EXAM
TCP/IP NETWORKING
Duration: 4 hours + including technical time for uploading solution.
With Solutions

Jean-Yves Le Boudec

January 2021

INSTRUCTIONS

1. All four problems have the same weight. They are independent and can be done in any order.
2. Briefly justify your answer. For grading, the justification is as important as the solution itself.
3. You may need to make additional assumptions in order to solve some of the questions. If this happens, please describe such assumptions explicitly.

Problem 1

Consider the network in Figure 1. H1 to H6 are hosts. R1 and R2 are routers. B1 to B6 are bridges. O1 to O12 are observation points. All machines are dual-stack.

N3 and N1 operate as NATs for both IPv4 and IPv6: the internal port is B3 [resp. B2] and the external port is the port towards B4 [resp. B6].

N2 operates as a NAT for IPv4 (with internal ports towards observation point O1) and as a router for IPv6.

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. There is no other system or interface than shown on the figure.

H1B3, B3H1, B3N3, N3B3, N3B4, B4N3, B4B5, B4B1, B4B6, B1N2, B1B6, B1B4, B1B5, N2B1, N2R1, R1N2, R1H6, H6R1 are MAC addresses.

n4, ..., k2 are ipv4 adresses and are noted in lowercase. N4, ..., K2 are ipv6 adresses and are noted in capital letter.

Some IPv4 and IPv6 addresses are known or partially known, as follow:

1. The IPv4 adress t4 is 10.213.141.212.
2. The IPv4 adress m1 is 8.2.3.7.
3. The IPv4 adress n4 is 192.168.47.134.
4. The IPv4 adress s2 is 192.168.47.253.
5. The IPv4 adress w1 is 192.168.47.46.
6. The IPv4 adress z6 is 8.2.7.11.
7. The IPv4 adress h1 is 10.213.141.163 .

1. The IPv6 adress V3 is fe80::18d6.
2. The IPv6 adress H1 is 3390:b:4:5:q::1.
3. The IPv6 adress T4 is 3390:b:4:5:e9::1.
4. The IPv6 adress W1 is fe80::a6c8.
5. The IPv6 adress X1 is 3390:b:4:5:e9:74b2::2.
6. The IPv6 adress N4 is fe80::ea71.

1. Give possible values for the IPv4 address x1 at N3 and for the network mask at this interface.
   This interface of N3 is on the same subnetwork as 10.213.141.163 and 10.213.141.212, therefore it should have the same subnet prefix. One possibility is that the network mask in this subnet is 255.255.255.0 and a possible IP address is x1 $=$ 10.213.141.num with $num \in \{0, \cdots, 255\} \notin \{163, 212\}$.
2. Give possible values for the IPv4 address v3 of H3 and for its network mask.
   H3 should be in the same subnet as H5, which is in the block of private addresses 192.168.0.0/16, therefore it should also be in that block. For example, if the network mask is 255.255.255.0, then a possible address is 192.168.47.134 (it should be different from 192.168.47.253 and 192.168.47.46).
3. H1 sends one UDP packet to H5, H3 sends one UDP packet to H6 and H2 sends one UDP packet to H4. We look for these UDP packets at observation points O5, O4, O10, O2 and O6, i.e. at the lines between bridges. At which of these observation points will we observe the packets?
   Possible answer 1: The spanning tree protocol will disable loops in the topology. Assume for example that B4 is elected as bridge. The links marked O4 and O10 are disabled as they are not on the shortest paths to B4. The packet H1 $\rightarrow$ H5 is seen at O5; the packet H3 $\rightarrow$ H6 is seen at O5 and O6; the packet H2 $\rightarrow$ H4 is seen at O6. (Optional: Here we are assuming that the learning phase of bridges has

stabilized, otherwise packets may be broadcast to the spanning tree and are visible at all non disabled links, i.e. at O5, O2 and O6.)
Possible answer 2: The spanning tree protocol will disable loops in the topology. Assume for example that B6 is elected as bridge and that the shortest path from B5 to B6 is via B4 (there is another shortest path and it could also have been chosen). The links marked O10 and O6 are disabled as they are not on the shortest paths to B4. The packet H1 $\rightarrow$ H5 is seen at O5; the packet H3 $\rightarrow$ H6 is seen at O4; the packet H2 $\rightarrow$ H4 is seen at O5 and O4. (Optional: Here we are assuming that the learning phase of bridges has stabilized, otherwise packets may be broadcast to the spanning tree and are visible at all non disabled links, i.e. at O5, O4 and O2.)

4. H1 and H5 send several small UDP packets each to H6, at about the same time, over IPv4, to udp port 510. By chance, the port source numbers are both equal to 16566. We observe these packets at observation points O7, O8 [resp. O12, O9], O1 and O11. What source and destination IP addresses and port numbers do we see ?

Packet H1 $\rightarrow$ H6:

at	srce IP addr	dest IP addr	srce port	dest port
O7	192.168.47.46	8.2.7.11	16566	510
O8	x1	8.2.7.11	some number, say p	510
O1	x1	8.2.7.11	p	510
O11	8.2.3.7	8.2.7.11	some number, say q	510

Packet H5 $\rightarrow$ H6:

at	srce IP addr	dest IP addr	srce port	dest port
O12	192.168.47.46	8.2.7.11	16566	510
O9	h1	8.2.7.11	some number, say $p' \neq p$	510
O1	h1	8.2.7.11	p'	510
O11	8.2.3.7	8.2.7.11	some number, say $q' \neq q$	510

5. (continued) At O7, O8 and O1, what source and destination MAC addresses do we see for the packet sent by H1 to H6 ?

at	srce MAC	dest MAC
O7	H1B3	N3B3
O8	N3B4	N2B1
O1	N3B4	N2B1

6. H3 sends one IP packet to H6 and one to H2, and sets the value of TTL at the source to 48. What is the value of the TTL at destination for each of these packets?
At H6: $48 - 3 = 45$, at H2: $48 - 4 = 44$

7. The IPv6 network prefix lengths at N1B6 and N2B1 are both equal to 72. The IPv6 address of interface N1B6 is 3390:b:4:5:q::1 where q is a block of 16 bits. Give all the possible values of q (use hexadecimal notation).
The prefix must be the same as N2B1 (and N3B4) and must therefore be equal to 3390:b:4:5::/72 . Furthermore, the address must differ from the address of N3B4 (always true because of the block 74b2) and N2B1 thus the valid values of q are $q = 00y_2y_3$ with $y_2, y_3 \in \{0, 1, ..., f\}$ and $y_2y_3 \neq e9$.

8. H1 does a traceroute to fe80::18d6. We look for the traceroute packets at observation points O7, O8, O5, O4, O10, O2, O6, O9 and O12. What are the source and destination IP addresses observed at these points ?
The destination in a link-local address that is not allocated in the LAN of the source. This packet can only be sent to a destination in the same LAN as the source. The source will perform NDP address resolution, it will fail and no packet will be sent. We will not see any traceroute packet at any of these observation points (but we will see a NDP address resolution packet at O7).

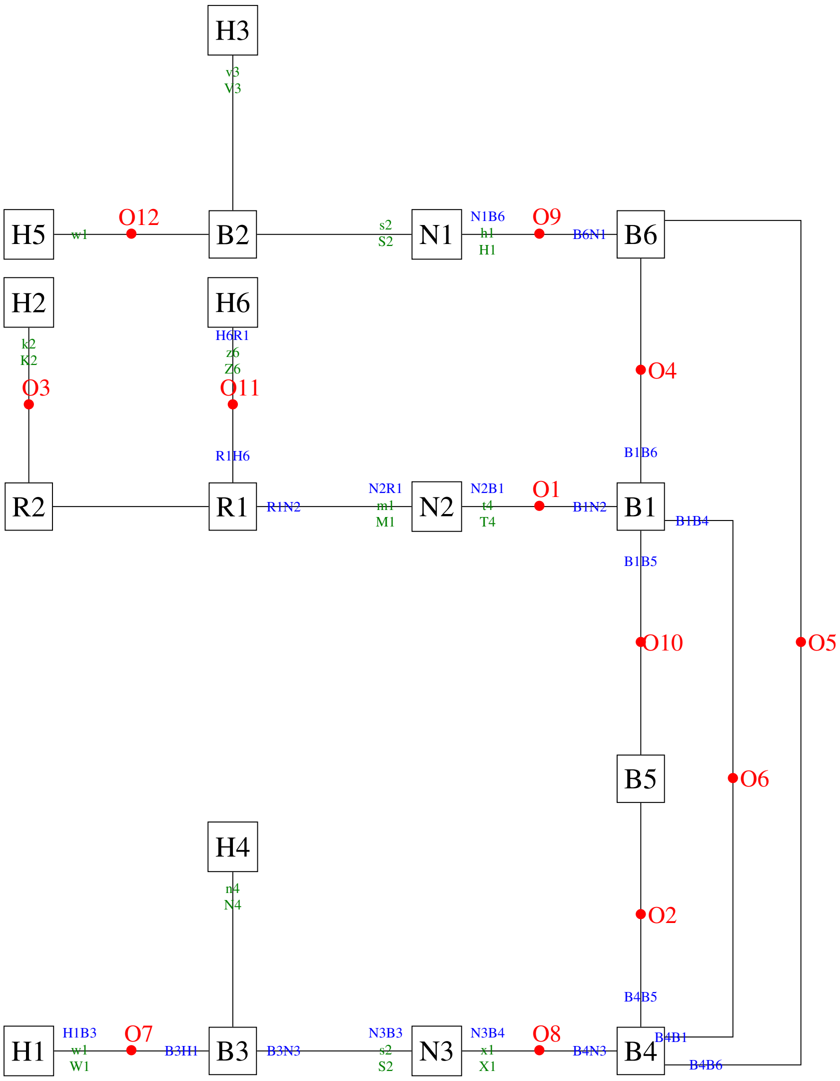

Figure 1: Network in Problem 1

PROBLEM 2

Consider the network in Figure 2. There are three ASs, S , V and R . Every node whose name begins with S belongs to AS S (same with V and R). Each AS uses OSPF with a single area and equal cost multipath. The OSPF cost of every link and every attached network is 1, unless otherwise specified.

The IP addresses at inter-router links are denoted as follows: at $S3$, the IP address of the interface towards $S1$ is $s3s1$, and similarly everywhere else (so for example the IP address at $R3$ of the interface towards $V1$ is $r3v1$, etc.)

The subnetwork 2614:9fca:1f7c::/48 is attached to node $S5$ and belongs to AS S .

All routers run BGP, unless otherwise specified. The BGP decision processes all 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, unless otherwise specified:

- All routers redistribute internal OSPF routes into BGP.
- Every BGP router injects its BGP routes into its forwarding table.
- The BGP routers do not redistribute external BGP routes into OSPF.
- No optional BGP attribute (such as MED, LOCAL-PREF etc.) is used in any BGP message.
- No aggregation is performed.
- There is no route announcement to 2614:9fca:1f7c::/48 or 2614:9fca:1f7c:cd56::/64 other than the ones described in the text and the ones that derive from that.
- The BGP policies accept all routes that are acceptable by the standard operation of BGP.

1. When OPSF and BGP have converged, what are all the BGP routes to destination 2614:9fca:1f7c::/48 stored by $V4$ in its RIB-In (give the source BGP router, the AS-PATH and NEXT-HOP attribute of the routes) ? What are the routes to 2614:9fca:1f7c::/48 advertized by $V4$ to other BGP routers ?
   $V4$ has received:
   - from $S3$ over E-BGP: 2614:9fca:1f7c::/48 NEXT-HOP = $s3v4$ AS-PATH = S
   - from $V2$ over I-BGP: 2614:9fca:1f7c::/48 NEXT-HOP = $s2v2$ AS-PATH = S

   No BGP route is received from $V5$ nor $V3$ because the best route selected by these (the routes via $V4$) is a route learnt over I-BGP. $V4$ keeps the two routes and selects as best route the route received from $S3$ because it is received via E-BGP and has priority over the other route, received via I-BGP. This route is learnt by E-BGP therefore $V4$ forwards it to all its I-BGP peers, i.e. it sends
   - to $V5, V3, V1$ and $V2$ over I-BGP: 2614:9fca:1f7c::/48 NEXT-HOP = $s3v4$ AS-PATH = S
2. Same question with $V5$.
   $V5$ has received:
   - from $R4$ over E-BGP: 2614:9fca:1f7c::/48 NEXT-HOP = $r4v5$ AS-PATH = R S
   - from $V4$ over I-BGP: 2614:9fca:1f7c::/48 NEXT-HOP = $s3v4$ AS-PATH = S
   - from $V2$ over I-BGP: 2614:9fca:1f7c::/48 NEXT-HOP = $s2v2$ AS-PATH = S

   No BGP route is received from $V3$ because the best route selected by $V3$ (the route via $V4$) is a route that $V3$ learnt over I-BGP.

$V5$ keeps the three routes and selects as best route the route received from $V4$ because it has the shortest AS-PATH and shortest path to NEXT-HOP. This route is learnt by I-BGP therefore $V5$ does not forward it with I-BGP. However, it forwards it with E-BGP, i.e. it sends

- to $R4$ over E-BGP: 2614:9fca:1f7c::/48 NEXT-HOP = $v5r4$ AS-PATH = V S

3. Router $V3$ has a packet to forward to destination address 2614:9fca:1f7c:cd56:ff1::9c53. Which path will this packet take inside AS V ? (Give the path as a sequence of routers, including the first router out of AS V).
BGP at $V3$ learns one route to 2614:9fca:1f7c::/48 from $V4$ and one from $V2$. They have equal AS-PATH and equal cost to next-hop, so $V3$ chooses the one with smallest BGP identifier; assume it is $V4$. The NEXT-HOP is $s3v4$ and the packet goes $V3-V5-V4-S3$. (If the smallest BGP identifier would be $V2$'s, the packet would go $V3-V1-V2-S2$).
4. Due to a configuration error, $R4$ announces to $V5$ the bogus route

   2614:9fca:1f7c:cd56::/64 AS-PATH = R NEXT-HOP = $r4v5$

   and this route is accepted by all BGP routers in AS V. After BGP and OSPF have stabilized again, router $V1$ has two packets to forward, one with destination address 2614:9fca:1f7c:cd57:b2eb::35c6, and one with destination address 2614:9fca:1f7c:cd56:ff1::9c53. Which paths will these packets follow from $V1$ to the exit of AS V ? (Give the path as a sequence of routers, including the first router out of AS V).
   The route to 2614:9fca:1f7c:cd56::/64 is different from the previously seen routes (though it overlaps). $V5$ propagates the route in the I-BGP mesh of AS V and since it is the only one, it is selected as best. Thus $V1$ has received two routes by BGP: one to 2614:9fca:1f7c::/48 , with NEXT-HOP = $s2v2$, one to 2614:9fca:1f7c:cd56::/64 , with NEXT-HOP = $r4v5$. Both are accepted and installed in the routing table.

   (a) The first packet matches 2614:9fca:1f7c::/48 and not 2614:9fca:1f7c:cd56::/64 . The route to 2614:9fca:1f7c::/48 is used and the packet goes $V1-V2-S2$.
   (b) The second packet matches both 2614:9fca:1f7c:cd56::/64 and 2614:9fca:1f7c::/48 . By longest prefix match the route to 2614:9fca:1f7c:cd56::/64 is used and the packet goes $V1-V3-V5-R4$.

5. $R4$ is patched and the configuration error that led to the bogus announcement is removed. Furthermore, we now modify the operation of AS V:

   - router $V1$ does not run BGP (but continues to run OSPF);
   - all other routers in AS V redistribute external BGP routes into OSPF with distance at the point of redistribution equal to 21.

   The rest is unchanged. After BGP and OSPF have stabilized: Router $V1$ has a packet to forward to destination address $2614 : 9fca : 1f7c : cd56 : ff1 :: 9c53$. Which path will this packet take inside AS V ? (Give the path as a sequence of routers, including the first router out of AS V). Compare to the paths that would have been followed if the configuration were the same as in Questions 1, 2 and 3.
   The packet follows the route to $2614 : 9fca : 1f7c :: /48$. $V1$ learns about the route $2614 : 9fca : 1f7c :: /48$ via OSPF. The prefix $2614 : 9fca : 1f7c :: /48$ is redistributed into OSPF by $V4$ and $V2$ with distance 21. $V1$ computes the shortest path and finds that it is $V1-V2$ and the next-hop is $v2v1$. The path of the packet is $V1-V2-S2$ ($-S5$) (same as before).
6. Router $V3$ now has packets to send to a large number of destination addresses such as $2614 : 9fca : 1f7c : cd56 : n : k : z : h$, where n, k, z and h each represents a sequence of four hex digits. Which paths will these packets take ? (Give the path as a sequence of routers, including the first router out of AS V). Compare to the paths that would have been followed if the configuration were the same as in Questions 1, 2 and 3.
   OSPF at $V3$ computes the shortest path and finds that there are two equal-cost paths, namely $V3-V5-V4-S3$ and $V3-V1-V2-S2$. In the routing table of $V3$, OSPF writes two equal cost

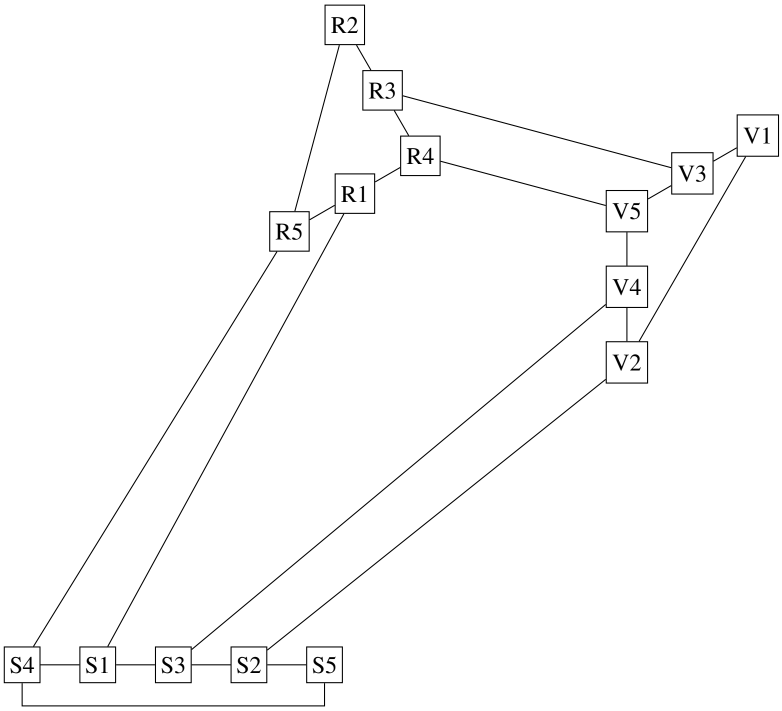

Figure 2: Network in Problem 2

routes, with next-hops $V5$ and $V1$.
By I-BGP, $V3$ also learns a route to $2614 : 9fca : 1f7c :: /48$, one from $V4$ and one from $V2$ and selects one of them. The corresponding route is installed in the routing table with an administrative distance larger than OSPF, so it is not used
So the routes written by OSPF are used. They have same distance to destination. Since ECMP is supported, the packets are sent with equal probability on the two paths $V3 - V5 - V4 - S3$ and $V3 - V1 - V2 - S2$.
With the previous configuration, only one of the two paths was used.

PROBLEM 3

Consider the network in Figure 3. $H0$, $H6$ and $H3$ are downloading very large contents from $H4$, simultaneously. The intermediate boxes are routers, unless otherwise specified. The line rates are shown in the figure with $m = 200$, $u = 1064$, and $K = 40432$, in Mb/s. The links are full duplex Ethernets. Let p [resp. w, s] be the rate at which $H0$ [resp. $H6$, $H3$] downloads contents from $H4$. All other flows (including ACKs) can be neglected.

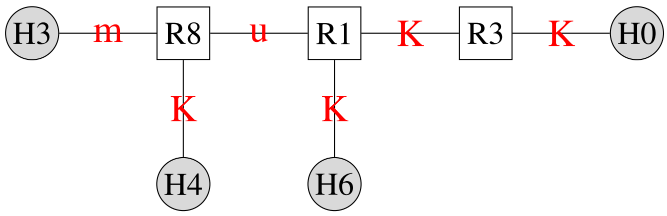

Figure 3: Network in Problem 3

1. Assume the rates are allocated by some central bandwidth manager. Which of the following rate allocations are Pareto-efficient?
   (a) $p = 399, w = 133, s = 100$.
   (b) $p = 798, w = 266, s = 200$.
   (c) $p = 532, w = 532, s = 200$.
   (d) $p = 798, w = 266, s = 100$.

   The constraints on the rates are

$$\begin{aligned} p + w &\leq u \\ s &\leq m \end{aligned}$$

   as the other constraints are redundant.
   (a) we can increase any of the three rates alone, thus this is not Pareto-efficient.
   (b) all constraints are satisfied with equality, thus we cannot increase any of the rates and the allocation is Pareto-efficient.
   (c) idem
   (d) s can be increased unilaterally thus this is not Pareto efficient.
2. Assume the rates are allocated by some central bandwidth manager according to max-min fairness. What are all the possible rate allocations?
   There is a unique max-min fair rate allocation. It is obtained by water-filling. In the first step, all rates are set to $m = 200$ and s is frozen. In the second step, the remaining rates are set to $u/2 = 532$. Thus the allocation is $p = u/2 = 532$, $w = u/2 = 532$, $s = 200$.
3. Assume the rates are allocated by some central bandwidth manager according to proportional fairness. What are all the possible rate allocations?
   There is a unique proportionally fair rate allocation. It is obtained by maximizing $\log(p) + \log(w) + \log(s)$ subject to the constraints

$$\begin{aligned} p + w &\leq u \\ s &\leq m \end{aligned}$$

   By symmetry (and since the optimum is unique) we must have $p = w$. The problem becomes: maximize $2\log(p) + \log(s)$ subject to

$$\begin{aligned} 2p &\leq u \\ s &\leq m \end{aligned}$$

The objective and the constraints are decoupled so the problem can be separated into two independent optimization problems, one with the variable p and one with the variable s. The solution is obtained when $p = u/2 = 532$ and $s = m = 200$. Thus the proportionally fair allocation is the same as the max-min fair allocation.

4. $H0$, $H6$ and $H3$ are now using TCP-Reno to perform their transfer from $H4$. The RTTs are: $\text{RTT}_{\text{H0}} = 113$, $\text{RTT}_{\text{H6}} = 15$ and $\text{RTT}_{\text{H3}} = 5$, in ms. All have same MSS equal to 1344, in bits. ECN is enabled in all hosts and all routers. What is the rate obtained by every source ?
Since the link $R8 - H4$, has a very large capacity, the constraint for $H3$, is only the capacity m of the link $H3 - R8$. Similarly, the constraint for $H0, H6$ is only the shared link $R1 - R8$.
The bottleneck link for $H3$ is used only by $H3$ therefore (Pareto efficiency of TCP), $s = m = 200$.
Hosts $H0$ and $H6$ have a single bottleneck link, hence we can assume all ECN marks occur at this link and both flows have the same marking probability. By the loss throughput formula, the rates are inversely proportional to RTTs. By Pareto-efficiency, $p + w = u$. Thus

$$\frac{p}{\frac{1}{\text{RTT}_{\text{H0}}}} = \frac{w}{\frac{1}{\text{RTT}_{\text{H6}}}} = \frac{p+w}{\frac{1}{\text{RTT}_{\text{H0}}} + \frac{1}{\text{RTT}_{\text{H6}}}} = \frac{u}{\frac{1}{\text{RTT}_{\text{H0}}} + \frac{1}{\text{RTT}_{\text{H6}}}}$$

and

$$p = u\frac{\frac{1}{\text{RTT}_{\text{H0}}}}{\frac{1}{\text{RTT}_{\text{H0}}} + \frac{1}{\text{RTT}_{\text{H6}}}} = 124.7, w = u\frac{\frac{1}{\text{RTT}_{\text{H6}}}}{\frac{1}{\text{RTT}_{\text{H0}}} + \frac{1}{\text{RTT}_{\text{H6}}}} = 939.3$$

5. $R3$ is now working as a web tunnel (web proxy) for $H0$, also using TCP Reno with ECN and with same MSS as all hosts. The RTT between $R3$ and $H4$ is the same as between $H6$ and $H4$.
In this question only, the link between $H0$ and $R3$ is wireless with a loss rate equal to 7.32 %, and the RTT between $H0$ and $R3$ is equal to 4.26 ms. (there is no congestion loss on this link, only wireless transmission losses). What is the rate achieved by $H0$?
On one hand, the rate p achieved by $H0$ is constrained by the achievable rate between $R3$ and $H0$. By the loss-throughput formula, the latter is equal to

$$\frac{1.22 \times 1344}{\sqrt{7.32 \times 0.01} \times 4.26 \times 10^{-3}} \text{ b/s} = 1422636 \text{ b/s} = 1.423 Mb/s$$

thus $p \leq 1.423$.
On the other hand, $R3$ will obtain for $H0$ a rate equal to the min of the above and of the rate x that $R3$ obtains from TCP. $R3$ is competing with $H6$ only (we can ignore $H3$ as before, because $200 << 40432$.
Option 1:
TCP maximises $U(x) + U(w)$ subject to the constraints

$$\begin{aligned} x &\leq 1.423 \\ w &\leq 40432 \\ x + w &\leq 1064 \end{aligned}$$

where the utility function U is the same for $R3$ and $H6$ because they have same RTT and same MSS. Since $40432 > 1064 + 1.423$, the second constraint can be ignored and, at the optimum we must have

$$x + w = 1064$$

Therefore TCP maximizes $U(x) + U(1064 - x)$ over $x \in [0; 1.423]$. Since U is concave, its derivative is decreasing and therefore it is positive when $x < 1064/2$ and negative when $x > 1064/2$. Since $1.423 < 1064/2$, the max is obtained when $x = 1.423$.
Option 2:

TCP aims at giving the same rate to both $R3$ and $H6$ because they have same RTT and same MSS; the common rate that TCP would like to give is $1064/2$. But the first constraint imposes $x < 1064/2$ therefore TCP will give 1.423 to $R3$ and the rest to $H6$.
In all cases, we obtain that the rate achieved by $H0$ is 1.423.

PROBLEM 4

1. The Agostinho Neto University (UAN) offers a repository of media files (lyrics, recorded audios and videos) to the public. The repository is hosted by a number of media servers. The UAN campus has access to IPv6 and to IPv4; the media servers use only IPv4 addresses. The media servers use old hardware and software, partly written in Cobol; migrating them to IPv6 would require expensive upgrades and is not an option. There are n media servers, and each of them has a DNS name of its own. The renderer application on the user device accesses the media files by using URLs, such as https://musica-hoji-ya-henda.uan.ao/2340993.flac where musica-hoji-ya-henda.uan.ao is the DNS name of the media server that hosts the file of interest.
   You joined the UAN Network Team as engineer and you are asked to design a solution for allowing external users that have only IPv6 connectivity to access the repository. Your solution should not modify the UAN media servers; it should also not modify the IPv6 configuration of external users. Please sketch at least two different solutions (one paragraph each).
   Solution 1: Address Translation. UAN installs a NAT64 between its IPv6 network and its IPv4 network. It must reserve a prefix of 96 bits taken from its IPv6 address space for embedding the v4 addresses of media servers. It must also reserve a block B of IPv4 addresses to map IPv6 external user addresses to IPv4 addresses that can be understood by the media servers. A private IPv4 address block can be used for that. The DNS must be updated accordingly: the IPv4-embedded IPv6 addresses of all media servers must be written as AAAA records in the DNS. When an external user issues a request to access a file such as https://musica-hoji-ya-henda.uan.ao/2340993.flac, its DNS resolver will obtain the IPv4 mapped IPv6 address of musica-hoji-ya-henda.uan.ao, which the NAT64 will automatically translate to the IPv4 address of musica-hoji-ya-henda.uan.ao. The external user's IPv6 address is also mapped to one of the IPv4 addresses in the reserved block. The mapping cannot be static as the IPv6 address space is larger than B; therefore the mapping must be dynamic and kept in memory. If the block B is large (for example if the entire $10/8$ block is used) there is no need to modify the port numbers as there are enough IPv4 addresses to represent all active users at any single point in time); else port numbers must be mapped as in a regular NAT.
   Solution 2: Server-side proxy. UAN installs a dual stack server (server side proxy) that terminates all TLS connections and TCP connections from external users, and pipes them into TCP connections with the media servers. The DNS must be updated accordingly: all names of all media servers must be associated with AAAA records that point to the server side proxy. The cryptographic keys associated with the media servers for TLS connections must be present in the server side proxy.
2. The network of the National Security Agency of Borduria (NSAB) uses addresses in the block 133.133.0.0/16 and is organized with subnet prefixes of length 24 bits. For example, the subnet 133.133.99.0/24 is used by the management team, the subnet 133.133.1.0/24 is used for the accounting department and the subnet 133.133.42.0/24 contains the information servers used by all NSAB departments. NSAB decided to deploy the SuperSecureIx operating system in all of their computers. This is a very secure operating system designed by the NSAB engineers. With SuperSecureIX, the IP address of a computer is assigned to a user and cannot be modified (DHCP is not used). Unfortunately, the NSAB engineers did not take the TCP/IP Networking course and, instead, read random things found on forums. As a result, they believe that, in class B networks, the network mask is 255.255.0.0 (Class B is an obsolete classification of IP address blocks) and since the NSAB addresses are in class B, they hardcoded this subnet mask in SuperSecureIx. As this is a very secure operating system, no one dares modify it.
   Immediately after deploying SuperSecureIX, many network problems appeared. Can you propose a solution ? Your solution should not involve any change to end-user devices but can involve the network infrastructure (which, fortunately, is not based on SuperSecureIX and can be modifed).
   The problem is with ARP: when a host A in one subnet sends a packet to a host B in another subnet of NSAB, the wrong mask lets A believe that B is in the same subnet. Therefore, A does ARP for

B's MAC address. Since B is not onlink with A, B does not see the response and the communication does not work. NSAB hosts cannot communicate with hosts in different NSAB subnets.
A Possible Solution: routers are configured with the correct net masks and perform proxy ARP for all NSAB addresses that are not on the same link from which the ARP is received. Thus, A's next-hop router will respond with its own MAC address and will receive the packet. Since it does not use the bogus OS, it can forward it to the next hop, and the final router on the path will deliver it to B.