

SCIPER: _____

First name: _____ Family name: _____

EXAM
TCP/IP NETWORKING
Duration: 3 hours

Jean-Yves Le Boudec

January 2020

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, except for a non-connected calculator.

- (c) B sends a sequence of IP packets to D . We assume all routing and bridging protocols have converged and all forwarding tables have been learnt. At which of the observation points $O2, O3, O4$ are the packets visible ? Justify your answer.

- (d) C downloads a huge file from a web server at E using HTTP over QUIC. C uses the local port 4567. The server port number for QUIC is 443. At the same time, B also downloads a file from E using HTTP over QUIC. By coincidence, B uses the same local port number, namely 4567. We observe the packet headers in the packets resulting from this transfer at $O5$ and $O6$, in the direction to E . 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 $O5$, towards E :				
IP source	IP dest	protocol	source port	dest. port

At observation point $O6$, towards E :				
IP source	IP dest	protocol	source port	dest. port

2. ISP6 delegates the prefix $2001:a:a:a::/76$ to the network of Problem 1 . Inside the network, all IPv6 subnet prefixes are $/80$.

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

- (b) Among the following addresses, which ones are possible for host B ? Put an X in the correct

boxes in the table below, with a short justification.

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

Justification:

- (c) *B* downloads a huge file from a web server at *A* using HTTP over TLS over TCP. *B* uses the local port 4567 and the server port at *A* is 443. At the same time, *D* also downloads a file from *A*, also using HTTP over TLS over TCP. By coincidence, *D* uses the same local port number, namely 4567. We observe the IP headers in the packets resulting from this transfer at *O1*, in the direction towards *A*. 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 <i>O1</i> , towards <i>A</i> :				
IP source	IP dest	prot	src. port	dest. port

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 $A1$, $A2$, $B1$, $B2$, $B3$, $B4$, $R1$, $R2$, $C1$ and $C2$. 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 shown on the figure.

The figure shows stub networks, at routers $A1$, $C2$, $R1$ and $R2$, with their IPv6 address prefixes. The lower case symbols such as $b1w$, $a1e$ also represent IPv6 addresses.

Routers $A1$, $A2$, $B1$, $B2$, $B3$, $B4$, $C1$ and $C2$ 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. Shortest AS-PATH
2. E-BGP is preferred over I-BGP
3. Shortest path to NEXT-HOP, according to IGP
4. Lowest BGP identifier is preferred (e.g. $A1$ is preferred over $A2$)

Furthermore, we assume that:

- No optional BGP attribute (such as MED, LOCAL-PREF etc.) is used in any BGP message.
- No aggregation of route prefixes is performed by BGP.
- The policy in A , B , C 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.
- Every router performs recursive forwarding-table lookup.
- Equal Cost Multi-Path routing is supported by all routers.

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 list of BGP routes received by $A2$ with destination = $2001:1::/32$? Which route is selected as best route by $A2$? Give your answer in the table below, with a short justification (put as many rows as necessary).

At $A2$:				
From BGP Peer	Destination Network	BGP Next-Hop	AS-Path	Best route ?
	2001:1::/32			
	2001:1::/32			
	...			
Justification:				

- (b) Still at time t_1 , what is the list of BGP routes received by $A2$ with destination = $4001:1::/32$? Which route is selected as best route by $A2$? Give your answer in the table below, with a short justification (put as many rows as necessary).

At $A2$:				
From BGP Peer	Destination Network	BGP Next-Hop	AS-Path	Best route ?
	4001:1::/32			
	4001:1::/32			
	...			
Justification:				

- (c) Still at time t_1 , what is the list of BGP routes received by $R2$ with destination = $4001:1::/32$? Which route is selected as best route by $R2$? Give your answer in the table below, with a short justification (put as many rows as necessary).

At $R2$:				
From BGP Peer	Destination Network	BGP Next-Hop	AS-Path	Best route ?
	4001:1::/32			
	4001:1::/32			
	...			
Justification:				

- (d) Still at time t_1 , $R2$ has a packet to forward with destination address $4001:1:2:3::1$. Which path will this packet take inside AS B ? At which routers along this path is recursive table lookup required for forwarding this packet ?

- (e) At time $t_2 > t_1$ the link $A1 - C1$ breaks. At time $t_3 > t_2$, the BGP protocol has converged again. what is the list of *valid* BGP routes received by $A2$ with destination = $4001:1::/32$? Which route is selected as best route by $A2$? Give your answer in the table below, with a short justification (put as many rows as necessary).

At $A2$:				
From BGP Peer	Destination Network	BGP Next-Hop	AS-Path	Best route ?
	4001:1::/32			
	4001:1::/32			
	...			
Justification:				

- (f) At time $t_3 > t_2$, the link $A1 - C1$ is repaired and all routing protocols have converged again. Router $C1$ is compromised and sends the bogus route

dest = $3001:1:2::/48$, as-path = C, next-hop = c1e

to $B4$. No other bogus message is sent by any other router. The bogus message is accepted by $B4$. At time $t_4 > t_3$, BGP has converged again and $B2$ has a packet to forward with destination address $3001:1:2:3::1$. By which link will this packet exit AS B ?

2. In this question we assume that the network is restarted, with the following changes in the configurations of $R1$:

- $R1$ does not run BGP (but continues to run OSPF).
- In addition, $R1$ is configured with a static default route to $B2$. This route co-exists in the forwarding table with the routes obtained from OSPF.

We also assume that there is no more bogus announcement. Recall that there is no redistribution of BGP into OSPF (but there *is* redistribution of internal OSPF destinations into BGP).

At time t_5 the routing protocols have converged.

$R1$ has a packet to forward with destination address $4001:1:1:3:3::1$. Which path will this packet take inside AS B ? At which routers along this path is recursive table lookup required for forwarding this packet?

3. In this question we assume that the network is restarted, with the following changes in the configurations of routers inside AS B :
- $R1$ and $R2$ do not run BGP (but continue to run OSPF).
 - There is no static default route in $R1$.
 - Routers $B1, B2, B3$ and $B4$ announce in OSPF the destination $::/0$ with cost = 5. This is treated by OSPF as an internal destination.

We also assume that there is no more bogus announcement. Recall that there is no redistribution of BGP into OSPF (but there *is* redistribution of internal OSPF destinations into BGP).

At time t_6 the routing protocols have converged. $R1$ has a large number of packets to forward with destination addresses $4001:1:x$ where x is a string of 96 bits. The value of x is different for every packet. Which way will these packets travel inside AS B ?

SCIPER: _____

First name: _____ Family name: _____

PROBLEM 3

Consider the network for problem 3 on the figure sheet.

- Hosts A , B and C are downloading content from server S . $R1$, $R2$ and $R3$ are routers, unless otherwise specified.
- The link rates are indicated on the figure. All links are full duplex with same rate in both directions.
- There is no other system than shown on the figure, and we neglect all flows other than between A , B , C and S . 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.
- The round trip times (RTTs) shown on the figure are between S and respectively A , B , C (for example, the RTT between S and B is 15 ms. The RTTs include all processing times and all queuing delays at routers.
- We neglect all overheads and assume that the link capacities can be fully utilized at bottlenecks.
- The MSS is the same for all flows and is equal to 1250 Bytes = 10^4 bits.
- We call x , y , z the rates of flows $S \rightarrow A$, $S \rightarrow B$, $S \rightarrow C$ in Mb/s.

1. What allocations (x, y, z) are Pareto-efficient ?

2. Assume that some bandwidth manager is used, which allocates rates to flows according to max-min fairness. What are the values of x, y, z ?

3. Same question with proportional fairness instead of max-min fairness.

4. We now assume that the three flows are using TCP RENO with ECN. What is the value of the rate of each flow ?

5. We continue to assume that the flows are using TCP with ECN. We observe the IP headers of packets on the link from $R1$ to A . Which proportion of packets do we see marked as “Congestion Experienced” ?

6. Assume now that $R3$ is an application layer gateway instead of a router, namely, the flow sent by S to C is relayed at the application level by $R3$ to C . We assume that all flows are using TCP with ECN. The round trip time from S to $R3$ is 15 msec; from $R3$ to C it is 80 msec; the other round trip times are unchanged. We assume that the application layer gateway has infinite processing power and storage. What is the value of the rate at which C receives data from S via $R3$?

SCIPER: _____

First name: _____ Family name: _____

PROBLEM 4

1. In a smart grid, a sensor *S* sends measurements every 20 msec to two data concentrators *D1* and *D2* (see figure). The sensor uses IP multicast, with source specific multicast. It uses the multicast address $ff35::4:3:2:1$. The data stream is unidirectional, from the sensor. The network is a single bridged LAN.
 - (a) We do a packet capture at the networking interface of the sensor *S* and at networking interface of *D1*. We observe only the packets that carry the unidirectional data stream sent by *S*. Which addresses do we see in the packets sent by *S*? Put the answer in the table below (see on the figure sheet for device addresses).

At sensor <i>S</i>:			
MAC source	MAC dest	IP source	IP dest

At Data Concentrator <i>D1</i>:			
MAC source	MAC dest	IP source	IP dest

- (b) We want that another machine (SCADA) also receives the measurements sent by the sensor. What is required for that at the sensor *S* and at the SCADA ?

2. Below is the python code of an application, one proposed by Homer and one proposed by Bart

```
#####  
# HOMER, Jan 2020, EPFL  
import socket  
HOST = 'localhost'  
PORT = 5002  
sock = socket.socket(socket.AF_INET6, socket.SOCK_STREAM)  
sock.setsockopt(socket.SOL_SOCKET, socket.SO_REUSEADDR, 1)  
sock.bind((HOST, PORT))  
sock.listen(1)  
while True:  
    connection, addr = sock.accept()  
    while True:  
        data = sock.recv(16).decode()  
        print("received:", data)  
        if data != b'':  
            sock.sendall(data.encode())  
        else:  
            print("No more data from", addr)  
            break  
    connection.close()
```

```
#####  
# BART, Jan 2020, EPFL  
import socket  
HOST = 'localhost'  
PORT = 5002  
sock = socket.socket(socket.AF_INET6, socket.SOCK_STREAM)  
sock.setsockopt(socket.SOL_SOCKET, socket.SO_REUSEADDR, 1)  
sock.bind((HOST, PORT))  
sock.listen(1)  
while True:  
    connection, addr = sock.accept()  
    while True:  
        data = connection.recv(16).decode()  
        print("received:", data)  
        if data != b'':  
            connection.sendall(data.encode())  
        else:  
            print("No more data from", addr)  
            break  
    connection.close()
```

(a) One of these two programs works, the other causes a run-time error. Which is the correct program ? Justify your answer.

(b) Say what is true about the version of the application that works (there is exactly one correct answer¹).

- i. it is a UDP server
- ii. it is a UDP client
- iii. it is a TCP server
- iv. it is a TCP client

¹For question 2 (b), 100% of the points of the question are obtained if only the correct answer is selected; if only one incorrect answer is selected, the grade is negative and is -33%; if zero answer or more than one answer is selected the grade is 0.

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

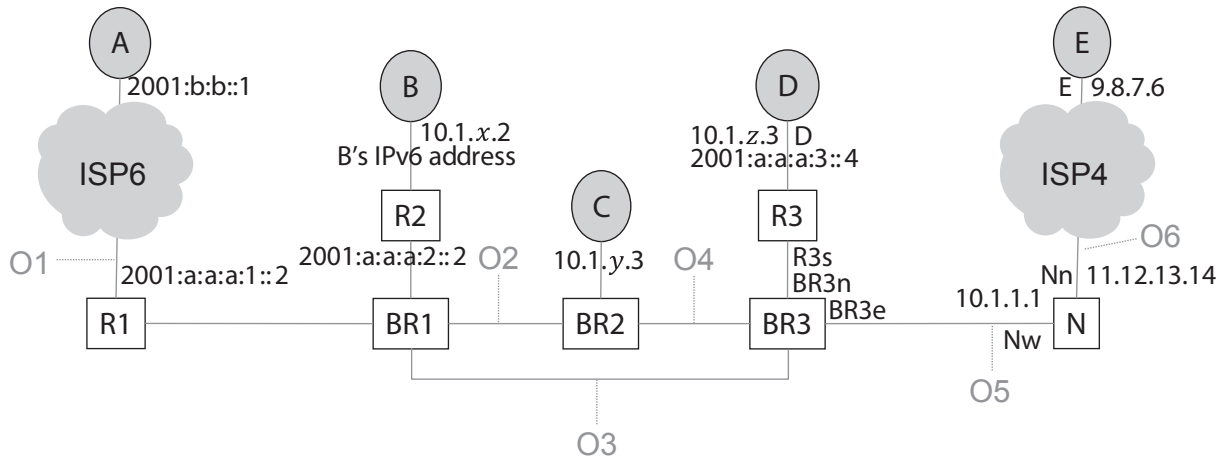
	Host may fragment	Router may fragment	Router may re-assemble
With IPv4			
With IPv6			

4. Both $H1$ and $H2$ are IPv6-only hosts (see figure). They communicate via IPv6, however $H1$'s local router A does not have native IPv6 access on its wide-area side. Instead, A receives IPv4 public access from an IPv4 provider and uses a tunnel broker offered by IPv6 provider P . P delegates to $H1$'s local area network the prefix $2001:1:2:3::/64$; A is the tunnel client and B is the tunnel server. The IPv6 address at A 's end of the tunnel is $2001:1:a:b::2$ and the IPv6 address at B 's end of the tunnel is $2001:1:a:b::1$. The IPv6 addresses of $H1$ and $H2$ are shown on the figure. $H1$ sends one UDP message to $H2$. The message is small and fits in one IP packet even after encapsulation. We observe the IP packet resulting from this activity at observation points 1, 2 and 3. Give the IP addresses and protocol / next header in the following table.

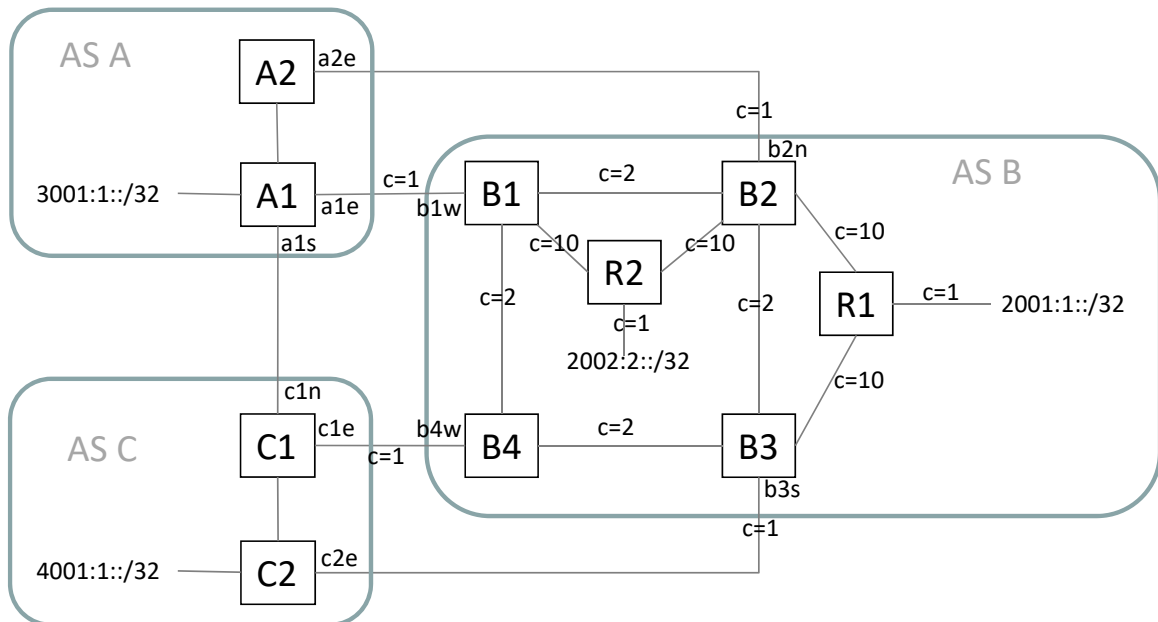
At observation point 1, towards $H2$:		
IPv6 source	IPv6 dest	Next Header
At observation point 2, towards $H2$:		
IPv4 source	IPv4 dest	Protocol
At observation point 3, towards $H2$:		
IPv6 source	IPv6 dest	Next Header

TCP/IP EXAM - FIGURES

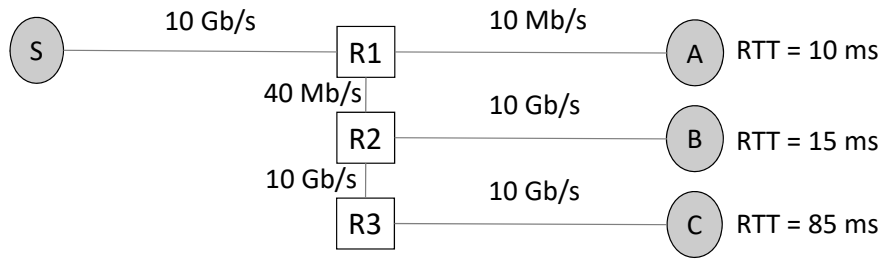
For your convenience, you can separate this sheet from the main document. Do not write your solution on this sheet, use only the main document. You do not need to return this sheet.




Problem 1




Problem 2



Problem 3



Problem 4, Question 1.



Problem 4, Question 4.