

SCIPER: _____

First name: _____ Family name: _____

EXAM
TCP/IP NETWORKING
Duration: 3 hours
With Solutions

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.

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 running the spanning tree protocol; N is an IPv4 NAT; $R1, R2$ and $R3$ are routers. Plain lines are physical connections. $O1$ to $O6$ are points where we observe traffic.

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

MAC addresses are denoted with e.g. $D, E, BR3e, \dots, Nn$. If you need to make assumptions about addresses, please write them explicitly.

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. ISP4 allocates to N -north the address $11.12.13.14$.

- (a) Give possible values of x, y, z (in the IPv4 addresses of B, C, D) and the network masks at N -west, B and C .

Solution. There are many possible solution; they all require that C and N -west are in the same subnet, and B, C are not and are in different subnets. One possible answer is $255.255.255.0$ for the network mask at all four locations, and $y = 1, x = 2, z = 3$.

- (b) D sends one UDP datagram to E . We observe the resulting packet at observation point $O5$. What are the MAC source and destination addresses ?

Solution. MAC source: $R3s$; MAC destination: Nw

- (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.

Solution. The spanning tree protocol has disabled one of the links. If the disabled link is $O3$, the packets are visible at $O2$ and $O4$ (and not at $O3$). If the disabled link is $O2$ or $O4$, the packets are visible at $O3$ only.

- (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
C 's IPv4 address	E 's IPv4 address	udp	4567	443
B 's IPv4 address	E 's IPv4 address	udp	4567	443

At observation point $O6$, towards E :				
IP source	IP dest	protocol	source port	dest. port
11.12.13.14	E 's IPv4 address	udp	e.g. 23456	443
11.12.13.14	E 's IPv4 address	udp	e.g. 23457	443

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 $2001:a:a:a:1::2$.

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

- (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		X
2001:a:a:a:2::3		X
2001:a:a:a:5::3	X	
2001:a:a:a:f::3	X	
2001:a:a:a:11::3		X

Justification: The address has to correspond to the delegated prefix $2001:a:a:a::/76$, i.e. should start with $2001:000a:000a:000a:000x:...$ where x is any hexadecimal digit. Since the subnet mask is 80 bits, the subnet prefix should be precisely $2001:000a:000a:000a:000x::/80$ with $x \in \{0, 1, 2, \dots, f\}$. Furthermore, the subnet must be different for B , D , $R1$ -north and $R2$ -south therefore x cannot be equal to 1, 2 or 3. Thus x should be in $\{4, 5, \dots, e, f\}$.

- (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 $O1$, towards A :				
IP source	IP dest	prot	src. port	dest. port
B's IPv6 address	A's IPv6 address	tcp	4567	443
D's IPv6 address	A's IPv6 address	tcp	4567	443

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 $A1, A2, \dots$

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 ?
B2	2001:1::/32	b2n	B	yes
A1	2001:1::/32	b1w	B	no
Justification: Because OSPF-internal is redistributed into BGP, B1 announces $2001:1::/32$, next-hop = b1w, AS-Path=B to A1 and B2 announces $2001:1::/32$, next-hop = b2n, AS-Path=B to A2. Also, inside the IBGP mesh of AS A, A1 announces $2001:1::/32$, next-hop = b1w, AS-Path=B to A2. A2 has two routes to $2001:1::/32$, the decision process prefers the former, because it is learnt by E-BGP.				

- (b) Still at time t_1 , what is the list of BGP routes received by A_2 with destination = 4001:1::/32 ? Which route is selected as best route by A_2 ? Give your answer in the table below, with a short justification (put as many rows as necessary).

At A_2 :				
From BGP Peer	Destination Network	BGP Next-Hop	AS-Path	Best route ?
B2	4001:1::/32	b2n	B C	no
A1	4001:1::/32	c1n	C	yes
Justification: A_2 receives the route 4001:1::/32, next-hop = b2n, AS-Path=B C from B1 and the route 4001:1::/32, next-hop = c1n, AS-Path=C from A1. The former is preferred because it has a shorter AS-Path.				

- (c) Still at time t_1 , what is the list of BGP routes received by R_2 with destination = 4001:1::/32 ? Which route is selected as best route by R_2 ? Give your answer in the table below, with a short justification (put as many rows as necessary).

At R_2 :				
From BGP Peer	Destination Network	BGP Next-Hop	AS-Path	Best route ?
B3	4001:1::/32	c2e	C	yes
B4	4001:1::/32	c1e	C	no
Justification: R_2 peers with all BGP routers inside AS B . B4 has selected as best route 4001:1::/32, next-hop = c1e, AS-Path= C (received from C1), for the same reasons as in a previous question (EBGP is preferred over IBGP). Since it is learnt by EBGP, this route is sent by B4 to R_2 (over a TCP connection). Similarly, R_2 receives from B3 the route 4001:1::/32, next-hop = c2e, AS-Path= C. B1 receives a route to 4001:1::/32 via EBGP from A1, but this route is not preferred because it has a longer AS path than the route that B1 learns from B3 or B4. Since such a route is learnt internally, it is not sent to R_2 . Same with B2. R1 has no external BGP connection therefore it cannot send any route to R_2 (because routes learnt by IBGP cannot be propagated by IBGP). In summary, R_2 has received only two routes. They have same AS path length, both come from IBGP, have the same distance to next-hop, therefore the one coming from the lowest BGP identifier is used, i.e., B3.				

- (d) Still at time t_1 , R_2 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 ?

Solution. The only matching route is 4001:1::/32, with next-hop=c2e; with recursive table lookup, R_2 uses the next-hop to c2e. There is only one shortest path, via B2, therefore R_2 sends the packet to B2.

At B2, the route to 4001:1::/32 also has next-hop=c2e; with recursive table lookup, B2 uses the next-hop to c2e, which is B3.

At B3, the route to 4001:1::/32 also has next-hop=c2e, which is on-link with B3 and the packet is sent directly to c2e.

In summary, the path is R_2 -B2-B3-c2e; recursive table lookup is performed at R_2 and B2.

- (e) At time $t_2 > t_1$ the link $A_1 - C_1$ breaks. At time $t_3 > t_2$, the BGP protocol has converged again. what is the list of *valid* BGP routes received by A_2 with destination = 4001:1::/32 ? Which route is selected as best route by A_2 ? 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 ?
B2	4001:1::/32	b2n	B C	yes
A1	4001:1::/32	b1w	B C	no

Justification: A1 notices the failure, removes its adjacency to C1, and marks all routes that were received from C1 as invalid. This triggers a new application of the decision process and the choices of best routes is modified, for example the route 4001:1::/32, next-hop=c1n, as-path=C is no longer best route and is replaced by 4001:1::/32, next-hop=b1w, as-path=B C. Since the invalid route was previously sent to A2, A1 now sends a WITHDRAW update to A2, who now prefers the route obtained from B2 (same AS path, EBGp is preferred).

- (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 ?

Solution. The bogus route is the only one available to the prefix 3001:1:2::/48 and is therefore selected by all BGP routers. B2 now has two overlapping routes:

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

3001:1::/32, next-hop = a2e, as-path=A

The packet matches both but, by the longest prefix match, the former route is applied. The packet is sent along one of the two shortest paths to c1e, i.e. it packet is sent along B2 – B1 – B4 or B2 – B3 – B4 and leaves AS B by the link B4 – C1.

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 ?

Solution. R1 does not have in its forwarding table any external prefix because it does not run BGP and BGP is not redistributed into OSPF. The only match is for the default route, therefore the packet is sent by R1 to B2. B2 is onlink so there is no recursive table lookup.

At B2, there is a forwarding table entry to 4001:1::/32 obtained from BGP; it is the same as before, the BGP next-hop is c2e, with recursive table lookup the packet is sent to B3.

At B3, same again, the packet is sent to c2e, which is onlink and does not require recursive table lookup.

In summary, the path is R1 – B2 – B3 – c2e with recursive table lookup at B2.

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 B_1, B_2, B_3 and B_4 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. R_1 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 ?

Solution. R_1 does not have in its forwarding table any external prefix because it does not run BGP and BGP is not redistributed into OSPF. However, by OSPF, R_1 obtains two equal cost next hops to $::/0$, one is B_2 and the other is B_3 . At R_1 the only match is for $::/0$ therefore the packet is sent by R_1 to either B_2 or B_3 , typically using a hash of the destination address, which are all different, therefore with equal probability. About 50% of the packets go via B_2 and 50% via B_3 .

At B_2 , there is a forwarding table entry obtained from BGP; the best route selected by BGP to $4001:1::/32$ is the one obtained by B_2 from B_3 , with next-hop= $c2e$ (the other one, obtained from B_4 is ranked equally but has a larger BGP identifier). Using recursive table lookup, these packets go to B_3 .

At B_3 , the packets are sent directly to C_2 (same reasons as in previous questions).

In summary, half of the packets follow the path $R_1 - B_2 - B_3 - c2e$ and half of the packets follow the path $R_1 - B_3 - c2e$.

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 . R_1, R_2 and R_3 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 ?

Solution. The constraints on the rates are $x, y, z \geq 0$ and

$$\begin{cases} x + y + z \leq 10000 \\ x \leq 10 \\ y + z \leq 40 \\ y \leq 10000 \\ z \leq 10000 \end{cases}$$

All constraints except the second and third are redundant. Therefore, the set of constraints is

$$\begin{cases} x \leq 10 \\ y + z \leq 40 \end{cases} \quad (1)$$

A Pareto-efficient allocation is one that cannot be increased unilaterally. We must have $x = 10$ since otherwise we can always increase x (which does not appear in the other constraint. With the same reasoning, we must have $y + z = 40$. Therefore, a Pareto-efficient allocation must satisfy

$$\begin{cases} x = 10 \\ 0 \leq y \leq 40 \\ z = 40 - y \end{cases} \quad (2)$$

Conversely, it is straightforward to verify that any allocation that satisfies these constraints is Pareto-efficient.

In summary, the Pareto-efficient allocations are those given by (??).

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 ?

Solution. The rates of the flows according to max-min fairness are obtained by water-filling. At the end of the first step all flows receive 10Mb/s and the flow $S \rightarrow A$ becomes frozen due to using a saturated link (between R_1 and A). At the end of step 2 y and z increase to 20, where they freeze due to the saturated link between R_1 and R_2 . Thus, the values are $x = 10, y = 20$ and $z = 20$.

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

Solution.

We have to solve the optimization problem:

$$\text{maximize } (\log x + \log y + \log z) \text{ over } x, y, z \geq 0$$

subject to all link constraints, i.e. subject to (??). The problem is symmetric in y, z and since we know that there is a unique optimum, the optimum must have $y = z$. Hence we can assume $y = z$ and we have to solve

$$\text{maximize } (\log x + 2 \log y) \text{ over } x, y \geq 0 \text{ subject to}$$

$$\begin{cases} x \leq 10 \\ 2y \leq 40 \end{cases}$$

The value $x = 10, y = 20$ is feasible and is an upper bound to any other feasible solution, therefore, since \log is an increasing function, it is the optimal solution. Finally we obtain $x = 10, y = 20, z = 20$, i.e. the same as the max-min fair allocation.

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

Solution.

(The rates are now affected by the RTTs compared to the previous question.) Here are two possible ways to solve this question.

Solution 1. We use the fact that TCP rates can be approximated as a utility-fair allocation where the utility of a flow with rate x is $\frac{\sqrt{2}}{\tau} \arctan \frac{x\tau}{\sqrt{2}}$ where τ is the RTT. We thus obtain the following optimization problem:

$$\text{maximize } \left(\frac{\sqrt{2}}{\tau_1} \arctan \frac{x\tau_1}{\sqrt{2}} + \frac{\sqrt{2}}{\tau_2} \arctan \frac{y\tau_2}{\sqrt{2}} + \frac{\sqrt{2}}{\tau_3} \arctan \frac{z\tau_3}{\sqrt{2}} \right) \text{ over } x, y, z \geq 0$$

subject to all link constraints, i.e., subject to $x \leq 10$ and $y + z \leq 40$. The groups of variables x and y, z can be separated and we obtain that, at the optimum, $x = 10$. To find y, z we have to solve

$$\text{maximize } \left(\frac{\sqrt{2}}{\tau_2} \arctan \frac{y\tau_2}{\sqrt{2}} + \frac{\sqrt{2}}{\tau_3} \arctan \frac{z\tau_3}{\sqrt{2}} \right) \text{ over } y, z \geq 0$$

subject to $y + z \leq 40$. Now we have a single constraint and we must have $y + z = 40$, i.e. $z = 40 - y$. We are thus left with a single variable problem:

$$\text{maximize } \left(\frac{\sqrt{2}}{\tau_2} \arctan \frac{y\tau_2}{\sqrt{2}} + \frac{\sqrt{2}}{\tau_3} \arctan \frac{(40 - y)\tau_3}{\sqrt{2}} \right) \text{ over } y \geq 0$$

subject to $y \leq 40$.

By taking the derivative of the above objective and setting it to zero, we obtain:

$$\frac{1}{1 + \frac{y^2\tau_2^2}{2}} - \frac{1}{1 + \frac{(40-y)^2\tau_3^2}{2}} = 0$$

After some algebra, we obtain

$$y\tau_2 = (40 - y)\tau_3$$

, i.e.

$$y = \frac{40\tau_3}{\tau_2 + \tau_3} = 34$$

It follows that the rates, in Mb/s, are $x = 10, y = 34$ and $z = 6$.

Solution 2. Using the throughput-loss formula and observing that the drop rates are the same for flows to B and C (which is true with ECN and a fair router) and the fact that TCP rate allocation is Pareto-efficient.

Obviously, $x = 10$ Mb/s since $S \rightarrow A$ does not compete with the other flows $S \rightarrow A, S \rightarrow B$ over a bottleneck link.

For y, z : by the loss-throughput formula, they are inversely proportional to RTTs (all other terms in the formula are identical for both); thus

$$y \times 15 = z \times 85 \quad (3)$$

Pareto-efficiency implies that

$$y + z = 40 \quad (4)$$

Combining the two we obtain $y = 34, z = 6$.

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’’ ?

Solution. If the flows were using TCP without ECN we could obtain by the loss throughput formula the loss ratio of the link i.e., the fraction q of the packets that is in average lost. Now, with ECN we could think that q is in average the fraction of packets that should be marked as congestion experienced in order to avoid having losses. Thus, the loss throughput formula tells us that

$$\theta = \frac{C \cdot MSS}{T\sqrt{q}}.$$

Therefore the fraction of marked packets is $q = \left(\frac{C \cdot MSS}{T\theta}\right)^2$. Replacing with values of flow $S \rightarrow C$, we obtain:

$$q = \left(\frac{1.22 \cdot 10^4 \text{bits}}{0.01 \text{sec} \cdot 10^7 \text{bits/sec}}\right)^2 = (1.22/10)^2 = 0.122 * 0.122 \approx 0.015 = 1.5\%$$

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 $R3$?

Solution. We have four flows in this case. Let x , y , z and t be the rates of flows $S \rightarrow A$, $S \rightarrow B$, $S \rightarrow R3$ and $R3 \rightarrow C$ (in Mb/s).

First let us compute the rate at which $R3$ can receive from S . We have a similar problem as in Question 4, with the difference that the RTTs for B and $R3$ are equal. Using the same reasoning, we obtain $x = 10$, $y = 20$ and $z = 20$.

Second, $R3$ receives data at rate 20Mb/s and the capacity of the link $R3 - C$ is much larger, therefore $R3$ will retransmit at the same rate to C . Thus C receives from S at a rate of 20Mb/s.

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 the sensor ? Put the answer in the table below (see on the figure sheet for device addresses).

At sensor S:			
MAC source	MAC dest	IP source	IP dest
S	33:33:0:2:0:1	fd24::1:1	ff35::4:3:2:1

At Data Concentrator $D1$:			
MAC source	MAC dest	IP source	IP dest
S	33:33:0:2:0:1	fd24::1:1	ff35::4:3:2:1

- (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 ?

Solution. SCADA needs to start a program that listens to the multicast group $ff35::4:3:2:1$. Nothing needs to be done at S .

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.

Solution. This is a TCP server; when it accepts a connection, it creates a new socket, here called `connection`. Data should be sent or received on the socket called `connection`, not on the socket called `sock`. Bart's program is correct. Homer's program receives a run-time error (e.g. "SOCKET NOT CONNECTED").

- (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

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	true	true	false
With IPv6	true	false	false

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

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.

the IP addresses and protocol / next header in the following table.


At observation point 1, towards <i>H2</i>:		
IPv6 source	IPv6 dest	Next Header
2001:1:2:3:a:b:c:d	3000:b::99	UDP

At observation point 2, towards <i>H2</i>:		
IPv4 source	IPv4 dest	Protocol
9.8.7.6	1.2.3.4	IPv6


At observation point 3, towards <i>H2</i>:		
IPv6 source	IPv6 dest	Next Header
2001:1:2:3:a:b:c:d	3000:b::99	UDP

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.