

SCIPER: \_\_\_\_\_

First name: \_\_\_\_\_ Family name: \_\_\_\_\_

---

EXAM  
TCP/IP NETWORKING  
Duration: 3 hours  
**With Solutions**

---

Jean-Yves Le Boudec

January 2019

## INSTRUCTIONS

1. 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.
2. All problems have the same weight.
3. You may need to make additional assumptions in order to solve some of the questions. If this happens, please describe such assumptions explicitly.
4. Figures are on a separate sheet, for your convenience.
5. No documents, no electronic equipments are allowed.
6. Justify every answer with a short explanation.

## PROBLEM 1

Consider the network for problem 1 in the figure sheet.  $A, B, C, D$  and  $E$  are hosts;  $BR1, BR2$  and  $BR3$  are bridges;  $N$  is an IPv4 NAT;  $R1, R2$  and  $R3$  are routers.  $O1$  to  $O6$  are observation points where we observe traffic in both directions of the link.

Hosts  $C$  and  $E$  are IPv6-only;  $D$  is IPv4-only;  $A$  and  $B$  are dual stack. All routers are dual-stack.

Some selected IP addresses are shown, as well as some selected MAC addresses (denoted with e.g.  $A, B, BR1w, \dots, R1e, \dots$ ). In some questions you might need to make assumptions about IP or MAC address values not shown in the figure.

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.

1. ISP6 delegates to Homer's network the prefix  $2001:a:a::/60$ . Inside Homer's network, all IPv6 netmasks are  $/64$ .

- (a) Give the uncompressed version of the address  $2001:a:a:1::2$ .

**Solution.**  $2001:000a:000a:0001:0000:0000:0000:0002$ .

- (b) Give a possible value of  $z$  (in the IPv6 address of  $R3$ 's east interface).

**Solution.**  $R3$ -east and  $B$  are in the same subnet therefore their network parts are identical, thus  $z = 2$ .

- (c) Among the following addresses, which ones are possible for host  $C$ ? Put an  $X$  in the correct boxes in the table below, with a short justification.

address	possible	not possible
$2001::3$		X
$2001:a:a:2::3$		X
$2001:a:a:5::3$	X	
$2001:a:a:f::3$	X	
$2001:a:a:11::3$		X

**Justification:** The address has to correspond to the delegated prefix  $2001:a:a::/60$ , i.e. should start with  $2001:000a:000a:000x:\dots$  where  $x$  is any digit. Since the subnet mask is 64 bits, the address network part should be precisely  $2001:000a:000a:000x::/64$  with  $x \in \{0, 1, 2, \dots, f\}$ . Furthermore, the subnet must be different for  $A, R1$ -east,  $B1$  and  $C$ , therefore  $x$  cannot be equal to 1, 2, 3 or 4.

- (d)  $A$  downloads a huge file from a web server at  $E$  using HTTP.  $A$  uses the local port 4567. At the same time,  $C$  also downloads a file from  $E$ , also using HTTP. By coincidence,  $C$  uses the same local port number, namely 4567. We observe the IP headers in the packets resulting from this transfer at  $O5$ , in the direction from  $E$  to  $A$  and from  $E$  to  $C$ . Give possible values of the protocol, the source and destination port numbers, the source and destination IP addresses and MAC addresses. Give the answers in the tables below.

At observation point $O5$ , from $E$ to $A$ and $C$ :						
MAC src.	MAC dest.	IP source	IP dest	prot	src. port	dest. port
$R1w$	$R3e$	$E$ 's IPv6 address	$A$ 's IPv6 address	tcp	e.g. 80	4567
$R1w$	$R2s$	$E$ 's IPv6 address	$C$ 's IPv6 address	tcp	e.g. 80	4567

- (e)  $A$  sends a ping message to  $B$ ,  $B$  sends a ping message to  $C$ , and  $C$  sends a ping message to  $A$ , all three using IPv6. At observation points  $O2, O3$  and  $O4$ , we observe the ping request

packets resulting from this activity. What are the MAC and IP source and destination addresses in such packets ? What is the Hop Count field, knowing that the HC value is equal to 64 in all IP packets generated by all hosts in this problem ? Put your answers in the table below, with a short justification.

Ping Requests seen at observation point O2:				
MAC source	MAC dest	IP source	IP dest	HC
B	R2s	B's IPv6 address	C's IPv6 address	64
R2s	R3e	C's IPv6 address	A's IPv6 address	63

Ping Requests seen at observation point O3:				
MAC source	MAC dest	IP source	IP dest	HC
R3e	B	A's IPv6 address	B's IPv6 address	63
B	R2s	B's IPv6 address	C's IPv6 address	64

Ping Requests seen at observation point O4:				
MAC source	MAC dest	IP source	IP dest	HC
This link is assumed to be disabled by the spanning tree protocol.				
Justification: The spanning tree protocol disables one of the links in the triangular topology of the bridges. The solution above is given assuming it is link O4 that is disabled.				

2. ISP4 allocates to  $N$  the address 11.10.9.8.

- (a) Give possible values of the network masks at  $N$ -west and at  $B$ .

**Solution.** They are both on the same subnet; thus they should have the same mask and the network part should be the same for both. One possible answer is 255.255.0.0 (i.e. 16 bits). Another possible answer is 255.255.240.0 (i.e. 20 bits).

- (b) Give a possible value for the complete IP address of  $A$ . Shortly justify your answer.

**Solution.** The address should be compatible with the answer to the previous question:  $A$ 's address should be in a different subnet than  $B$ . A possible answer is 10.1.1.7. Other possible answers are 10.x.y.7 with  $x \neq 9$  and  $x \in \{0, 1, \dots, 255\}$ . If your subnet mask at the previous question was 255.255.240.0 then any combination of  $x \in \{0, 1, \dots, 255\}$  and  $y \in \{16, 17, \dots, 255\}$  is possible.

- (c)  $A$  downloads a huge file from a web server at  $D$  using HTTP.  $A$  uses the local port 4567. At the same time,  $B$  also downloads a file from  $D$ , also using HTTP. By coincidence,  $B$  uses the same local port number, namely 4567. We observe the IP headers in the packets resulting from this transfer at  $O6$  and  $O1$ , in the direction from  $D$  to  $A$  and  $B$ . Give possible values of the protocol, the source and destination port numbers and the source and destination IP addresses. Give the answers in the tables below.

At observation point O6, from $D$ to $A$ and $B$ :				
IP source	IP dest	protocol	source port	dest. port
10.9.1.1	A's IPv4 address	tcp	80	4567
10.9.1.1	B's IPv4 address	tcp	80	4567

At observation point O1, from $D$ to $A$ and $B$ :				
IP source	IP dest	protocol	source port	dest. port
D's IPv4 address	11.10.9.8	tcp	80	e.g. 23456
D's IPv4 address	11.10.9.8	tcp	80	e.g. 23457



SCIPER: \_\_\_\_\_

First name: \_\_\_\_\_ Family name: \_\_\_\_\_

## PROBLEM 2

Consider the network for problem 2 in the figure sheet. There are three ASs,  $A$ ,  $B$  and  $C$  with routers  $A1, A2, A3, A4, B1, B2, B3, C1, C2, C3$  and  $C4$ . 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 for three links shown of the figure, where the cost is 10.

The figure shows two stub networks, at routers  $A3$  and  $A4$ , with their IPv6 address prefixes. The lower case symbols such as  $b1w$ ,  $b2se$  also represent IPv6 addresses.

Routers  $A1, A2, B1, B2, B3, C1$  and  $C2$  use BGP with their external neighbours and as required with their internal neighbours. Unless otherwise specified, the other routers do not use BGP. We assume that the BGP decision process use the following criteria in decreasing order of priority.

1. Shortest AS-PATH
2. Lowest MED, if taken seriously by this network
3. E-BGP > I-BGP
4. Shortest path to NEXT-HOP, according to IGP
5. Lowest BGP identifier

Furthermore, we assume that no optional BGP attribute (such as MED, LOCAL-PREF etc.) is used in any BGP message and that no aggregation is performed.

In the entire problem we assume that each AS redistributes internal OSPF destinations into BGP.

1. At time  $t_0$ , we assume that each AS is configured to redistribute E-BGP routes into OSPF. When redistributing E-BGP into OSPF, the OSPF cost of the redistributed route is set to the OSPF cost to the BGP next-hop plus 100.

Furthermore, the policy in  $A, B, C$  is such that all available routes are propagated to neighbouring ASs.

(a) At time  $t_1 > t_0$ , BGP and OSPF have converged in all ASs. At this time:

- i. Which routes are selected by BGP at  $B3$  and  $C1$  ? Give the answers in the tables below, with a short justification.

At B3	DESTINATION NETWORK	BGP NEXT-HOP	AS-PATH
	2001:1::/32	a2ne	A
	2001:2::/32	a2ne	A
At C1	DESTINATION NETWORK	BGP NEXT-HOP	AS-PATH
	2001:1::/32	a2e	A
	2001:2::/32	a2e	A

**Justification:**

- Due to redistribution from OSPF into BGP, A1 will announce (via E-BGP) to B1 about the two stub networks 2001:1::/32 and 2001:2::/32. Similarly, A2 will announce to B2 about the two stub networks. B1 and B2 will announce to B3 about the two stub networks, using I-BGP. The decision process at B3 will prefer the route obtained from B2, because the IGP distance to the NEXT-HOP is smaller.
- C1 will learn the two stub networks via either I-BGP updates from C2 or E-BGP updates from B3. According to the rule of shortest AS path, C1 will choose a2e of A2 as the BGP NEXT-HOP.

- ii. In the forwarding tables of B3 and C3, what are the entries for the destinations 2001:1/32 and 2001:2/32? Give the answers in the tables below, with a short justification.

At B3	DESTINATION NETWORK	NEXT-HOP	DISTANCE
	2001:1::/32	b2e	102
	2001:2::/32	b2e	102
At C3	DESTINATION NETWORK	NEXT-HOP	DISTANCE
	2001:1::/32	c1e,c4n	112
	2001:2::/32	c1e,c4n	112

**Justification:**

- Both of the stub networks are redistributed by B1 and B2, each with a distance 101. By Dijkstra's algorithm, the shortest path from B3 is via B2, and the distance is 101+1=102.
- Inside AS C, the two networks are redistributed only by C2, each with a distance 101 (C1's best routes are learned by I-BGP and therefore are not redistributed). At C3, there are two Equal-Cost choices for the NEXT-HOP: c1e and c4n. The distance at C3 is 101 (between a2e and c2w) plus 11 (aggregated between c2n and c3w, or between c2e and c3s), which is 112.

(b) At time  $t_2 > t_1$ , router  $A_2$  crashes. At time  $t_3 > t_2$ , BGP and OSPF have converged again in all ASs and  $A_2$  is not yet repaired. At this time:

i. Which routes are selected by BGP at  $B_3$  and  $C_1$ ? Give the answers in the tables below, with a short justification.

At B3	DESTINATION NETWORK	BGP NEXT-HOP	AS-PATH
	2001:1::/32	a1e	A
	2001:2::/32	a1e	A
At C1	DESTINATION NETWORK	BGP NEXT-HOP	AS-PATH
	2001:1::/32	b3e	A,B
	2001:2::/32	b3e	A,B

**Justification:**

- Since  $A_2$  crashes, the route it announced to  $B_2$  is withdrawn by  $B_2$ . As this was the best route at  $B_2$ ,  $B_2$  sends a withdraw update to  $B_3$ .  $B_3$  has now lost its best route, and picks as best the next available route, which is the one learned via  $B_1$ .
- A similar process will occur in AS C. Eventually,  $C_1$  picks as best route the one learned from  $B_3$ .

ii. In the forwarding tables of  $B_3$  and  $C_3$ , what are the entries for the destinations 2001:1/32 and 2001:2/32? Give the answers in the tables below, with a short justification.

At B3	DESTINATION NETWORK	NEXT-HOP	DISTANCE
	2001:1::/32	b2e	103
	2001:2::/32	b2e	103
At C3	DESTINATION NETWORK	NEXT-HOP	DISTANCE
	2001:1::/32	c1e	102
	2001:2::/32	c1e	102

**Justification:**

- In order to reach the BGP NEXT-HOP at  $A_1$ ,  $B_3$  needs to reach  $B_1$ . In AS B, the NEXT-HOP of  $B_3$  should be  $B_2$ , according to the cost. The total distance to the two stub networks in AS A is 101 (between a1e and b1w) plus 2 (aggregated between b1s and b3s), which is 103.
- In order to reach the stub networks,  $C_3$  needs to reach  $C_1$ . The total distance to the two stub networks is 101 (between b3e and c1n) plus 1 (between c1e and c3w), which is 102.

- (c) At time  $t_4 > t_3$  router  $A2$  is repaired. At time  $t_5 > t_4$ ,  $B3$  is compromised and now sends to  $C1$  the (bogus) BGP announcements

DESTINATION = 2001:1::/32, AS-PATH = B

DESTINATION = 2001:2:baba::/48, AS-PATH = B

Also assume that  $C1$  accepts these BGP announcements and that there are no other bogus announcements than these. At time  $t_6 > t_5$  BGP and OSPF have stabilized again in all ASs. At this time, how is an IPv6 packet routed from  $C3$  and  $C4$  to 2001:1:bebe:bebe::1 and to 2001:2:baba:baba::1? Give only the exit point out of AS  $C$  by putting an X in the table below, with a short justification.

From $C3$ to	Exit point out of AS $C$		
	c1n	c2nw	c2w
2001:1:bebe:bebe::1	X		
2001:2:baba:baba::1	X		

From $C4$ to	Exit point out of AS $C$		
	c1n	c2nw	c2w
2001:1:bebe:bebe::1			X
2001:2:baba:baba::1	X		

**Justification:**

1. The bogus route to 2001:1::/32 is accepted by  $C1$  as best route (under the condition that AS paths have the same length, E-BGP is preferred). The correct route via  $A2$  to 2001:1::/32 is selected as the best by  $C2$ , for the same reasons. From  $C3$ , the distance to  $b3e$  is less than that to  $a2e$ . Therefore, the packet to 2001:1:bebe:bebe::1 will go via  $C1$  (and then  $B3$ ). In contrast, from  $C4$ , the distance to  $a2e$  is less than that to  $b3e$ . Therefore, the packet to 2001:1:bebe:bebe::1 will go via  $C2$  (and then  $A2$ ).

2. The bogus route to 2001:2:baba::/48 is accepted by  $C1$  and by  $C2$ , as it is the only route to this prefix. Note that, 2001:2:baba::/48 is considered as a different destination than 2001:2::/32. By longest prefix match, the packets to 2001:2:baba:baba::1 match the bogus route, and in both cases go via  $C1$  (and then  $B3$ ).



2. We consider now a different scenario. At time  $t_0$ , we assume (as in the previous scenario) that each AS redistributes E-BGP routes into OSPF and that the OSPF cost of the redistributed route is set to the OSPF cost to the BGP next-hop plus 100. However, we now assume that the policy at  $B$  is such that routes imported from  $A$  are not propagated to  $C$ , and similarly, at  $C$ , routes imported from  $A$  are not propagated to  $B$ . In this scenario,  $B3$  does not send any bogus BGP announcement. We assume that router  $A2$  works normally until time  $t'_1 > t_0$ , then crashes and is not repaired. At time  $t'_2 > t'_1$ , BGP and OSPF have converged in all ASs. At this time: Which routes are selected by BGP at  $B3$  and  $C1$ ? Give the answers in the tables below, with a short justification.

At B3	DESTINATION NETWORK	BGP NEXT-HOP	AS-PATH
	2001:1::/32	ale	A
	2001:2::/32	ale	A
At C1	DESTINATION NETWORK	BGP NEXT-HOP	AS-PATH
	2001:1::/32	no route	no route
	2001:2::/32	no route	no route
<b>Justification:</b> 1. For AS B, the results are the same as in previous question 1-(b)-i.  2. C1 and C2 do not learn any route from AS B, due to the specified policy. Therefore, the two stub networks are unreachable.			

3. We consider yet another scenario. At time  $t_0$ , we assume that all routers, including  $A2$ , work normally. We also assume that the policy at  $B$  and  $C$  is the same as in the previous question. However, the configuration of AS  $C$  is modified: E-BGP is not redistributed into OSPF (but continues to be redistributed at  $A$  and  $B$ ). In this scenario,  $B3$  does not send any bogus BGP announcement. Explain what can be done in AS  $C$  to maintain full connectivity, in particular, we would like all routers in AS  $C$  to be able to forward packets to all destinations in  $A$  (propose only one solution). Note that, in this question, you may make changes to the assumptions that we put at the beginning of Problem 2.
- Solution.** We could configure  $C1$  or  $C2$  as the default gateway of  $C3, C4$ . Alternatively, and this is more robust, we could run BGP also in  $C3, C4$  (with injection and recursive lookup).



First name: \_\_\_\_\_ Family name: \_\_\_\_\_

**PROBLEM 3**

Consider the network for problem 3 on the figure sheet.

- The four boxes are routers. The capacity of the links between them is shown on the figure. The links are full duplex with same rate in both directions.
- There are 4 unidirectional flows, as shown on the figure. There is no other system and no other flow than shown on the figure. There is no other capacity constraint than the 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.

1. Assume the rates  $x_0, x_1, x_2, x_3$  of the four flows are allocated according to max-min fairness. Compute the values of  $x_0, x_1, x_2$  and  $x_3$ .

**Solution.** Apply the water-filling algorithm. Rates are in Gb/s.

Step 1: We give the same value to all rates and maximize it. We obtain the value of 1; the first link is saturated; it is used by flows 0 and 1 therefore  $x_0 = x_1 = 1$  and we freeze these variables to these values.

Step 2: We impose  $x_0 = x_1 = 1$ , give the same value to the other rates and maximize it. We obtain the value of 5; the second link is saturated; it is used by the unfrozen flows 2 and 3; therefore  $x_2 = x_3 = 5$ .

The max-min fair allocation is  $x_0 = 1, x_1 = 1, x_2 = 5, x_3 = 5$ .

TBD

2. Which of the following allocations, in Gb/s, are Pareto-efficient ? Justify your answer.

(a)  $x_0 = 1, x_1 = 1, x_2 = 5, x_3 = 5$

(b)  $x_0 = 1, x_1 = 1, x_2 = 3, x_3 = 5$

(c)  $x_0 = 1, x_1 = 1, x_2 = 4, x_3 = 6$

**Solution.** (a) is Pareto-efficient because it is max-min fair. We can also check directly: If we want to increase  $x_0$  we must decrease  $x_1$  because of the first link, which is saturated; and vice versa. If we want to increase  $x_2$  we must decrease  $x_1$  or  $x_3$  because of the second link, which is saturated; idem for  $x_3$ .

(b) is not Pareto-efficient because we can increase  $x_2$  without modifying the other rates (and obtain allocation (a)).

(c) is Pareto-efficient because every link is saturated, hence we cannot increase any single rate.

**Comment:** observe that (a) is Pareto-efficient although the third link is not saturated (but every flow goes through one link that is saturated).

3. Is the following allocation, in Gb/s, proportionally fair ? Justify your answer.

$$x_0 = 1, x_1 = 1, x_2 = 4, x_3 = 6$$

**Solution.** No.

**Solution 1.** Intuitively, flow 0 uses more resources and should get less than flow 1. Formally, consider the modified allocation, where we decrease  $x_0$  and increase  $x_1$  and  $x_3$ :

$$x'_0 = 1 - \delta, x'_1 = 1 + \delta, x'_2 = 4, x'_3 = 6 + \delta$$

For any  $\delta \in (0, 1]$ , this modified allocation is feasible. The sum of the rates of change is

$$\sum_i \frac{x'_i - x_i}{x_i} = \delta - \delta + \frac{\delta}{6} = \frac{\delta}{6} > 0$$

hence it is not proportionally fair (if it would be proportionally fair, the sum of the rates of change would be  $\leq 0$  for any feasible  $x'$ ).

**Solution 2.** From the previous question we know that this allocation is not Pareto-efficient. Since the proportional allocation must be Pareto-efficient, the answer is no.

**Solution 3.** Consider the two Pareto efficient allocations, (a) and (c), of the previous question, where (c) coincides with the given allocation in this question. Both of them satisfy the link capacity constraints. Then, we compute the sum of logarithms of the source rates: for (a) it is 1.3979 and for (c) it is 1.38. The proportionally fair allocation maximizes the sum of logarithms of the source rates, therefore, it cannot be the given one.

4. In this question flows 2 and 3 are shut down (i.e.  $x_2 = x_3 = 0$ ). Flows 0 and 1 use TCP Reno with ECN. Queuing at all routers is FIFO with RED enabled. The round trip times are:
- 300 ms for flow 0,
  - 100 ms for flow 1.

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 0 and 1.

**Solution.** The rates of the second and third links are much larger so there will be quasi no congestion indication coming from these links and we can ignore them, i.e. everything is as if flows 0 and 1 are sharing a single link of capacity 2Gb/s.

TCP provides a Pareto-efficient allocation therefore the sum of the two rates will be (approximately) 2Gb/s:

$$x_0 + x_1 = 2$$

By TCP Reno's loss throughput formula, the rate of each flow is inversely proportional to its RTT:

$$\frac{x_0}{\frac{1}{300}} = \frac{x_1}{\frac{1}{100}}$$

Combing the two gives, in Gb/s:

$$x_0 = 0.5, \quad x_1 = 1.5$$

5. Assume now that flows 0 and 1 use TCP Cubic with ECN instead of TCP Reno with ECN; the rest is as in the previous question. Can you guess how the allocation of rates would differ? We don't ask you to compute the rates; simply put an *X* in the correct boxes in the table below, with a short justification.

**Solution:**

the rate achieved with TCP Cubic and ECN is ...	flow 0	flow 1
higher	X	
same		
lower		X

**Justification:**

The rates depend on the bias of TCP against large RTTs. At large RTT and large throughputs, Cubic is less sensitive to RTT than Reno. We need to see if we are in the region of large RTT and large throughput, where Reno does differ from Cubic. We see this on the slide given in the exam booklet; here the RTTs of interest are 100 and 300 msec. We see on the slide that for such RTTs and for throughputs larger than  $\approx 50$  Mb/s we are in the region where Cubic differs. Therefore we expect the rates of flows 0 and 1 to be more similar than with TCP Reno, namely smaller for flow 1 and larger for flow 0.

First name: \_\_\_\_\_ Family name: \_\_\_\_\_

**PROBLEM 4**

- A, B and S are IPv4-only hosts. A and B communicate with S via IPv6 using 464XLAT, as shown on the figure. The address block 2001:a.b.c.d::/80 is allocated to the CLAT to represent hosts such as A and B. The address block 2001:baba::/96 is used by PLAT to represent the IPv4 internet. The block 200.0.0/24 is used by PLAT to represent hosts such as A and B.

A opens one TCP connection to S; the source port number used by A is 4444 and the destination port number is 443. B also opens one TCP connection to S; the source port number used by B is, by coincidence, also 4444, and the destination port number is also 443. We observe the IP and TCP packet headers for the packets resulting from this activity at observation points 1 and 2, in the direction from A and B to S. What are the protocol, IP source and destination addresses and port numbers ? Put your answer in the table below, with a short explanation of any assumption you are making.

<b>At observation point 1, from A and B to S:</b>				
IP source	IP dest	prot	src. port	dest. port
2001:a:b:c:d::a01:203	2001:baba::102:304	TCP	4444	443
2001:a:b:c:d::a01:204	2001:baba::102:304	TCP	4444	443
<b>At observation point 2, from A and B to S:</b>				
IP source	IP dest	prot	src. port	dest. port
200.0.0.10	1.2.3.4	TCP	5001	443
200.0.0.11	1.2.3.4	TCP	5002	443

**Justification:**  
 Here we assume that the PLAT translates source IP address 2001:a:b:c:d::a01:203 and source port 4444 to source IP address 200.0.0.10 and port number 5001. Similarly, it translates source IP 2001:a:b:c:d::a01:204 and source port 444 to source IP address 200.0.0.11 and port number 5002.

2.  $R1$  is a router and uses RIP, a distance-vector routing protocol that implements the distributed Bellman-Ford algorithm. All link costs are equal to 1. At time  $t_0$ , the routing table at  $R1$  contains

Destination Network	Next-Hop	Distance
9/8	2.2.2.2	2
9.9/16	2.1.1.1	3

The router  $R2$  is a neighbour of  $R1$  and has IP address 2.2.2.2. At  $t_1 > t_0$ ,  $R1$  receives from  $R2$  the distance-vector message

Destination = 9/8, distance = 5  
 Destination = 9.9/16, distance = 5

No other message is received between  $t_0$  and  $t_1$ . Just after processing this message, what is the state of the routing table at  $R1$ ? Give your solution by filling the empty cells in the table below, together with a short justification.

Destination Network	Next-Hop	Distance
9/8	2.2.2.2	6
9.9/16	2.1.1.1	3

Justification:

In distributed Bellman-Ford algorithm distances and next hops are updated as follows when node  $i$  receives an update from node  $j$ :

if  $j == \text{nexthop}(i)$

$q(i) = A(i, j) + q(j)$

else

$q(i) = \min(q(i), A(i, j) + q(j))$

end

if the 'else' condition above is satisfied and  $A(i, j) + q(j) > q(i)$ , then  $\text{nexthop}(i)$  is set to  $j$ .

Applying this algorithm gives the above table.



3. Homer uses a media server to stream music in his house. The audio server uses IP multicast, with source specific multicast. It uses the multicast address `ff35::1:2:3:4`. The music stream is unidirectional, from the media server to whomever is listening.

Lisa and Bart receive the music stream, each on their own separate device. Homer also would like to receive the audio stream. Say what happens at the IP layer when Homer decides to receive the audio stream. In particular, among the following machines: the media server, Lisa, Bart and Homer's machines, say which have to send specific IP layer control messages for Homer to be able to receive the audio stream. (Application layer messages are not considered here).

**Solution.** With IP multicast, it is the receiver that needs to join the group. Here, only Homer's device will send MLD messages. The media server does nothing, it is not aware at the IP layer of the fact that Homer joins.

4. Say what is true about IP fragmentation (put true/false in the cells below).

	Host may fragment	Router may fragment	Router may re-assemble
With IPv4	true	true	false
With IPv6	true	false	false

5. Sovkom's network uses OSPF with shortest paths. There are many equal-cost shortest path from several sources to destination. Could this have some negative side-effects ? If so, explain what could be done to avoid such side-effects.

**Solution.** All the equal cost paths (next hops) computed using Dijkstra's algorithm implemented in OSPF appear in the routing table of the router. The router then does per-flow load balancing to ensure that packets of the same flow (e.g. same TCP connection) are sent to the same next hop. This is done to improve TCP performance by avoiding reordering of packets at the destination. It is typically achieved using hash functions on flow identifiers (source and destination IP's and port numbers).

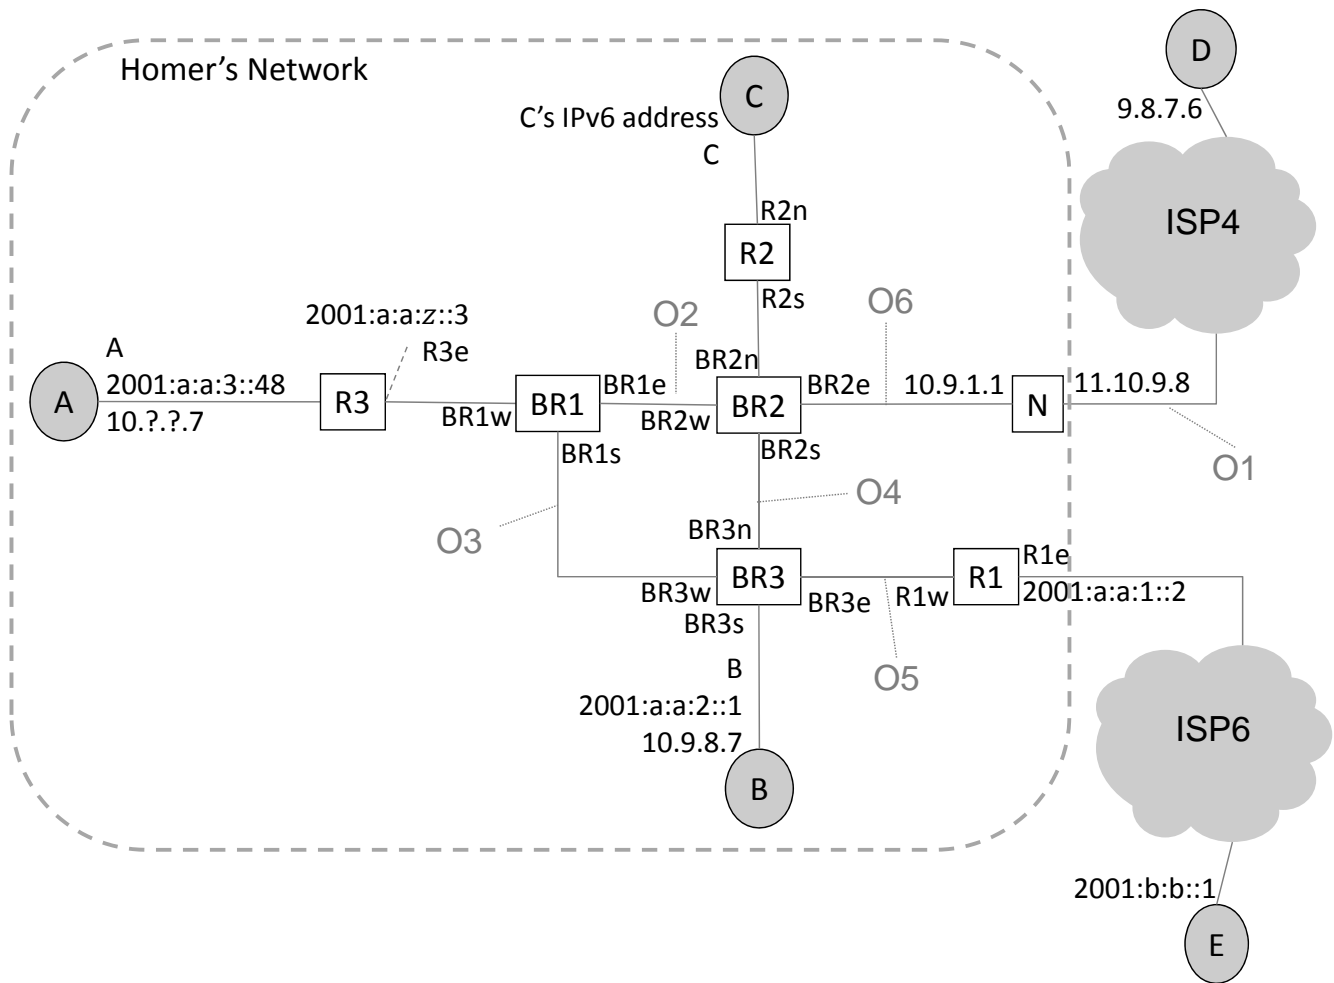


---

## TCP IP EXAM - FIGURES

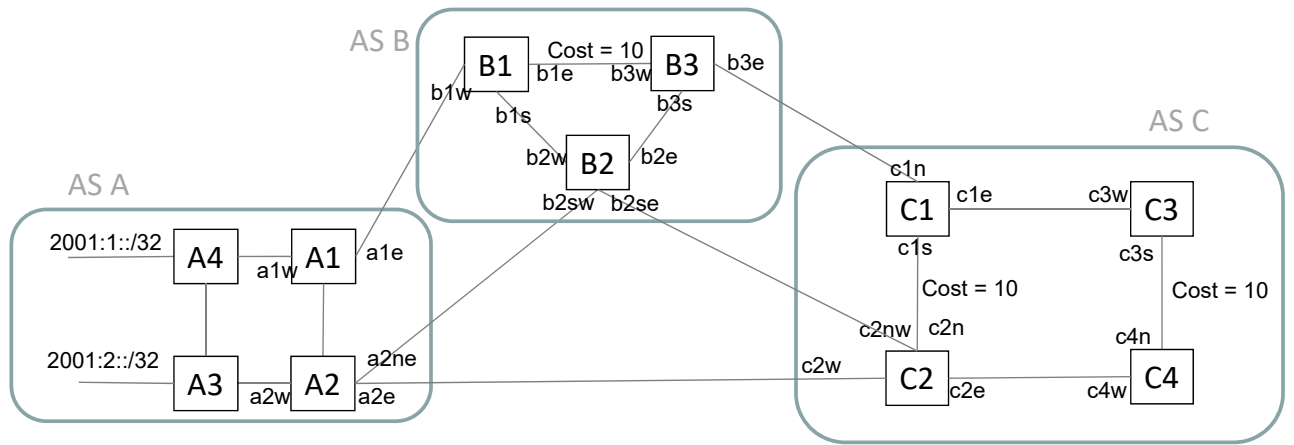
For your convenience, you can separate this sheet from the main document. Do not write your solution on this sheet, use only the main document. You do not need to return this sheet.

---




### Problem 1

---



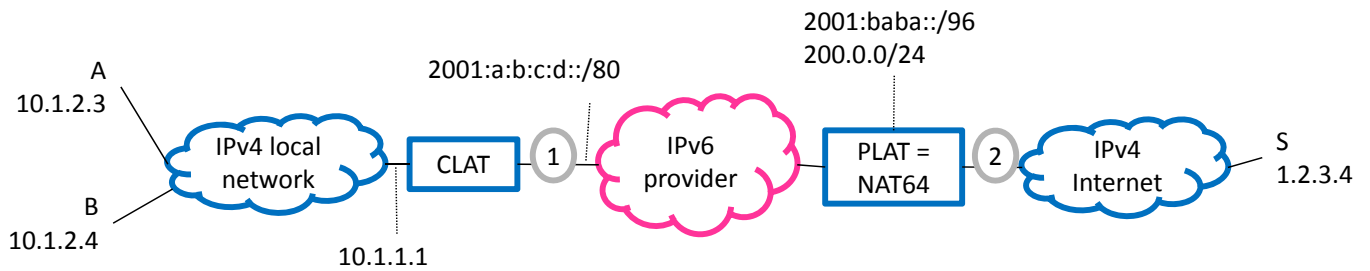
**Problem 2**

---



**Problem 3**

---



**Problem 4, Question 1.**

---