SCIPER: ____________________

First name: ____________________ Family name: ____________________

EXAM
TCP/IP NETWORKING
Duration: 3 hours
With Solutions

Jean-Yves Le Boudec

January 2016

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 will probably need to make additional assumptions in order to solve some of the questions. In that case, please write those assumptions down explicitly.
4. Figures are on a separate sheet, for your convenience.
5. No documents, no electronic equipments are allowed.
6. Justify every answer with a short explanation.

Problem 1

Consider the network in Figure 1. A, B, C, D and E are hosts; BR is a bridge; X and $Y1, Y2, Y3, Y4$ are network boxes that can be configured in different ways, as explained next. $O1, O2$ and $O3$ are observation points where we observe traffic in both directions of the link. Some selected IPv4 addresses are shown, as well as some selected MAC addresses (denoted with A, B, C, D, E, Xe and $Y1w$). You may need to specify other IP or MAC addresses. All links are full duplex Ethernet. We assume that all machines are correctly configured (unless otherwise specified) and proxy ARP is not used.

1. In this question $X, Y1, Y2, Y3$ and $Y4$ are routers, running RIP with all link costs equal to 1.

   (a) Give one possible value for the netmask at A and one for the netmask at B.

   Solution. For example, `255.0.0.0` would be a valid netmask for A as well as B.

   (b) A sends one ping message to B, B sends one ping message to C, C sends one ping message to D, D sends one ping message to E, and E sends one ping message to A. We assume that A and B have just booted and have empty ARP tables. We observe the traffic resulting from this activity at observation points $O1, O2$ and $O3$. At which of these points are the ARP Requests issued by A and B visible?

   Solution. $O1$ only, *i.e.*, the whole broadcast domain of the A's LAN.

   What is the target IP address in the ARP request issued by A ? by B?

   Solution. by A: the IP address of B, i.e. `10.2.1.2`. By B: the IP address of B's default gateway, which should be the west interface of X.

   We observe the ping request packets resulting from this activity at observation points $O1, O2$ and $O3$. What are the MAC and IP source and destination addresses in such packets ? What is the TTL field, knowing that the TTL value is equal to 64 in all IPv4 packets generated by all hosts in this problem ? Put your answers in the tables below. Recall that to denote the MAC address of, for example, the "south-side" interface of $Y4$, you should write $Y4s$, etc.

At observation point $O1$:				
MAC source	MAC dest	IP source	IP dest	TTL
B	Xw	10.2.1.2	9.3.1.2	64
Xw	A	9.5.1.2	10.1.1.2	61

At observation point $O2$:				
MAC source	MAC dest	IP source	IP dest	TTL
$Y4w$	$Y1e$	9.5.1.2	10.1.1.2	63

At observation point $O3$:				
MAC source	MAC dest	IP source	IP dest	TTL
$Y1s$	$Y2n$	10.2.1.2	9.3.1.2	62

2. In this question X is a NAT box; the WAN port is Xe.
   $Y1, Y2, Y3, Y4$ are routers, running RIP with all link costs equal to 1.

   (a) Are the netmasks at A and B obtained in the previous question still valid ?

   Solution. Yes. A and B still belong to the same broadcast domain.

(b) A sends one ping message to B, B sends one ping message to C, C sends one ping message to D, D sends one ping message to E, and E sends one ping message to A. We observe the ping request packets resulting from this activity at observation points $O1, O2$ and $O3$. What are the MAC and IP source and destination addresses in such packets ? Put your answers in the tables below.

At observation point $O1$:			
MAC source	MAC dest	IP source	IP dest
B	Xw	10.2.1.2	9.3.1.2
Xw	A	9.5.1.2	10.1.1.2

At observation point $O2$:			
MAC source	MAC dest	IP source	IP dest
$Y4w$	$Y1e$	9.5.1.2	9.2.1.2

At observation point $O3$:			
MAC source	MAC dest	IP source	IP dest
$Y1s$	$Y2n$	9.2.1.2	9.3.1.2

3. In this question BR is a VLAN switch, which gives different VLAN labels to A and B. Furthermore, $X, Y1, Y2, Y3, Y4$ are routers (none of them is a NAT). $Y1, Y2, Y3, Y4$ run RIP with all link costs equal to 1. X does not run RIP. None of the boxes uses proxy ARP.

(a) Are the netmasks at A and B obtained in question 1 still valid ?

Solution. No.

(b) Give a possible configuration of the routing table at X that enables full connectivity in this network.

Destination Network	Interface	Gateway
10.1/16	Xw VLAN A	onlink
10.2/16	Xw VLAN B	onlink
9/8	Xe	onlink

(c) A sends one ping message to B, B sends one ping message to C, C sends one ping message to D, D sends one ping message to E, and E sends one ping message to A. We observe the ping request packets resulting from this activity at observation point $O1$. What are the MAC and IP source and destination addresses in such packets ? Put your answers in the table below.

At observation point $O1$:			
MAC source	MAC dest	IP source	IP dest
A	Xw	10.1.1.2	10.2.1.2
Xw	B	10.1.1.2	10.2.1.2
B	Xw	10.2.1.2	9.3.1.2
Xw	A	9.5.1.2	10.1.1.2

4. In this question $Y1, Y2, Y3, Y4$ are bridges. X is a router.

A sends one ping message to B, B sends one ping message to C, C sends one ping message to D, D sends one ping message to E, and E sends one ping message to A. We observe the ping request packets resulting from this activity at observation points $O2$ and $O3$. What are the MAC and IP source and destination addresses in such packets ? Put your answers in the tables below.

Solution. The bridges run the spanning tree algorithm, which disables one of the ports. Depending on which port is disabled, there are 4 possible answers.

First, we assume that the link between Y3 and Y4 is disabled.

At observation point $O2$:			
MAC source	MAC dest	IP source	IP dest
D	E	9.4.1.2	9.5.1.2
E	Xe	9.5.1.2	10.1.1.2
At observation point $O3$:			
MAC source	MAC dest	IP source	IP dest
Xe	C	10.2.1.2	9.3.1.2
D	E	9.4.1.2	9.5.1.2

Otherwise, if link between $Y1$ and $Y2$ is disabled:

At observation point $O2$:			
MAC source	MAC dest	IP source	IP dest
Xe	C	10.2.1.2	9.3.1.2
E	Xe	9.5.1.2	10.1.1.2
At observation point $O3$:			
MAC source	MAC dest	IP source	IP dest

Otherwise, if link between $Y1$ and $Y4$ is disabled:

At observation point $O2$:			
MAC source	MAC dest	IP source	IP dest
At observation point $O3$:			
MAC source	MAC dest	IP source	IP dest
Xe	C	10.2.1.2	9.3.1.2
E	Xe	9.5.1.2	10.1.1.2

Otherwise, if link between $Y2$ and $Y3$ is disabled:

At observation point $O2$:			
MAC source	MAC dest	IP source	IP dest
C	D	9.3.1.2	9.4.1.2
E	Xe	9.5.1.2	10.1.1.2
At observation point $O3$:			
MAC source	MAC dest	IP source	IP dest
C	D	9.3.1.2	9.4.1.2
Xe	C	10.2.1.2	9.3.1.2

SCIPER: ____________________

First name: ________________________ Family name: ________________________

PROBLEM 2

Consider the network in Figure 2. Boxes $B1, ...B7, C1, D1$ are routers.

All routers in AS B, namely $B1, B2, ...B7$, run an IPv4 distance vector routing protocol with route poisoning and with infinity $= 256$. Unless otherwise specified, the cost of a link between two routers is 1 and the cost from a router to a directly attached network is 1.

There is one IPv4 network between consecutive routers on the figure, with subnet prefix of length 24. All networks shown on the figure, including `30.20.0.0/24` and `40.20.0.0/24` are considered interior by all routers in B. All routers in AS B originate their directly attached networks into the distance vector routing protocol.

Routers do not perform aggregation, unless explicitly mentioned.

1. Assume that BGP is not yet enabled in routers of AS B. Give a possible value of the routing table at router $B2$, at a time t_1 such that the interior routing protocol has stabilized. Give the values in the table below (do not give the value of the "interface" field).

At $B2$		
Destination Network	Next-Hop	Distance

Solution.

Destination Network	Next-Hop	Distance
2.1.2.0/24	on link	1
2.3.4.0/24	2.2.3.3	2
2.4.5.0/24	2.2.3.3	3
2.5.6.0/24	2.2.3.3	4
2.2.3.0/24	on link	1
2.1.7.0/24	2.1.2.1	2
30.20.0.0/24	2.1.2.1	2
2.6.7.0/24	2.1.2.1	3
40.20.0.0/24	2.1.2.1	3

2. BGP is now enabled inside AS B, in routers $B1$ and $B7$ only. Other routers in AS B do not run BGP. At time $t_2 > t_1$, $B1$ receives from $C1$ the announcement

```
66.66/16, AS path = C E F, NEXT-HOP=30.20.0.3
66.66.0/17, AS path = C E F, NEXT-HOP=30.20.0.3
```

No other BGP announcements were received by $B1$ before this one. Explain what BGP protocol actions are performed by $B1$ upon receiving these routes. Say in particular to which BGP routers, if any, $B1$ will send announcements as a result.

Solution. At B1, the Adj-RIB-in is updated to include the two routes. As these are the only two routes for the networks 66.66/16 and 66.66.0/17, they are both promoted to best by the BGP decision process. Then, the Adj-RIB-out is updated to include both the routes and B1 will send announcements to its I-BGP peer, i.e., B7.

At time $t_3 > t_2$, BGP has stabilized and $B1$ redistributes the network prefixes it has learnt from BGP into the interior routing protocol, with cost $= 100$ for networks learnt via I-BGP, and cost $= 50$ for networks learnt via E-BGP. Note that $B7$ does not (yet) redistribute.

$B1$ does its job and sends an interior routing update to all its neighbors inside AS B. Assume that no other redistribution occurs and the interior routing protocol stabilizes again at time t_4. Give a possible value of the routing table at router $B2$ at t_4.

At $B2$		
Destination Network	Next-Hop	Distance

Solution. At t3, the only routes that B1 has learnt via BGP are for networks 66.66/16 and 66.66.0/17 and both have been learnt through E-BGP. Hence, the interior routing protocol updates sent will be with distance metric 50.

Destination Network	Next-Hop	Distance
2.1.2.0/24	on link	1
2.3.4.0/24	2.2.3.3	2
2.4.5.0/24	2.2.3.3	3
2.5.6.0/24	2.2.3.3	4
2.2.3.0/24	on link	1
2.1.7.0/24	2.1.2.1	2
30.20.0.0/24	2.1.2.1	2
2.6.7.0/24	2.1.2.1	3
40.20.0.0/24	2.1.2.1	3
66.66/16	2.1.2.1	51
66.66.0/17	2.1.2.1	51

3. Then $B7$ also redistributes BGP into the interior routing protocol with cost $= 100$ for networks learnt via I-BGP, and cost $= 50$ for networks learnt via E-BGP. We assume that at time $t_5 > t_4$ both the interior routing protocol and BGP have stabilized again and at that time $B7$ receives from $D1$ the announcement:

```
66.66/16, AS path = D F, NEXT-HOP=40.20.0.4
```

   Assume no other routing message than previously shown was sent by either $C1$ or $D1$. Explain what BGP protocol actions are performed by $B7$ upon receiving these routes. Say in particular to which BGP routers, if any, $B7$ will send announcements as a result. Also say what BGP computations will be performed at the BGP routers that receive these announcements, if any.

   Solution. B7 will promote this route to the best (shorter AS path). The only BGP peer of B7 is B1. B7 will update its Adj-RIB-out with this best route and send an announcement to B1. Then, B1 also promotes the route received from B7 through I-BGP (as it has a shorter AS path) as the best path.

4. At time $t_6 > t_5$ BGP has stabilized again inside ASs B, C and D and at that time, $B1$ redistributes the network prefixes it has learnt from BGP into the interior routing protocol, with cost $= 100$ for networks learnt via I-BGP, and cost $= 50$ for networks learnt via E-BGP. Note that $B7$ does not (yet) redistribute.Then $B1$ does its job and sends an interior routing update to both of its neighbors inside AS B. Assume that no other redistribution occurs no other interior routing message was generated in the time interval $[t_5, t_6]$.

   Explain which computations $B2$ and $B7$ will perform upon receiving the interior routing update from $B1$.

   Solution. $B2$ and $B7$ will receive a internal routing update for 66.66/16 with cost 100. As this update is received from $B1$, which is the next hop for the network 66.66/16 at both $B2$ and $B7$, the update is accepted without any computation. Consequently, the distance metric of the route for network 66.66/16 is updated to 100 at $B2$ and 101 at $B7$.

5. Then $B7$ redistributes BGP into the interior routing protocol with cost $= 100$ for networks learnt via I-BGP, and cost $= 50$ for networks learnt via E-BGP and let $t_7 > t_6$ be a time at which the interior routing protocol has stabilized again. Give a possible value of the routing table at $B1$ at t_7.

At $B1$		
Destination Network	Next-Hop	Distance

Solution.

Destination Network	Next-Hop	Distance
2.1.2.0/24	on link	1
2.1.7.0/24	on link	1
30.20.0.0/24	on link	1
2.2.3.0/24	2.1.2.2	2
2.3.4.0/24	2.1.2.2	3
2.4.5.0/24	2.1.2.2	4
2.6.7.0/24	2.1.7.7	2
2.5.6.0/24	2.1.7.7	3
40.20.0.0/24	2.1.7.7	2
66.66/16	2.1.7.7	51
66.66.0/17	30.20.0.3	50

$B1$ has a packet to forward with destination address `66.66.1.2`. Which path will this packet follow ?

Solution. The longest prefix match results in the routing rule for 66.66.0/17. So, the packet will be sent to $C1$. After this, it will follow the path through AS's E and F.

6. At time $t_8 > t_7$ the link between $B1$ and $B2$ breaks; the loss of the link is detected immediately by $B2$. $B2$ does its job and sends an interior routing update to its neighbor $B3$. Assume that no other message was generated in the time interval $[t_6, t_8]$.

   Explain which computation $B3$ will perform upon receiving the interior routing update from $B2$ and give a possible value of the routing table at $B3$ just after performing these computations.

At $B3$		
Destination Network	Next-Hop	Distance

Solution. $B3$ sets the distance metric of all routes with next-hop B2 to 256 (unreachable).

Destination Network	Next-Hop	Distance
2.2.3.0/24	on link	1
2.1.2.0/24	2.2.3.2	256
2.1.7.0/24	2.2.3.2	256
30.20.0.0/24	2.2.3.2	256
40.20.0.0/24	2.2.3.2	256
66.66/16	2.2.3.2	256
66.66.0/17	2.2.3.2	256
2.6.7.0/24	2.2.3.2	256
2.3.4.0/24	on link	1
2.4.5.0/24	2.3.4.3	2
2.5.6.0/24	2.3.4.3	3

Let $t_9 > t_8$ be the time at which the routing protocol stabilizes again (the link between $B1$ and $B2$ is definitively lost). Give a possible value of the routing table at $B3$ at time t_9.

At $B3$		
Destination Network	Next-Hop	Distance

Solution.

Destination Network	Next-Hop	Distance
2.3.4.0/24	on link	1
2.4.5.0/24	2.3.4.4	2
2.5.6.0/24	2.3.4.4	3
2.6.7.0/24	2.3.4.4	4
2.1.7.0/24	2.3.4.4	5
30.20.0.0/24	2.3.4.4	6
40.20.0.0/24	2.3.4.4	5
66.66/16	2.3.4.4	54
66.66.0/17	2.3.4.4	55

7. Assume we change the configuration of $B1$ and $B7$ such that they do not re-distribute BGP into the interior routing protocol. Propose one alternative method by which the routes learnt via BGP by $B1$ and $B7$ can be also learnt by all routers inside AS B. Are there any drawbacks to your proposed alternative ?

   Solution. 1. By configuring appropriate default routes on all other routers. Drawback: Cannot cope with link-failures.

   1. By running BGP on all routers. Drawback: Expensive

SCIPER: ____________________

First name: ________________________ Family name: ________________________

PROBLEM 3

Consider the network in Figure 3 on the figure sheet.

- Hosts $A, B1, B2, B3, B4$ 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 or Bi and S. We also neglect the impact of the acknowledgement flows in the reverse direction.
- The round trip time between A and S is 100 ms, and the round trip time between Bi and S is 20 ms. These numbers include all processing times.
- 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.

1. Assume that some bandwidth manager is used, which allocates rates to flows according to max-min fairness. What are the values of the rates of the flows $S \to A$ and $S \to Bi$?
   Solution. The unique max-min fair solution is obtained by water filling. Here, the bottleneck is the link S-$R1$ which carries 5 flows. Hence, the max-min fare solution is given by assigning $x_i = 21/5 = 4.2$ Mb/s for flows $i = 1, ..., 5$.
2. Same question with proportional fairness instead of max-min fairness.
   Solution. Same solution as max-min allocation. This is true due to the fact that all the constraints apart from the bottleneck S-$R1$ are inactive, and $\sum_{i=1}^{5} \log(x_i)$ subject to $\sum_{i=1}^{5} x_i \leq 21$ is maximized when $x_i = 21/5 = 4.2$ Mb/s. Indeed, the maximum is attained when the inequality is satisfied with equality, and due to symmetry and uniqueness of allocation, the solution is attained for equal rates.

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

   Solution. There are at two possible ways to solve this question.

   Solution 1. We use the fact that TCP maximizes $\sum_{i=1}^{5}(\sqrt{2}/\tau_i)\arctan(x_i\tau_i/\sqrt{2})$ subject to $\sum_{i=1}^{5} x_i = 21$. For the flows to Bi, by symmetry, $y \triangleq x_1 = x_2 = x_3 = x_4$ as the RTT $\tau_i = 0.02$ sec for $i = 1, ..., 4$. For the flow to A, let $x \triangleq x_5$ denote the corresponding flow. We have the constraint $4y + x = 21$. Hence, we should solve the optimization

$$\max \left[4\frac{\sqrt{2}}{0.02} \arctan\left(\frac{0.02y}{\sqrt{2}}\right) + \frac{\sqrt{2}}{0.1} \arctan\left(\frac{0.1(21-4y)}{\sqrt{2}}\right)\right]$$

   subject to $0 \leq y \leq 21/4$. By taking derivatives it is easy to see that the maximum is attained when

$$0.02y = 0.1(21 - 4y),$$

   which yields $y = 21/4.2 = 5$ Mb/sec. It follows that $x = 21 - 4y = 1$ Mb/sec.

   Solution 2. Same result can be attained using the throughput-loss formula assuming the drop rate is the same for flow x and y, namely using $x/y = 0.02/0.1 = 0.2$. Namely, $y = 5x$. Now we use the constraint $4y + x = 21$ and obtain that $x = 1$ Mb/sec as before.

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

   Solution. If the two flows (x and y) 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 $q = (\frac{C \cdot MSS}{T\theta})^2$. Replacing with values of flow y (flow x gives the same result), we obtain:

$$q = \left(\frac{1.22 \cdot 10^4 bits}{0.02sec \cdot 5 \cdot 10^6 bits/sec}\right)^2 = (1.22/10)^2 = 0.122 * 0.122 \approx 0.015 = 1.5\%$$

5. We continue to assume that the flows are using TCP and ECN, but now A cheats and uses a smart piece of software that allows it to open several TCP connections and use them in parallel in order to transfer one single file. Assume that A uses k TCP connections. What should the value of k be in order for A to obtain the same throughput as $B1$?
   Solution. Similarly to question (3), there are two possible ways.
   Solution 1. We have the constraint $4y + kx = 21$, and the maximum of the total utility is attained when
   $$0.02y = 0.1\frac{21-4y}{k}.$$
   This implies $y = 21/(4 + 0.2k)$. Therefore, the total throughput from S to A is given by
   $$kx = 21 - 4y = \frac{4.2k}{4 + 0.2k}.$$
   Hence, in order to have the same throughput to A as to $B1$, we require $y = kx$, or $21 = 4.2k$, which gives $k = 5$ flows.
   Solution 2. We use the throughput-loss formula assuming the drop rate is the same for flow x and y, giving $y = 5x$. Since we require $kx = y$, the value of k is then trivially 5.
6. Assume now that $R2$ is an application layer gateway instead of a router. We assume that all flows are using TCP with ECN. We assume that A is not cheating and is using only one TCP connection with the application layer gateway. The round trip time from S to $R2$ is 20 msec; from $R2$ to A it is 90 msec; from S to Bi it is 20 msec (unchanged). We assume that the application layer gateway has infinite processing power and storage. What is the value of the rate of each flow ?
   Solution. When an application layer gateway is used instead of a router $R2$, the congestion control mechanism is broken into two parts. The first part controls the network before $R2$. The second part controls the network after $R2$. Since the RTT in the first part is the same for all 5 flows at the bottleneck, the obtained rates for every flow will be the max-min allocation, namely $21/5 = 4.2$ Mb/sec. Note that due to large capacity of the second part of network (from $R2$ to A) this rate can be supported there without any decrease.

SCIPER: ____________________

First name: ______________________ Family name: ______________________

PROBLEM 4

Consider the network in Figure 4.

- $D1$, $D2$ and $S2$ are dual-stack hosts. In particular, $S2$ can be configured either as a NAT64 or as a application layer gateway; $D1$ and $D2$ are DNS servers that serve requests for A and AAAA records over both IPv4 and IPv6;
- $R11$ and $R21$ are IPv4-only routers; $R13$ and $R23$ are IPv6-only routers;
- $R12$ and $R22$ are dual-stack routers.

1. What are the 3rd and 4th bits of B's IPv6 address ? (the first bit is the leftmost one).
   Solution. The first hexa digit of B's IPv6 address comprises the bits 1 to 4. Hence, since $(2)_{16} = (0010)_2$, the 3rd bit is 1 and the 4th is 0.
   Give the non-compressed form of B's IPv6 address.
   Solution. `2001:000a:000a:000a:0000:0000:0000:0002`
   Give a possible value for the IPv6 network mask at B.
   Solution. The network mask can be `/64` or longer. The two LANs, `2001:a:a:a::` and `2001:a:a:b::` need to have different prefixes.
2. A is a dual-stack host that receives only an IPv4 address from the network. A uses 6to4 to connect to the IPv6 Internet. Give a possible value for A's 6to4 IPv6 address.
   e.g., `2002:0303:0303:abcd:EUI`$_A$
   What should A's default gateway's IPv6 address be ?
   It should be a 6to4 address derived from the IPv4 prefix 192.88.99.0/24, which is the block of IPv4 addresses reserved for access to the IPv6 Internet; for example (as long as with correct prefix) `2002:c058:6301::`
3. B is an IPv6-only host. A sends one UDP packet to B. We observe this packet on the Ethernet link between A and $R11$. We find that the Ethernet header contains protocol type = 0800, *i.e.*, this is an IPv4 packet. Say what the IPv4 source and destination addresses in this packet are; what is the protocol field in the IPv4 header ?
   IPv4_{Src} = `3.3.3.3`
   IPv4_{Dst} = `192.88.99.1`
   Protocol = 41 (Because we have an IPv6 packet encapsulated in an IPv4 packet.)
   B responds to A with one single UDP packet. We observe this packet on the Ethernet link between B and $R23$. Which protocol type should we see in Ethernet header? Say what the IP (v4 or v6) source and destination addresses in this packet are. Please be consistent with your previous answer.
   Protocol (Ethertype) = `0x86DD` (Because Ethernet frame carries an IPv6 packet.)
   IPv6_{Src} = `2001:a:a:a::2`
   IPv6_{Dst} = `2002:0303:0303:abcd:EUI`$_A$
   Will this packet go via $R12$ or via $R22$? Justify your answer.
   Both $R12$ and $R22$ inject the prefix 2002::/16 into their IPv6 routing tables. Inside ISP2, the network prefix 2002::/16 will be learned via both BGP and IGP. The choice will depend on specific assumptions. For example, if ISP2 have no special policy, then the route learned via E-BGP from ISP1 will be chosen which implies $R12$. As another example, if ISP2 gives local routes higher weight, then the route learned via IGP will be chosen which implies $R22$.

4. S is an IPv4-only host that runs a web server. Now, $S2$ is a NAT64. B runs a web browser and sends an HTML request to S. No application layer gateway is used.

   (a) Consider the first TCP SYN packet related to this activity, sent by B to S; we observe this packet on the Ethernet link between B and $R23$. Which protocol type should we see in Ethernet header? Say what the IP (v4 or v6) source and destination addresses in this packet are.

   Here, NAT64 is used (S2). IPv6 addresses are mapped to valid IPv6 addresses from the dedicated block `64:ff9b::/96`. So we see:
   Protocol (Ethertype) = `0x86DD` (IPv6 packet)
   IPv6_{Src} = `2001:a:a:a::2`
   IPv6_{Dst} = `64:ff9b::202:202`

   (b) We also observe this packet on the Ethernet link between $R21$ and S. Which protocol type should we see in Ethernet header? Say what the IP (v4 or v6) source and destination addresses in this packet are. Specify any assumption you may need to make.
   We need to assume that S2 owns an IPv4 address pool to map IPv6 addresses. These IPv4 addresses cannot be addresses used by IPv4 hosts, so it cannot be `2.4.4.1`. Assume $S2$ owns the address pool `120.130.26.0/24`. We could see (here, HostID=33 for example):
   Protocol (Ethertype) = `0x0800` (IPv4 packet)
   IPv4_{Src} = `120.130.26.33`
   IPv4_{Dst} = `2.2.2.2`

   (c) S responds to B with a TCP SYN ACK packet. Will this packet go via $R22$ or $S2$? How does R21 know where to route this packet ? Justify your answer.
   Packet sent by S has: (should be consistent with previous IP address(es))
   IPv4_{Src} = `2.2.2.2`
   IPv4_{Dst} = `120.130.26.33`
   The packet will go via S2 (NAT64). $R21$ knows where to send the packet based on its destination address ($S2$ announces the prefix for the IPv4 pool it has at its disposal, *i.e.,* `120.130.26.0/24`).

   (d) S now downloads one object to B using HTML. One of the packets sent by S to B during the download is lost on the link $R23 - B$. There is no other packet loss in this scenario. Explain by which mechanism in which machines the loss will be repaired.
   Solution. The loss will be repaired by the TCP connection involving S and B, *i.e.,* it will be detected by S and the missing data will be retransmitted by S. ($S2$ will not be involved.)

5. A runs a web browser and sends an HTML request to S. Consider the first TCP SYN packet related to this activity, sent by A to S; we observe this packet on the Ethernet link between A and $R11$. Which protocol type should we see in Ethernet header? Say what the IP (v4 or v6) source and destination addresses in this packet are.
This is now IPv4 to IPv4 communication over an IPv4 network:
Protocol (Ethertype) = `0x0800` (IPv4 packet)
IPv4_{Src} = `3.3.3.3`
IPv4_{Dst} = `2.2.2.2`

6. Assume now that $S2$ is configured as an application layer gateway, and not a NAT64. B would like to download a video file from the web server at S; this server corresponds to the DNS name `ba.ba`. To this end, B would like to use $S2$ as a web proxy but we would not like to configure anything special (other than the normal IPv6 configuration) at B.

   (a) Is this possible ? Justify your answer.
   Solution. Yes. For example, if the DNS server $D2$ is configured to return the IPv6 address of $S2$ when queried about `ba.ba` in IPv6. Then any HTTP request for `ba.ba` from B is automatically relayed to $S2$, which acts as an ALG46.

   (b) We assume B found a way to use $S2$ as a web proxy and downloads a video file from S using HTTP. Once the transfer is successfully established, we observe the packets resulting from this activity on the links $R21-S2$ and $S2-R23$, in the direction $S \to B$. What are the IP addresses, protocol type/next header and TTL/HL that we see ? Put your answers in the table below.
   Solution.

   $R21 - S2$ ($S \to B$): IP source: `2.2.2.2` (S), IP dest: `2.4.4.1` ($S2$ v4),
   Protocol: `TCP`, TTL = 63 (past $R21$).

   $S2 - R23$ ($S \to B$): IP source: `2001:a:a:b::1`, IP dest: `2001:a:a:a::2`,
   Next Header: `TCP`, HL = 64.

   (c) One of the packets sent by S to B during the download is lost on the link $R23 - B$. There is no other packet loss in this scenario. Explain by which mechanism in which machines the loss will be repaired.
   Solution. The loss will be repaired by the TCP connection involving $S2$ and B, *i.e.,* it will be detected by $S2$ and the missing data will be retransmitted by $S2$. (S will not be involved.)

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. Do not return this sheet.

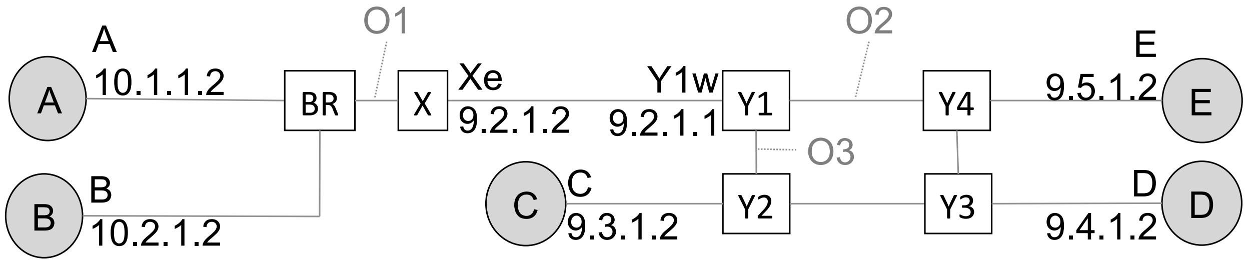

Figure 1: The network used in Problem 1, showing some selected addresses. You may need to specify other addresses.

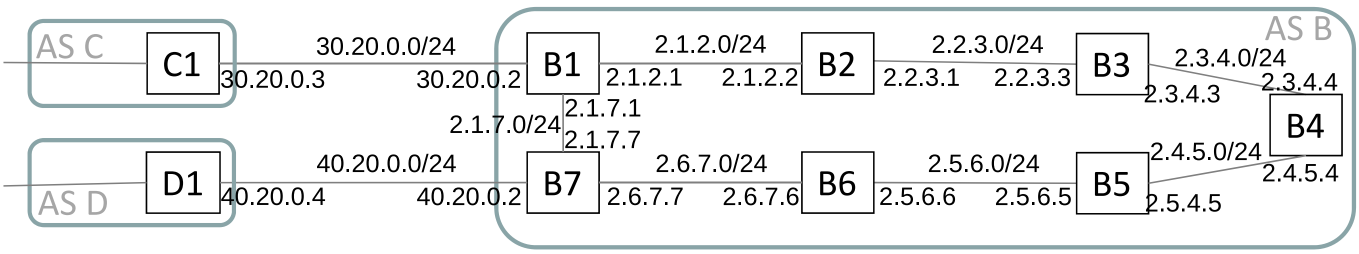

Figure 2: The network used in Problem 2.

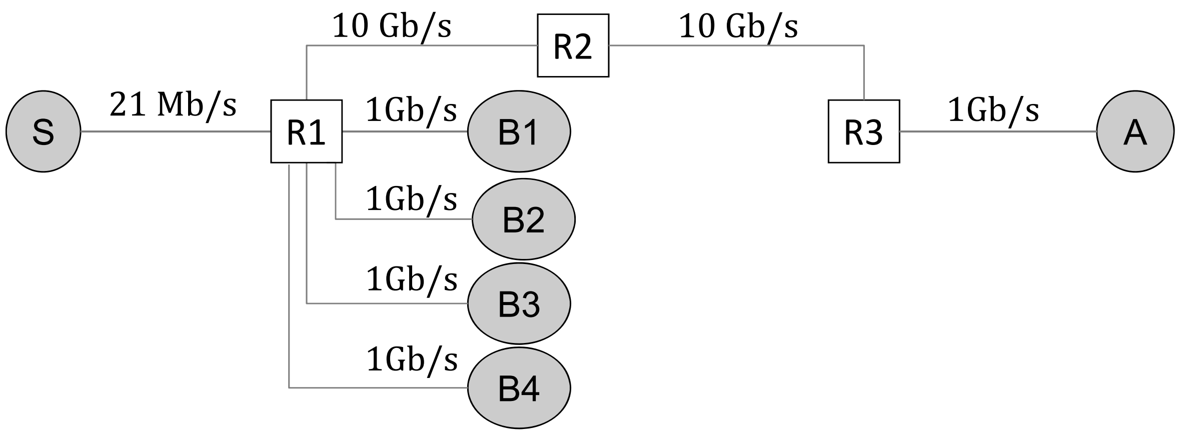

Figure 3: The network used in Problem 3.

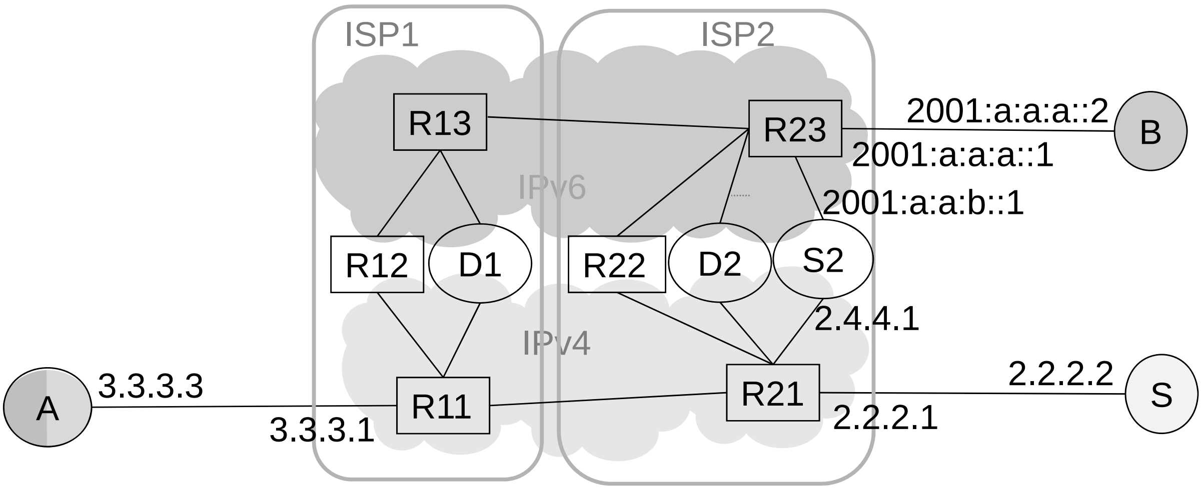

Figure 4: The network used in Problem 4.