
EXAM
TCP/IP NETWORKING
Duration: 3 hours
With Solutions

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	<input type="checkbox"/>	<input checked="" type="checkbox"/>
255.255.0.0	<input checked="" type="checkbox"/>	<input type="checkbox"/>
255.255.240.0	<input checked="" type="checkbox"/>	<input type="checkbox"/>

The subnet mask should be such that B and R1n are in one subnet, A and R1s in another one, and R1e and Nw in yet another one.

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	<input type="checkbox"/>	<input checked="" type="checkbox"/>
68	<input type="checkbox"/>	<input checked="" type="checkbox"/>
80	<input checked="" type="checkbox"/>	<input type="checkbox"/>

B and R1n should be in one subnet, A and R1s in another one, R1e and Nw in another one and Ne in yet another one.

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).

B sends an ARP request in order to determine the MAC address of R1 and R1 replies to B. Then B sends the UDP packet to A.

at	srce MAC addr	dest MAC addr	type	srce IP addr	dest IP addr	TTL
O1	R1s	A	IPv4	10.3.3.1	10.2.4.3	63
O2	B	ff-ff-ff-ff-ff-ff	ARP	N/A	N/A	N/A
O2	R1n	B	ARP	N/A	N/A	N/A
O2	B	R1n	IPv4	10.3.3.1	10.2.4.3	64

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

The bridges setup a tree by electing a root and tracing all shortest paths to the root. Assume the elected root is B1. The link B2-B3 is de-activated and the tree is B2-B1-B3. The first ARP packet is broadcast and is sent to all bridges along the tree, so it is visible at O4 and O5. The ARP reply is unicast from R1 to B and is visible at O4 and O5 as well. Same for the UDP packet.

If the root is B2, the link B1-B3 is deactivated. The first ARP packet is visible at O3 and O5, the second ARP packet and the UDP packet are visible at O3 only since bridges are assumed to be in steady state hence have learnt where to forward the MAC addresses of B and R1.

If the root is B3, the link B1-B2 is deactivated. The first ARP packet is visible at O3 and O4, the second ARP packet and the UDP packet are visible at O3 only.

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.

Since N is a NAT for IPv4, the IP addresses of A and B are mapped to the public IP address of N. The port numbers used by A and B must also be mapped to different values.

From S to A					
At	srce IP addr	dest IP addr	protocol	srce port	dest port
O7	9.8.7.6	6.5.4.3	tcp	443	e.g. 7777
O6	9.8.7.6	10.2.4.3	tcp	443	5432

From S to B					
At	srce IP addr	dest IP addr	protocol	srce port	dest port
O7	9.8.7.6	6.5.4.3	tcp	443	e.g. 7778
O6	9.8.7.6	10.3.3.1	tcp	443	5432

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 ?

The packet is re transmitted by the TCP source, i.e. S.

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.

Since N is a router for IPv6, the IP addresses and port numbers are the same as at the source.

From S to A					
At	srce IP addr	dest IP addr	protocol	srce port	dest port
O7	2001:baba:bebe::1	3456:1:2:3:2::2	udp	443	9876
O6	2001:baba:bebe::1	3456:1:2:3:2::2	udp	443	9876

From S to B					
At	srce IP addr	dest IP addr	protocol	srce port	dest port
O7	2001:baba:bebe::1	3456:1:2:3:3::2	udp	443	9876
O6	2001:baba:bebe::1	3456:1:2:3:3::2	udp	443	9876

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 ?

The data is re transmitted by the QUIC source, i.e. by S.

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 ?

There are two TCP connections, one between A and N and one between N and S. The missing data is re-transmitted by the source of the connection between A and N, i.e by N.

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	A2	2001:1:1::/48	d1s	D
A2	D1	2001:1:1::/48	d1s	D
R1	A2	2001:1:1::/48	d1s	D
R2	A2	2001:1:1::/48	d1s	D

Justification: Because OSPF-internal is redistributed into BGP, D1 announces 2001:1:1::/48, next-hop = d1s, AS-PATH=D to A2 and 2001:1:1::/48, next-hop = d1w, AS-PATH=D to C1. Eventually, all routers in AS A may receive routes to 2001:1:1::/48 directly from D or from B. The latter routes, if any, have larger AS-PATH and are not preferred. Therefore, A2 prefers the route 2001:1:1::/48, next-hop = d1s, AS-PATH=D received from D1 and A1, R1 and R2 prefer the route 2001:1:1::/48, next-hop = d1s, AS-PATH=D received from A2.

- (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	A1	2001:1:1::/48	a1n	A D

Justification: B1 may receive routes from A1 (by E-BGP, with AS-PATH = A D) or from B2 (by I-BGP, with AS-PATH = A D or AS-PATH = C D, depending on the choice made by B2). They all have same AS-PATH length and E-BGP is preferred, so B1 prefers the route received from A1.

- (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 ?
A2	2001:1:1::/48	a2n	A D	yes
B1	2001:1:1::/48	a1n	A D	no
C1	2001:1:1::/48	c1w	C D	no

Justification: The route received from B1 is not selected because it has same AS-PATH length but comes from I-BGP. B2 can now choose between 2001:1:1::/48, next-hop = c1w, AS-PATH= C D received from C1 and 2001:1:1::/48, next-hop = a2n, AS-PATH=A D received from A2. They have same AS-PATH length and are both from E-BGP. Since the identifier A2 is smaller than C1, the route received from A2 is preferred.

- (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) ? **Both will reach the destination because all routers on the path have injected their BGP NEXT-HOP into the forwarding tables and recursive lookup is enabled.**

Packet from B1: B1 - A1 - R1 - R2 - A2 - D1

Packet from B2: B2 - A2 - D1

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.

In questions 2 and 3, there are two possible correct answers, depending on the order the BGP announcements were sent by the routers.

Solution 1:

At	From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH
A1	B1	2001:1:1::/48	b1s	B C D
A2	B2	2001:1:1::/48	b2s	B C D

Justification: Assume that A2 receives a route from D1 with PATH=D and B2 receives a route from C1 with PATH=C D. Both these routes are the best ones for the two routers so they need to be announced. If B2 announces its route to A2 first, then A2 will promote this route to best one, due to higher LOCAL-PREF. A1 will also receive the route from B1 and promote it to best for the same reason. A2 will not announce a best route to B2, because its best route includes AS B in the path, so that would create a loop. Therefore, the best route at B2 remains the one with PATH=C D.

Solution 2:

At	From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH
A1	A2	2001:1:1::/48	d1s	D
A2	D1	2001:1:1::/48	d1s	D

Justification: Assume now that A2 announces its best route to B2 first, then B2 will promote this route to best one, due to a smaller BGP Identifier $A2 < C1$ (the rest of the rules are the same). B2 will not announce this route back to A2, because that would create a loop, so the best route at A2 remains the one received from D1. A2 will propagate this route to A1, which promotes it to best one, as it does not receive any route from B1.

- (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).

Solution 1:

At R1 :				
From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH	Best route ?
A1	2001:1:1::/48	b1s	B C D	yes
A2	2001:1:1::/48	b2s	B C D	no

Justification: R1 peers with all BGP routers inside AS A. R1 receives from A2 only its best route, i.e. $2001:1:1::/48$, next-hop = b2s, AS-PATH= B C D. From A1, R1 receives the route $2001:1:1::/48$, next-hop = b1s, AS-PATH= B C D. Both routes have same AS path length, both come from I-BGP; the route learnt from A1 has shorter distance to next-hop and is preferred.

Solution 2:

At R1 :				
From BGP Peer	Destination Network	BGP NEXT-HOP	AS-PATH	Best route ?
A2	2001:1:1::/48	d1s	D	yes
Justification: In this case, R1 receives only the best route from A2. It does not receive any route from A1, because its best route was learnt from I-BGP.				

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 received 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) ?

Solution 1:

By OSPF, R2 learns that the prefix $2001:1:1::/48$ is reachable at A1 and at A2. OSPF also computes the shortest paths to both and finds that the final distances to $2001:1:1::/48$ are both equal to 102. Since ECMP is supported, packets go via both possible paths: some packets will go via R2 - A2 - B2 - C1 - D1 and others via R2 - R1 - A1 - B1 - B2 - C1 - D1.

Solution 2:

In the second case, there is only one path to reach the destination network, so all the packets will follow the path R2-A2-D1.

(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 OSPF 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) ?

Solution 1:

R1 injects the bogus prefix $2001:1::/32$ into OSPF. Since routers do not aggregate, both prefixes $2001:1:1::/48$ and $2001:1::/32$ are propagated into OSPF and BGP. However, the former is shorter and longest prefix match is used in forwarding tables. The considered packets match both destinations but only the more specific one is used; therefore, the packets are not affected and will reach D1, over the same paths as before.

Solution 2:

In this case, there is no path that leads to $2001:1:1::/48$ through R1, so again using longest prefix match, all packets will reach the destination over the same path as before.

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.

- $x_1 = 160, x_2 = 80, x_3 = x_4 = 40$
- $x_1 = x_2 = x_3 = x_4 = 50$
- $x_1 = x_2 = x_3 = x_4 = 80$
- $x_1 = 80, x_2 = 160, x_3 = 80, x_4 = 80$

Solution. The Pareto-efficient allocations are (a) and (d). For these two allocations, the link constraints are satisfied with equality so it is not possible to unilaterally increase any rate.

The allocations (b) and (c) are not Pareto-efficient. For (b), any rate can be increased unilaterally. For (c), x_2 can be increased without the need of decreasing any of x_1, x_3, x_4 .

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 ?

Solution. There is a unique max-min fair allocation and it is Pareto efficient. It can be computed via water-filling. At step 1 we assign $x_1 = x_2 = x_3 = x_4 = 80$ Mb/s. The link between A and B saturates and thus x_1, x_3, x_4 freeze at 80 Mb/s. At step 2 x_2 increases to 160 Mb/s and then the link between B and C saturates. Thus, the max-min fair allocation is $x_1 = x_3 = x_4 = 80$ Mb/s and $x_2 = 160$ Mb/s.

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 ?

Solution. There is a unique proportionally fair allocation and it is Pareto efficient. The allocation is obtained by solving the following optimization problem:

$$\max \log x_1 + \log x_2 + \log x_3 + \log x_4$$

subject to:

$$x_1 + x_2 \leq 240$$

$$x_1 + x_3 + x_4 \leq 240$$

$$x_i \geq 0, i = 1, 2, 3, 4$$

The first constraint should be satisfied with equality because otherwise we can increase x_2 until we obtain equality. Similarly, the second constraint should be satisfied with equality because otherwise we can increase x_3 or x_4 until we obtain equality.

Therefore,

$$x_2 = 240 - x_1$$

$$x_3 + x_4 = 240 - x_1$$

Note that x_3 and x_4 have exactly the same constraints and we know that there is a unique proportionally fair allocation. Therefore we must have $x_3 = x_4$, because if $x_3 \neq x_4$ we can exchange x_3 and x_4 and obtain another proportionally fair allocation, which is not possible.

Thus we must have

$$x_3 = x_4 = \frac{240 - x_1}{2}$$

Thus, the optimization problem can be transformed as

$$\max \log x_1 + \log(240 - x_1) + 2 \log \frac{240 - x_1}{2}$$

$$x_1 \geq 0$$

$$x_1 \leq 240$$

This is a problem in one single variable, so all we need is to study the function $f : x_1 \mapsto \log x_1 + \log(240 - x_1) + 2 \log \frac{240 - x_1}{2}$ over $x_1 \in [0; 240]$. The derivative of f is

$$f'(x_1) = \frac{1}{x_1} - \frac{1}{240 - x_1} - 2 \frac{1}{240 - x_1} = \frac{1}{x_1} - 3 \frac{1}{240 - x_1} = \frac{240 - x_1 - 3x_1}{x_1(240 - x_1)} = \frac{240 - 4x_1}{x_1(240 - x_1)}$$

x_1	0	60	240	
$f'(x_1)$		+	0	-
$f(x_1)$			↗	↘

Thus the maximum of f is reached at $x_1 = 60$; it follows that the proportionally fair allocation is $x_1 = 60$, $x_2 = 180$ Mb/s and $x_3 = x_4 = 90$ Mb/s.

We can observe that flow 1 obtains less than in the max-min fair allocation, as expected as flow 1 uses two resources.

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 is 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.

Solution. The UDP source does not adapt its rate. TCP flows adapt to use the leftover capacity at every link, which is 180 Mb/s. Flow 2 does not compete with another TCP source and since TCP allocations are Pareto efficient, it must fully utilize the first link. Thus, $x_2 = 180$ Mb/s.

By using the loss-throughput formula for TCP Reno, we obtain

$$x_3 = \frac{MSS \cdot 1.22}{0.300 \sqrt{q}}$$

and

$$x_4 = \frac{MSS \cdot 1.22}{0.100 \sqrt{q}}$$

where q is the percentage of the ECN-marked packets, which we can assume is the same for both x_3 and x_4 . Therefore, $x_4 = 3x_3$. Since TCP Reno allocations are Pareto-efficient, $x_3 + x_4 = 140$ and as a result, $x_4 = 105$ Mb/s and $x_3 = 35$ Mb/s. To conclude, $x_2 = 180$ Mb/s, $x_3 = 35$ Mb/s and $x_4 = 105$ Mb/s.

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 ?

Flow 1 now consists of two piped connections: let us call them flow 1a (S1 - B) and flow 1b (B- D1). Flows 1b, 3 and 4 compete for 240 Mb/s; they have the same RTT and MSS, thus if there is no other constraint they obtain the same rate, namely 80 Mb/s.

Flows 1a and flow 2 also compete for 240 Mb/s and have the same RTT and MSS, thus if there is no other constraint they obtain the same rate, namely 120 Mb/s.

However flow 1a and flow 1b are piped; if flow 1a has a larger rate than flow 1b, the buffer at B increases and when it becomes full, B cannot consume data arriving on the TCP connection 1a. The window flow control of TCP ensures that S1 does not send faster than B can consume. Over the long run, flow 1a is limited by the rate limit of flow 1b, namely $x_{1a} = 80$ Mb/s. The end-to-end rate of flow 1 is therefore $x_1 = 80$ Mb/s. Flow 2 will increase its rate to obtain Pareto efficiency, i.e. $x_2 = 160$ Mb/s.

PROBLEM 4

- 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.
 - 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. **With IP multicast, it is the receiver that needs to join the group. Network Monitor 2 sends an MLD message to the group of all on-link multicast routers. Flight Info Server does nothing, it is not aware at the IP layer of the addition of a receiver to the multicast group. Nothing is received by Display 1, Display 2 and Network Monitor 1.**
 - 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
S	33-33-00-06-00-07	fd24:b1b1:b0b0::1:1	ff35::6:7

At Display 1			
MAC source	MAC destination	IP source	IP destination
S	33-33-00-06-00-07	fd24:b1b1:b0b0::1:1	ff35::6:7

The addresses are the same at both observation points. The destination MAC address is derived from the destination IP address.

- 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 ?
At the Flight Info Server there is no change. At Display 1, the only change is that the source MAC address is replaced by the MAC address of the last router on the path.
- 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 ?

The packet is associated with VRF 1 and receives inner label 24 and outer label 50. The label switching tables at the LSRs indicate that the packet travels via A - M - Q - B. At B the outer label is removed and the remaining label, namely 24, indicates VRF 1. The virtual routing table indicates that the outgoing port is f, therefore the packet exits the network on port f.

The last MPLS router is B and it arrives on port b. The MPLS labels are: outer label 48 and inner label 24.

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 customer'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	<input checked="" type="checkbox"/>	<input type="checkbox"/>
2b00:1400:b080:7069::1	<input checked="" type="checkbox"/>	<input type="checkbox"/>
2b00:1401:1080:7060::1	<input type="checkbox"/>	<input checked="" type="checkbox"/>
2b00:1401:1080:7069::1	<input type="checkbox"/>	<input checked="" type="checkbox"/>

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.

The address prefix allocated by Simpscom is obtained by concatenating the 28 bits of the reserved prefix with the 32 bits of our IPv4 address. In hexa, our IPv4 address is 0b08:0706 therefore the allocated prefix is 2b00:1400:b080:7060::/60. Depending on how we organize subnet prefixes in our home network, any address of the form

2b00:1400:b080:706x:xxxx:xxxx:xxxx:xxxx

is valid, where x stands for any hexa digit.


- (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	2003:bebe:baba:b0b0::23	A's address, for example 2b00:1400:b080:7060::1	TCP
O2	85.2.3.4	11.8.7.6	IPv6


At O1 we see an IPv6 packet that contains a TCP payload. At O2 we see an IPv4 packet that contains an IPv6 packet.

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.