Same question for A1. Give your answers in the table below, with a short justification.

At	From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH
A1		2001:1:1::/48		
A2		2001:1:1::/48		
Justification:				

(b) At time t_2, what is the list of BGP routes received by R1 with destination = 2001:1:1::/48 ? Which route is selected as best route by R1 ? Give your answer in the table below, with a short justification (put as many rows as necessary).

At R1 :		
From BGP Peer	Destination Network	BGP NEXT-HOP
	2001:1:1::/48 2001:1:1::/48 ...	
Justification:		

3. In this question we assume that the network is restarted, with the following changes in the configurations of routers inside AS A:
   - R1 and R2 do not run BGP (but run OSPF).
   - A1 and A2 redistribute routes learnt by E-BGP into OSPF. At A1, the OSPF cost from A1 to such a re-distributed route is set to 101. At A2, the OSPF cost from A2 to such a re-distributed route is set to 102.
   - As in the previous question: all routers inside AS A tag all incoming E-BGP announcements incoming from AS B with LOCAL-PREF = 100; other incoming E-BGP announcements are tagged with LOCAL-PREF = 0 and BGP routers in other ASs tag incoming E-BGP announcements with LOCAL-PREF = 0.

   At time t_3, BGP and OSPF have converged in all ASs.

   (a) At time t_3, R2 has a large number of packets to send, with various destination addresses that all fall in the block `2001:1:1::/48`. Assuming these packets are not lost due to transmission errors or buffer overflows, will they reach the router D1 ? If so, over which path (given as a sequence of routers) ?

   (b) At time $t_4 > t_3$, R1 is re-configured and re-started. Unfortunately, there is a configuration error in R1: the OSPF interface database of R1 is configured with the prefix `2001:1::/32` (i.e the OPSF software believes that the network `2001:1::/32` is directly attached at R1. At time $t_5 > t_4$, OSPF and BGP have converged again and R2 has a large number of packets to send, with various destination addresses that all fall in the block `2001:1:1::/48`. Assuming these packets are not lost due to transmission errors or buffer overflows, will they reach the router D1 ? If so, over which path (given as a sequence of routers) ?

SCIPER: _____

First name: _____ Family name: _____

PROBLEM 3

Consider the network for problem 3 on the figure sheet.

- A, B and C are routers. The capacity of the links between them is 240 Mb/s. The links are full duplex with same rate in both directions.
- There are 4 unidirectional flows, as shown on the figure:
  - Flow 1 goes from S1 to D1
  - Flow 2 goes from S1 to D2
  - Flows 3 and 4 go from S2 to D1

  There is no other system and no other flow than shown on the figure. There is no other capacity constraint than the two link capacities shown on the figure. We also neglect the impact of the acknowledgement flows in the reverse direction.
- We neglect all overheads and assume that the link capacities can be fully utilized at bottlenecks.
- We call x_1 [resp. x_2, x_3, x_4] the rate of flow 1 [resp, 2,3,4].

1. Which of the following allocations, in Mb/s, are Pareto-efficient ? Justify your answer.
   (a) $x_1 = 160,\ x_2 = 80,\ x_3 = x_4 = 40$
   (b) $x_1 = x_2 = x_3 = x_4 = 50$
   (c) $x_1 = x_2 = x_3 = x_4 = 80$
   (d) $x_1 = 80,\ x_2 = 160,\ x_3 = 80,\ x_4 = 80$

2. Assume the rates x_1, x_2, x_3, x_4 of the four flows are allocated according to max-min fairness. What are all the possible allocations ? Which of those are Pareto-efficient ?

3. Assume the rates x_1, x_2, x_3, x_4 of the four flows are allocated according to proportional fairness. What are all the possible allocations ? Which of those are Pareto-efficient ?

4. In this question flow 1 is using UDP and sends at a constant rate equal to 100 Mb/s. Flows 2, 3, 4 use TCP Reno with ECN. Queuing at all routers if FIFO with RED enabled. The round trip times are:

   - 300 ms for flows 2 and 3,
   - 100 ms for flow 4.

   These numbers include all processing times. The MSS is the same for all flows and is equal to 1250 Bytes = 10^4 bits. We assume that the offered window is very large. Compute the rates of flows 2, 3 and 4.

5. In this question we assume the following modifications:
   - All flows use TCP.
   - B now acts as an application layer gateway for flow 1.
   - The connection S1 - B has the same RTT as flow 2.
   - Flows 3 and 4 have the same RTT, which is also the same as for the connection B - D1.

   What is the rate achieved by flow 1 (from S1 to D1 ?). What is the rate achieved by flow 2 ?

SCIPER: _____

First name: _____ Family name: _____

PROBLEM 4

1. The private network of the airport uses IP multicast to send a status information from a machine called Flight Info Server to displays and network monitors, as shown in the figure sheet. The Flight Info Server uses source specific multicast and sends status information to the IP address `ff35::6:7`. Unless otherwise specified, there are multiple routers in the network.

   (a) Display 1, Display 2 and Network Monitor 1 are all receiving the status information. Network Monitor 2 is started and is configured to also receive the status information. Say what happens at the network layer when Network Monitor 2 decides to receive the status information: say in particular what happens at the IP layer at the machines Display 1, Display 2, Network Monitor 1, Network Monitor 2 and Flight Info Server.

   (b) Assume, in this question only, that the network N is a bridged LAN. We observe the packets that carry status information at the Flight Info Server and at Display 1. Which MAC and IP addresses do we see in the packets ? Write your answer in the table below.

At Flight Info Server			
MAC source	MAC destination	IP source	IP destination

At Display 1			
MAC source	MAC destination	IP source	IP destination

(c) Assume now that N is a large network with many routers and subnetworks, and that Display 1 is not in the same subnetwork as the Flight Info Server. Is there any change to the answers of the previous question ?

(d) Assume in addition that all routers in N use BIER for multicast routing. Say which statement is true (only one answer is correct).
When a BIER router has one multicast packet with a BIER header to forward, it forwards...

- ☐ one single copy of the packet;
- ☐ one copy of the packet to every destination BIER router (i.e. if there are n destination BIER routers it sends n copies of the packet);
- ☐ one copy of the packet to every next-hop that is on the path of a destination BIER router (i.e. if there are m paths leading to some destination BIER routers it sends m copies of the packet).

The usual rules of quizzes apply: (i) if you cross the correct box you obtain the full score; (ii) if you cross zero or two boxes your score is 0 for this item; (iii) if you cross exactly one box and it is incorrect your score is negative, namely a penalty equal to half the score for this item.

2. Consider the network for Problem 4, Question 2 in the figure sheet. An IPv4 packet is received by LER A on port e with destination address 10.2.2.4. Which path does this packet follow ? We observe this packet on the link through which it arrives at the last MPLS router on its path. Which MPLS labels do we see ?

3. Our home network is connected to IPv4's Simpscom network (see figure sheet). Simpscom uses 6rd to provide IPv6 access. The parameters used by Simpscom are as follows. The prefix 2b00:1400::/28 is reserved for 6rd. The IPv6 prefix allocated to a customer network (such as our home network) is the concatenation of this prefix and the customers's IPv4 address provided by Simpscom. This makes a 60-bit prefix, which can be used freely by the home network. The address 85.2.3.4 is reserved by Simpscom to represent the IPv6 internet.

   (a) Among the following addresses, say which ones are possible choices for the IPv6 address of A (select all the valid answers).

address	valid	invalid
`2b00:1400:b080:7060::1`	☐	☐
`2b00:1400:b080:7069::1`	☐	☐
`2b00:1401:1080:7060::1`	☐	☐
`2b00:1401:1080:7069::1`	☐	☐

   The usual rules of quizzes apply: (i) if you cross the correct box in one row you obtain the full score for this row; (ii) if you cross zero or two boxes in one row your score is 0 for this row; (iii) if you cross exactly one box in a row and it is incorrect your score is negative for this row, namely a penalty equal to half the score for this row.

   (b) A downloads a file from B using TLS over TCP. We observe the packets flowing from B to A at observation points O1 and O2. In the IP header of the captured packets, what IP addresses (source and destination) and protocol types do we see ?

From B to A:			
At	IP source addr	IP dest addr	protocol
O1			
O2			

TCP IP EXAM - FIGURES

For your convenience, you can separate this sheet from the main document. Do not write your solution on this sheet, use only the main document. You do not need to return this sheet.

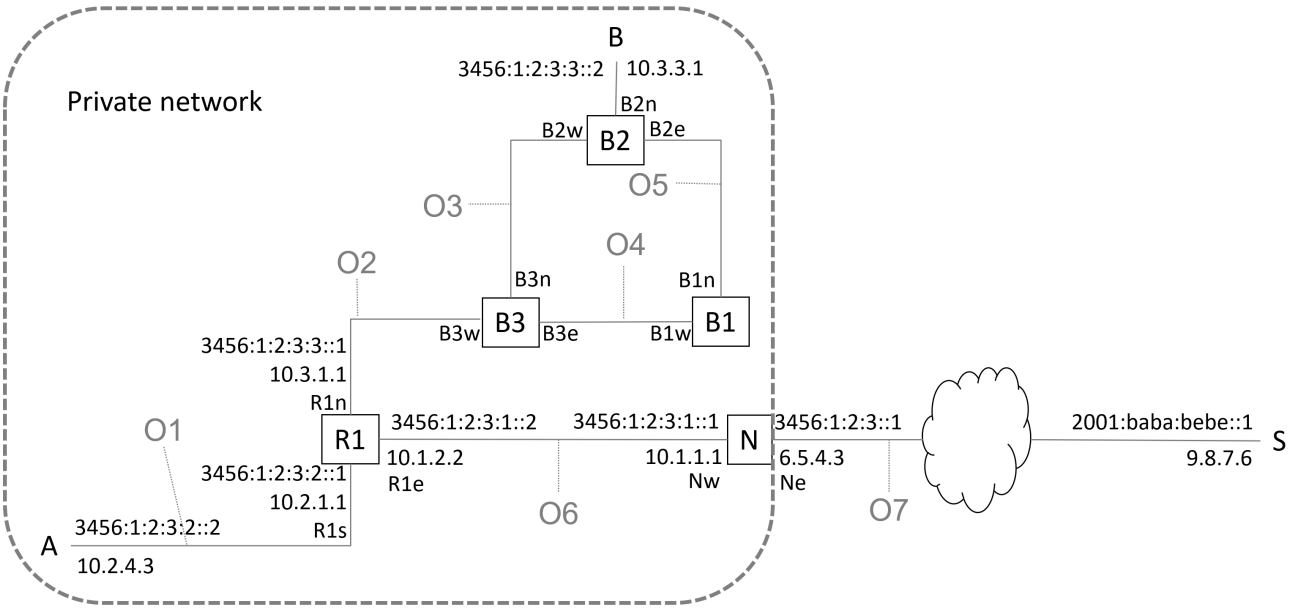

Problem 1

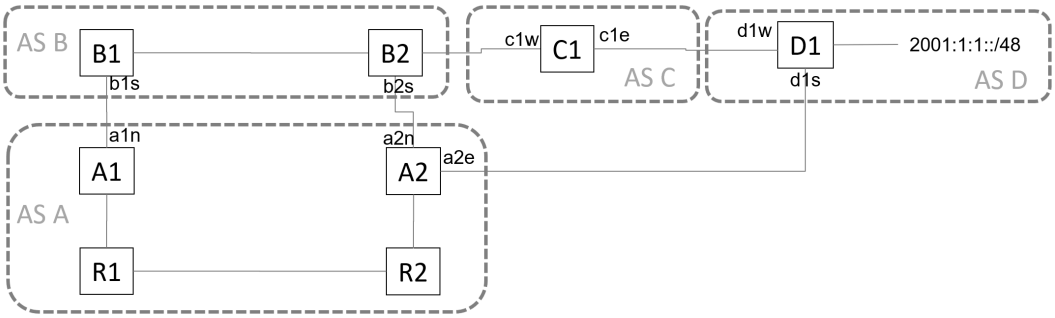

Problem 2

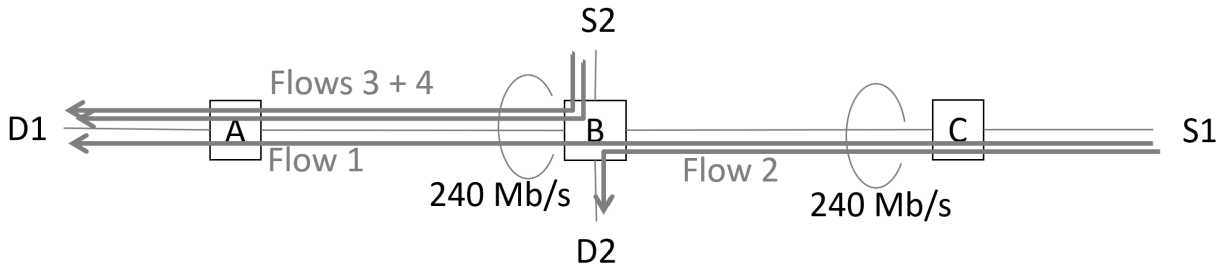

Problem 3

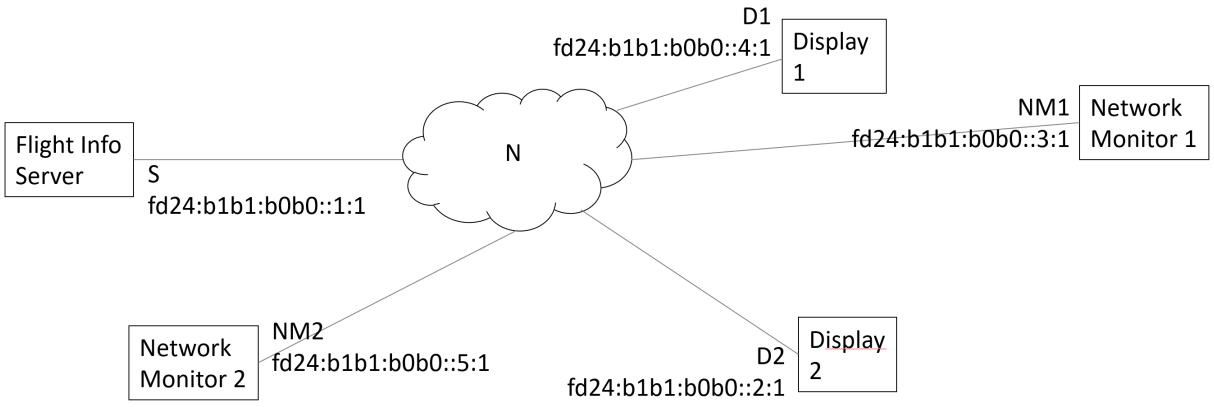

Problem 4, Question 1. S, NM1, NM2, D1 and D2 are MAC addresses.

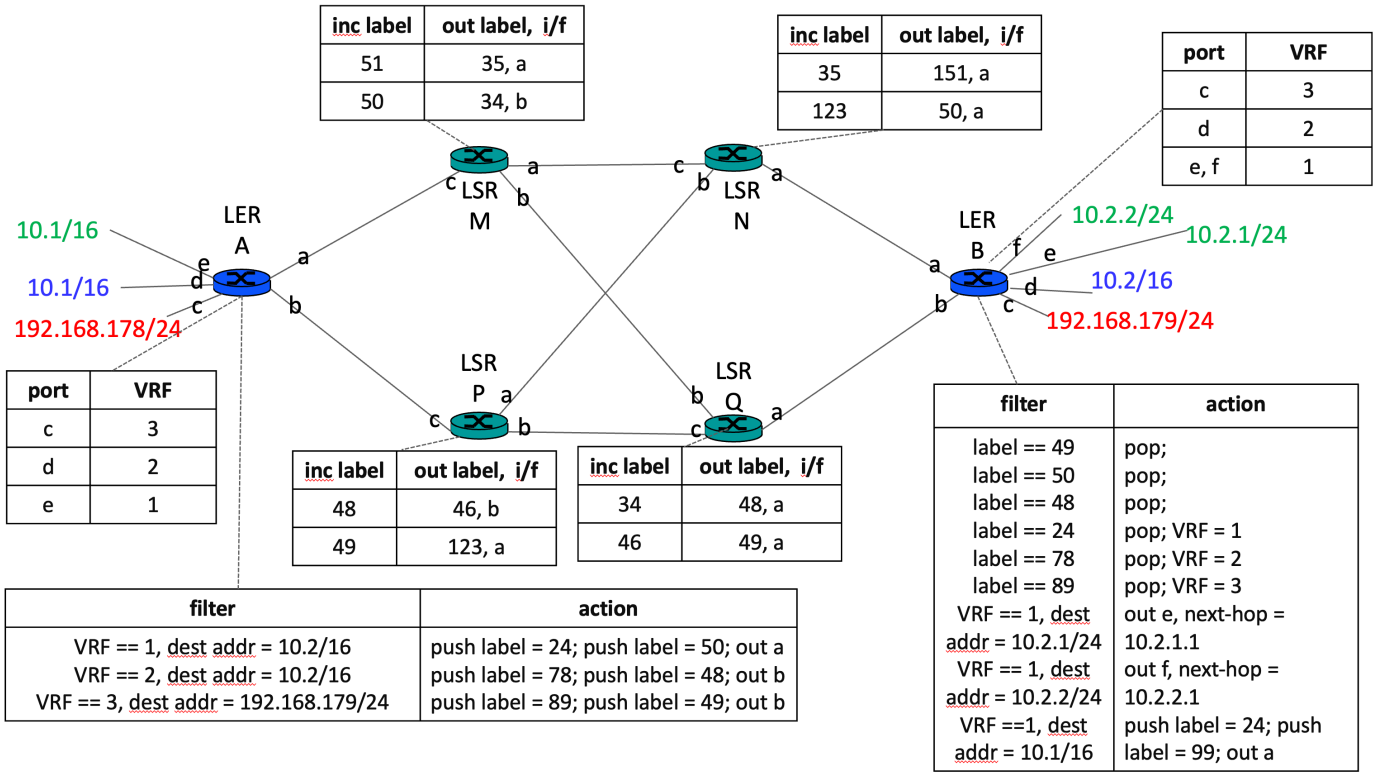

Problem 4, Question 2.

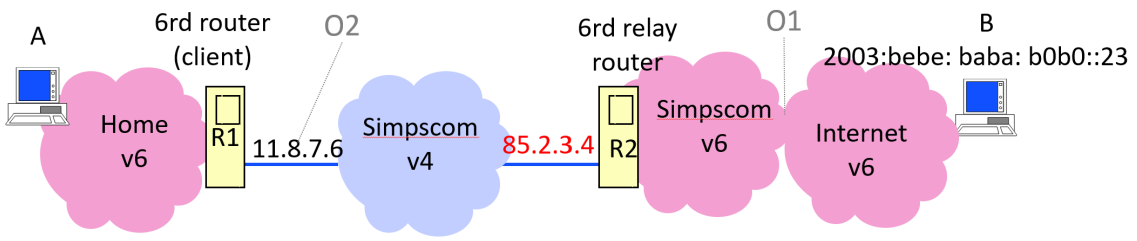

Problem 4, Question 3.