Exercises
TCP/IP Networking

Jean-Yves Le Boudec

Fall 2009

Exercises marked with a ⋆ were given at exams in the past.

1 Module 1: TCP/IP Architecture

Exercise 1.1 *Elaine is setting in front of lrcpc3 and connects to machine ’ezinfo.ethz.ch’ by Telnet. A clairvoyant angel has read all the frames passing on the network. Here is the first packet resulting from this activity:*

```
ETHER:  ----- Ether Header -----
ETHER:
ETHER:  Packet 1 arrived at 19:03:32.39
ETHER:  Packet size = 60 bytes
ETHER:  Destination = ff:ff:ff:ff:ff:ff
ETHER:  Source      = 0:0:c0:b8:c2:8d
ETHER:  Ethertype = 0806
ETHER:
ARP:  ----- ARP/RARP Frame -----
ARP:
ARP:  Hardware type = 1
ARP:  Protocol type = 0800 (IP)
ARP:  Length of hardware address = 6 bytes
ARP:  Length of protocol address = 4 bytes
ARP:  Opcode 1 (ARP Request)
ARP:  Sender’s hardware address = 0:0:c0:b8:c2:8d
ARP:  Sender’s protocol address = 128.178.156.7, lrcpc3.epfl.ch
ARP:  Target hardware address = ?
ARP:  Target protocol address = 128.178.156.1, in-inr-e4.epfl.ch
```

1. *What is this frame used for in this exchange?*
2. *What stations receive this frame? What stations reply to it?*
3. *How can we determine if a frame is an ARP frame?*

Exercise 1.2 *Among the packets observed, shortly afterwards, we find the following two:*

```
ETHER:  ----- Ether Header -----
ETHER:
ETHER:  Packet 2 arrived at 19:03:32.39
ETHER:  Packet size = 74 bytes
ETHER:  Destination = 0:0:c:2:78:36
ETHER:  Source      = 0:0:c0:b8:c2:8d
ETHER:  Ethertype = 0800
ETHER:
IP:   ----- IP Header -----
IP:
IP:   Version = 4
IP:   Header length = 20 bytes
IP:   Type of service = 0x00
IP:         xxx. .... = 0 (precedence)
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length = 60 bytes
IP:   Identification = 2947
IP:   Flags = 0x0
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live = 64 seconds/hops
IP:   Protocol = 17
IP:   Header checksum = c2ba
IP:   Source address = 128.178.156.7
IP:   Destination address = 128.178.15.8, IP:   No options
IP:
UDP:  ----- UDP Header -----
UDP:
UDP:  Source port = 1267
UDP:  Destination port = 53 (DNS)
UDP:  Length = 40
UDP:  Checksum = B672
UDP:
DNS:  ----- DNS:   -----
DNS:
DNS:  ""
DNS:



ETHER:  ----- Ether Header -----
ETHER:
ETHER:  Packet 3 arrived at 19:03:32.40
```

```
ETHER:  Packet size = 202 bytes
ETHER:  Destination = 0:0:c0:b8:c2:8d, Western Digital
ETHER:  Source      = 0:0:c:2:78:36, Cisco
ETHER:  Ethertype = 0800
ETHER:
IP:   ----- IP Header -----
IP:
IP:   Version = 4
IP:   Header length = 20 bytes
IP:   Type of service = 0x00
IP:         xxx. .... = 0 (precedence)
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length = 188 bytes
IP:   Identification = 38579
IP:   Flags = 0x0
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live = 58 seconds/hops
IP:   Protocol = 17
IP:   Header checksum = 3d0a
IP:   Source address = 128.178.15.8,
IP:   Destination address = 128.178.156.7,
IP:   No options
IP:
UDP:  ----- UDP Header -----
UDP:
UDP:  Source port = 53
UDP:  Destination port = 1267
UDP:  Length = 168
UDP:  Checksum = 0000
UDP:
DNS:  ----- DNS:   -----
DNS:
DNS:  ""
DNS:
```

1. *What has happened?*
2. *What is lrcpc3's IP address? and ezinfo.ethz.ch's? What is the source IP address of packet 3. Which is the source MAC?*
3. *What is UDP port 53 reserved for? 1267? How can a UDP packet be recognised?*
4. *Comment on the value of the TTL fields.*
5. *Comment on the UDP checksum.*

Exercise 1.3 *The following packets are then observed.*

1. *What has happened?*
2. *What is the length of the TCP payload contained in packets 5 to 9?*
3. *What will the value of the sequence and acknowledgement fields be in the next packet sent by lrcpc3?*

```
ETHER:  ----- Ether Header -----
ETHER:
ETHER:  Packet 4 arrived at 19:03:32.40
ETHER:  Packet size = 60 bytes
ETHER:  Destination = 0:0:c:2:78:36, Cisco
ETHER:  Source      = 0:0:c0:b8:c2:8d, Western Digital
ETHER:  Ethertype = 0800 (IP)
ETHER:
IP:   ----- IP Header -----
IP:
IP:   Version = 4
IP:   Header length = 20 bytes
IP:   Type of service = 0x00
IP:         xxx. .... = 0 (precedence)
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length = 44 bytes
IP:   Identification = 2948
IP:   Flags = 0x0
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live = 64 seconds/hops
IP:   Protocol = 6 (TCP)
IP:   Header checksum = cec2
IP:   Source address = 128.178.156.7, lrcpc3.epfl.ch
IP:   Destination address = 129.132.2.72, ezinfo.ethz.ch
IP:   No options
IP:
TCP:  ----- TCP Header -----
TCP:
TCP:  Source port = 1268
TCP:  Destination port = 23 (TELNET)
TCP:  Sequence number = 2591304273
TCP:  Acknowledgement number = 0
TCP:  Data offset = 24 bytes
TCP:  Flags = 0x02
TCP:        ..0. .... = No urgent pointer
TCP:        ...0 .... = No acknowledgement
TCP:        .... 0... = No push
TCP:        .... .0.. = No reset
TCP:        .... ..1. = Syn
TCP:        .... ...0 = No Fin
```

```
TCP:  Window = 512
TCP:  Checksum = 0x2fc4
TCP:  Urgent pointer = 0
TCP:  Options: (4 bytes)
TCP:    - Maximum segment size = 448 bytes
TCP:
TELNET:  ----- TELNET:   -----
TELNET:
TELNET:  ""
TELNET:

ETHER:  ----- Ether Header -----
ETHER:
ETHER:  Packet 5 arrived at 19:03:32.50
ETHER:  Packet size = 60 bytes
ETHER:  Destination = 0:0:c0:b8:c2:8d, Western Digital
ETHER:  Source      = 0:0:c:2:78:36, Cisco
ETHER:  Ethertype = 0800 (IP)
ETHER:
IP:   ----- IP Header -----
IP:
IP:   Version = 4
IP:   Header length = 20 bytes
IP:   Type of service = 0x00
IP:         xxx. .... = 0 (precedence)
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length = 40 bytes
IP:   Identification = 33316
IP:   Flags = 0x0
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live = 119 seconds/hops
IP:   Protocol = 6 (TCP)
IP:   Header checksum = 2126
IP:   Source address = 129.132.2.72, ezinfo.ethz.ch
IP:   Destination address = 128.178.156.7, lrcpc3.epfl.ch
IP:   No options
IP:
TCP:  ----- TCP Header -----
TCP:
TCP:  Source port = 23
TCP:  Destination port = 1268
TCP:  Sequence number = 2068544000
TCP:  Acknowledgement number = 2591304274
TCP:  Data offset = 20 bytes
```

```
TCP:  Flags = 0x12
TCP:        ..0. .... = No urgent pointer
TCP:        ...1 .... = Acknowledgement
TCP:        .... 0... = No push
TCP:        .... .0.. = No reset
TCP:        .... ..1. = Syn
TCP:        .... ...0 = No Fin
TCP:  Window = 3000
TCP:  Checksum = 0x4477
TCP:  Urgent pointer = 0
TCP:  No options
TCP:
TELNET:  ----- TELNET:   -----
TELNET:  ""
TELNET:

ETHER:  ----- Ether Header -----
ETHER:
ETHER:  Packet 6 arrived at 19:03:32.50
ETHER:  Packet size = 60 bytes
ETHER:  Destination = 0:0:c:2:78:36, Cisco
ETHER:  Source      = 0:0:c0:b8:c2:8d, Western Digital
ETHER:  Ethertype = 0800 (IP)
ETHER:
IP:   ----- IP Header -----
IP:
IP:   Version = 4
IP:   Header length = 20 bytes
IP:   Type of service = 0x00
IP:         xxx. .... = 0 (precedence)
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length = 40 bytes
IP:   Identification = 2949
IP:   Flags = 0x0
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live = 64 seconds/hops
IP:   Protocol = 6 (TCP)
IP:   Header checksum = cec5
IP:   Source address = 128.178.156.7, lrcpc3.epfl.ch
IP:   Destination address = 129.132.2.72, ezinfo.ethz.ch
IP:   No options
IP:
TCP:  ----- TCP Header -----
TCP:
```

```
TCP:  Source port = 1268
TCP:  Destination port = 23 (TELNET)
TCP:  Sequence number = 2591304274
TCP:  Acknowledgement number = 2068544001
TCP:  Data offset = 20 bytes
TCP:  Flags = 0x10
TCP:        ..0. .... = No urgent pointer
TCP:        ...1 .... = Acknowledgement
TCP:        .... 0... = No push
TCP:        .... .0.. = No reset
TCP:        .... ..0. = No Syn
TCP:        .... ...0 = No Fin
TCP:  Window = 30619
TCP:  Checksum = 0xd894
TCP:  Urgent pointer = 0
TCP:  No options
TCP:
TELNET:  ----- TELNET:   -----
TELNET:
TELNET:  ""
TELNET:

ETHER:  ----- Ether Header -----
ETHER:
ETHER:  Packet 7 arrived at 19:03:32.56
ETHER:  Packet size = 78 bytes
ETHER:  Destination = 0:0:c:2:78:36, Cisco
ETHER:  Source      = 0:0:c0:b8:c2:8d, Western Digital
ETHER:  Ethertype = 0800 (IP)
ETHER:
IP:   ----- IP Header -----
IP:
IP:   Version = 4
IP:   Header length = 20 bytes
IP:   Type of service = 0x00
IP:         xxx. .... = 0 (precedence)
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length = 64 bytes
IP:   Identification = 2950
IP:   Flags = 0x0
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live = 64 seconds/hops
IP:   Protocol = 6 (TCP)
IP:   Header checksum = ceac
```

```
IP:   Source address = 128.178.156.7, lrcpc3.epfl.ch
IP:   Destination address = 129.132.2.72, ezinfo.ethz.ch
IP:   No options
IP:
TCP:  ----- TCP Header -----
TCP:
TCP:  Source port = 1268
TCP:  Destination port = 23 (TELNET)
TCP:  Sequence number = 2591304274
TCP:  Acknowledgement number = 2068544001
TCP:  Data offset = 20 bytes
TCP:  Flags = 0x18
TCP:        ..0. .... = No urgent pointer
TCP:        ...1 .... = Acknowledgement
TCP:        .... 1... = Push
TCP:        .... .0.. = No reset
TCP:        .... ..0. = No Syn
TCP:        .... ...0 = No Fin
TCP:  Window = 30719
TCP:  Checksum = 0x7ebf
TCP:  Urgent pointer = 0
TCP:  No options
TCP:
TELNET:  ----- TELNET:   -----
TELNET:
TELNET:  ""
TELNET:


ETHER:  ----- Ether Header -----
ETHER:
ETHER:  Packet 8 arrived at 19:03:32.67
ETHER:  Packet size = 60 bytes
ETHER:  Destination = 0:0:c0:b8:c2:8d, Western Digital
ETHER:  Source      = 0:0:c:2:78:36, Cisco
ETHER:  Ethertype = 0800 (IP)
ETHER:
IP:   ----- IP Header -----
IP:
IP:   Version = 4
IP:   Header length = 20 bytes
IP:   Type of service = 0x00
IP:         xxx. .... = 0 (precedence)
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length = 40 bytes
IP:   Identification = 33319
```

```
IP:   Flags = 0x0
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live = 119 seconds/hops
IP:   Protocol = 6 (TCP)
IP:   Header checksum = 2123
IP:   Source address = 129.132.2.72, ezinfo.ethz.ch
IP:   Destination address = 128.178.156.7, lrcpc3.epfl.ch
IP:   No options
IP:
TCP:  ----- TCP Header -----
TCP:
TCP:  Source port = 23
TCP:  Destination port = 1268
TCP:  Sequence number = 2068544001
TCP:  Acknowledgement number = 2591304298
TCP:  Data offset = 20 bytes
TCP:  Flags = 0x10
TCP:        ..0. .... = No urgent pointer
TCP:        ...1 .... = Acknowledgement
TCP:        .... 0... = No push
TCP:        .... .0.. = No reset
TCP:        .... ..0. = No Syn
TCP:        .... ...0 = No Fin
TCP:  Window = 2976
TCP:  Checksum = 0x4478
TCP:  Urgent pointer = 0
TCP:  No options
TCP:
TELNET:  ----- TELNET:   -----
TELNET:
TELNET:  ""
TELNET:


ETHER:  ----- Ether Header -----
ETHER:
ETHER:  Packet 9 arrived at 19:03:32.72
ETHER:  Packet size = 84 bytes
ETHER:  Destination = 0:0:c0:b8:c2:8d, Western Digital
ETHER:  Source      = 0:0:c:2:78:36, Cisco
ETHER:  Ethertype = 0800 (IP)
ETHER:
IP:   ----- IP Header -----
IP:
IP:   Version = 4
IP:   Header length = 20 bytes
```

```
IP:   Type of service = 0x00
IP:         xxx. .... = 0 (precedence)
IP:         ...0 .... = normal delay
IP:         .... 0... = normal throughput
IP:         .... .0.. = normal reliability
IP:   Total length = 70 bytes
IP:   Identification = 33322
IP:   Flags = 0x0
IP:         .0.. .... = may fragment
IP:         ..0. .... = last fragment
IP:   Fragment offset = 0 bytes
IP:   Time to live = 119 seconds/hops
IP:   Protocol = 6 (TCP)
IP:   Header checksum = 2102
IP:   Source address = 129.132.2.72, ezinfo.ethz.ch
IP:   Destination address = 128.178.156.7, lrcpc3.epfl.ch
IP:   No options
IP:
TCP:  ----- TCP Header -----
TCP:
TCP:  Source port = 23
TCP:  Destination port = 1268
TCP:  Sequence number = 2068544001
TCP:  Acknowledgement number = 2591304298
TCP:  Data offset = 20 bytes
TCP:  Flags = 0x18
TCP:        ..0. .... = No urgent pointer
TCP:        ...1 .... = Acknowledgement
TCP:        .... 1... = Push
TCP:        .... .0.. = No reset
TCP:        .... ..0. = No Syn
TCP:        .... ...0 = No Fin
TCP:  Window = 3000
TCP:  Checksum = 0xb907
TCP:  Urgent pointer = 0
TCP:  No options
TCP:
TELNET:  ----- TELNET:   -----
TELNET:
TELNET:  ""
TELNET:
```

Exercise 1.4 1. *Let us consider an IP packet transmitted on an Ethernet. Is it possible that the destination MAC address in the packet is different from the MAC address of the destination?*

2. *A router receives an IP packet on one of its Ethernet interfaces. How can it determine the Ethernet connected system that just sent the packet?*

3. *How many IP addresses does an IP router have?*
4. *Does a bridge need an IP address?*

⋆ **Exercise 1.5** *At ETHZ the IP addresses are of the form 129.132.x.x and use 6 bits for the host part. Thus the prefix length is 26 bits.*

1. *The prefixes of the* `Globi` *and* `Fritz` *subnets are respectively 129.132.43.128/26 and 129.132.43.192/26. For each of the following addresses, say whether they belong to the* `Globi` *subnet , to the* `Fritz` *subnet, or none.*
   (a) $129.132.43.213$
   (b) $129.132.43.25$
   (c) $129.132.43.150$
2. *What should normally be the subnet mask at every ETHZ host ?*
3. *Assume that an IP host A on the* `Globi` *subnet is not well configured and has a subnet mask equal to 255.255.0.0. Explain what happens when such a host wants to send an IP packet to a destination host B (consider both cases where B is on the* `Globi` *subnet or not).*
4. *We assume that we cannot change the configuration of host A. Propose a solution to the problem.*

⋆ **Exercise 1.6** 1. *Consider the network in Figure 1. Only the systems shown on the figure exist in the network. The box in the middle, labeled "NB" is a multi-function network box, which can be configured either as a router or a bridge. It also runs a web server.*
In this question, we assume that NB is configured to work as a bridge. Figure 1 shows the IP addresses and MAC addresses of all interfaces. The network mask on all machines is $255.255.255.0$.

(a) Are the IP addresses plausible, or would you change anything ? (justify your answer)

(b) Does NB need IP addresses, or could we remove them ? (justify your answer)

(c) We assume that the ARP cache at machine `pc33.sovkom.ch` *is empty. We start a TCPDump somewhere on the LAN between* `pc33.sovkom.ch` *and NB (at the place called "Observation Point").*
Then a user at `pc33.sovkom.ch` *executes a command, as shown below:*

```
pc33# telnet batz.sovkom.ch daytime
Trying 192.168.38.1 ...
Connected to batz.sovkom.ch.
Escape character is '^]'.
Tue Nov 29 14:21:34 2005
Connection closed by foreign host.
pc33#
```

(The user sends one request to the server `batz.sovkom.ch` *using telnet, i.e. using TCP, to destination port 13–the port number reserved for the daytime service, obtains one answer from the server, and the TCP connection is closed.)*
For each of the packets that can be observed, give the values of the following fields:

- *MAC source address*
- *MAC destination address*
- *IP source address*

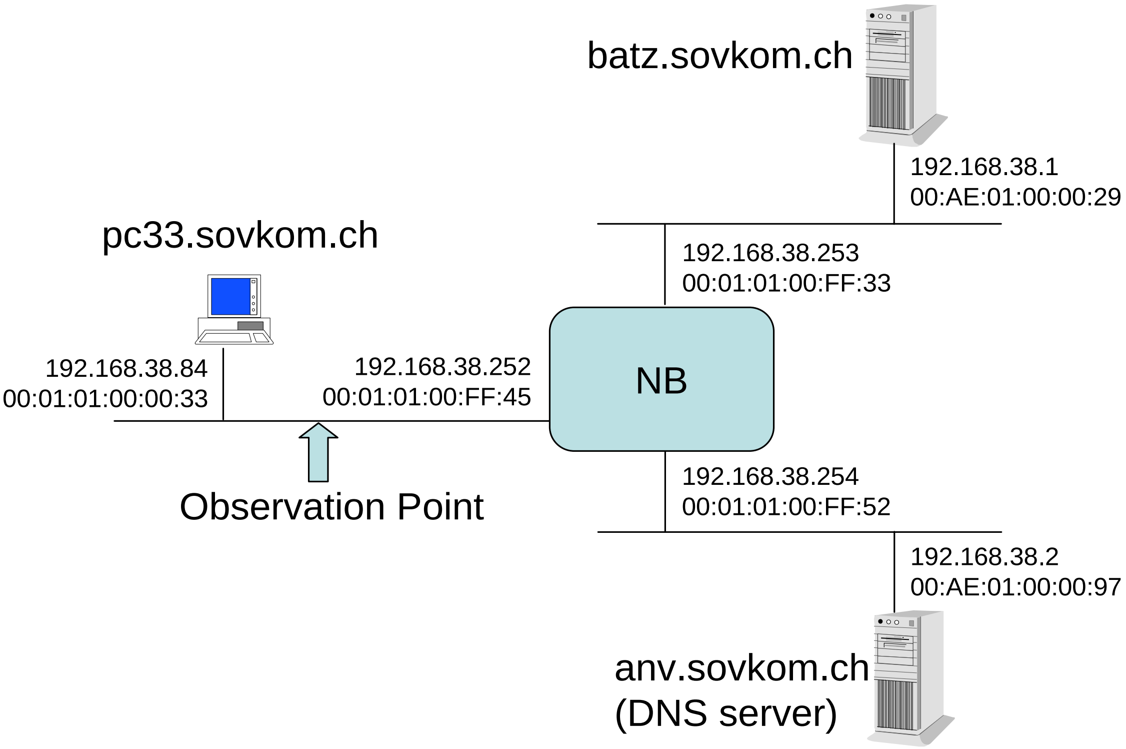

Figure 1: The network of Problem 1, with NB configured as bridge

- *IP destination address*
- *protocol type*
- *if applicable, TCP or UDP source and destination ports*

If some of the values cannot be determined exactly, explain what possible values would be. If two different packets give the same set of values, give it only once.

2. *Now we assume that NB is configured as a router. The addresses are now as shown in Figure 2 Answer the same three questions (a) to (c) as in the previous case.*

⋆ **Exercise 1.7** *Consider the configuration in Figure 3.*

- *PCA is a host running a web browser. WSS is a web server, responding to URL=www.sovkom.an, and with IP address S.*
- *WP is a firewall acting as web proxy (application level gateway). PCA is configured to use WP by default.*
- *ROUT1 and ROUT2 are routers, BRI is a bridge.*
- *All links are Ethernet. All links between those systems are shown on the figure.*
- *MAC and IP addresses are shown on the figure.*
- *We assume that PCA already knows the IP addresses of WP and WSS. Thus, there will be no call to DNS. We also assume that ARP tables are already populated with the correct values, so there will be no ARP messages.*

Assume that PCA sends an HTTP request with target URL = www.sovkom.an in order to transfer one single file. Assume that we do a TCPDump at points 1 to 5, and we capture a copy of all packets that correspond

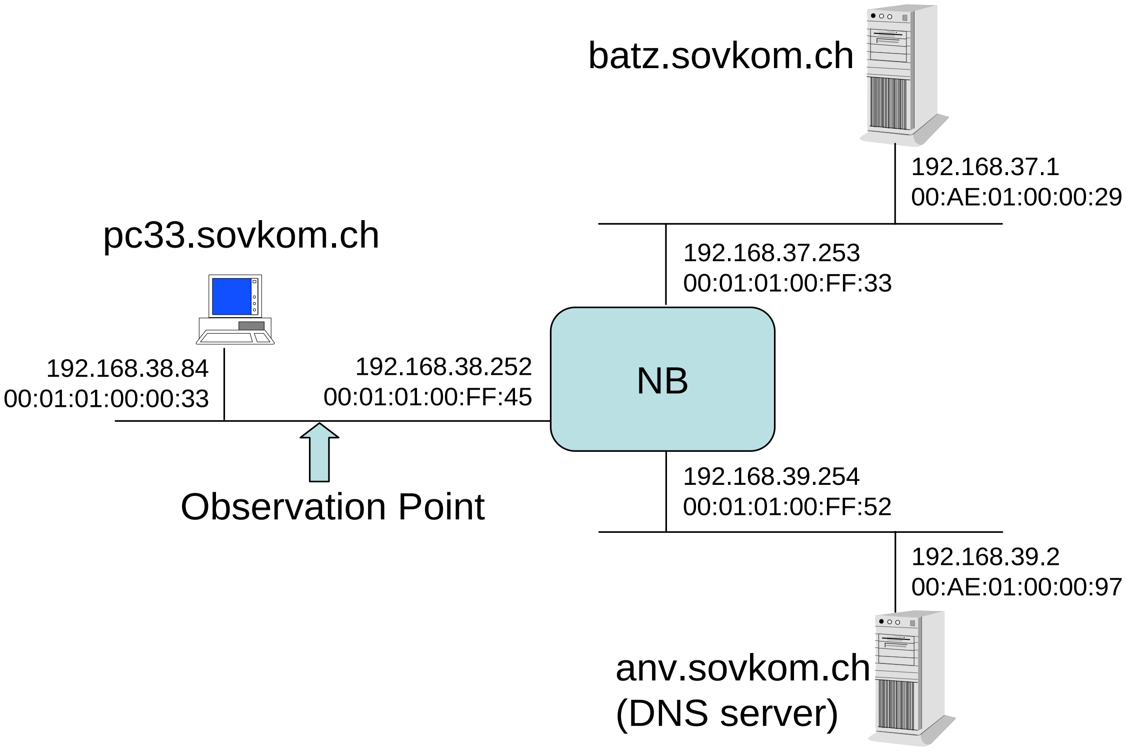

Figure 2: The network of Problem 1, with NB configured as router

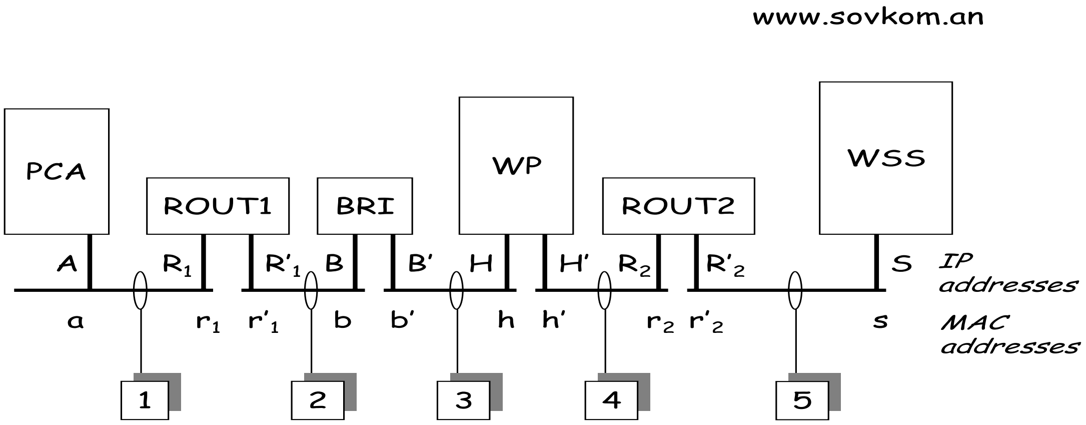

Figure 3: The configuration with web browser on PCA, web proxy WP and web server on WSS. MAC addresses are called $a, r_1, \ldots$. IP addresses are called $A, R_1, \ldots$.

to the transfer of the file and that travel in the direction from PCA to WSS. Give the IP and MAC source and destination addresses that can be read in the packets. Give the solution in the table below.

	MAC source	*MAC destination*	*IP source*	*IP destination*
at 1				
at 2				
at 3				
at 4				
at 5				

Exercise 1.8 *The following trace was captured with TCPDump. It was generated by a Web session. The S, P, and F letters indicate the corresponding TCP flags. Starting with the second packets in each direction, Sequence and Acknowldegement numbers are given by their offset from the values in the first packets. The notation 1:449(448) means that the packet sequence number is 1, it carries 448 bytes of data, and the last byte has sequence number 449 -1.*

1. *Explain the use of the flags.*
2. *Some packets have been mis-ordered. At which lines is that visible ?*
3. *Which lines are retransmissions ?*
4. *Show the states (from the slide "TCP Finite State Machine") for each of the two ends of the connection, for the first and last 10 lines of the trace, after the packet has been received.*

```
15:07:24.544104 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: S
3695173790:3695173790(0) win 512 <mss 448>
15:07:24.584105 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: S
1043264000:1043264000(0) ack 3695173791 win 4096
15:07:24.584105 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 1 win 14247
15:07:24.594105 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . 1:449(448) ack 1
win 14335
15:07:24.704107 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . ack 449 win 3648
15:07:24.704107 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: P 449:500(51) ack 1
win 14335
15:07:24.764108 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: P 1:182(181) ack 500
win 4096
15:07:24.764108 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 182 win 14159
15:07:24.804109 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 182:630(448) ack
500 win 4096
15:07:24.804109 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 630 win 14023
15:07:24.814109 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 630:1078(448) ack
500 win 4096
15:07:24.814109 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 1078 win 14023
15:07:24.814109 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 1078:1526(448) ack
500 win 4096
15:07:24.814109 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 1526 win 14023
15:07:24.834109 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 1526:1974(448) ack
500 win 4096
15:07:24.834109 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 1974 win 14023
15:07:24.834109 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 1974:2422(448) ack
500 win 4096
15:07:24.834109 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 2422 win 14023
15:07:24.844109 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 2422:2870(448) ack
500 win 4096
15:07:24.844109 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 2870 win 14023
15:07:24.844109 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 2870:3318(448) ack
500 win 4096
15:07:24.844109 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 3318 win 13711
```

```
15:07:24.844109 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 3318 win 14335
15:07:24.844109 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 3318:3766(448) ack
500 win 4096
15:07:24.844109 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 3766 win 14023
15:07:24.844109 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 3766:4214(448) ack
500 win 4096
15:07:24.844109 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 4214 win 14023
15:07:24.864110 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 4214:4662(448) ack
500 win 4096
15:07:24.864110 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 4662 win 14023
15:07:24.874110 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 4662:5110(448) ack
500 win 4096
15:07:24.874110 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 5110 win 14023
15:07:24.914111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 5110:5558(448) ack
500 win 4096
15:07:24.914111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 5558 win 14023
15:07:24.914111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 5558:6006(448) ack
500 win 4096
15:07:24.914111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 6006 win 14023
15:07:24.914111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 6006:6454(448) ack
500 win 4096
15:07:24.914111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 6454 win 14023
15:07:24.914111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 6454:6902(448) ack
500 win 4096
15:07:24.914111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 6902 win 14023
15:07:24.924111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 6902:7350(448) ack
500 win 4096
15:07:24.924111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 7350 win 14023
15:07:24.924111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 7350:7798(448) ack
500 win 4096
15:07:24.924111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 7798 win 14023
15:07:24.924111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 7798:8246(448) ack
500 win 4096
15:07:24.924111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 8246 win 14023
15:07:24.934111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 8694:9142(448) ack
500 win 4096
15:07:24.934111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 8246 win 14023
15:07:24.944111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 9142:9590(448) ack
500 win 4096
15:07:24.944111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 8246 win 14023
15:07:24.954111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 9590:10038(448)
ack 500 win 4096
15:07:24.954111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 8246 win 14023
15:07:24.954111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 10038:10486(448)
ack 500 win 4096
15:07:24.954111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 8246 win 14023
15:07:24.964111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 10934:11382(448)
ack 500 win 4096
15:07:24.964111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 8246 win 14023
15:07:24.964111 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 11830:12278(448)
ack 500 win 4096
15:07:24.964111 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 8246 win 14023
15:07:24.974112 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 11382:11830(448)
ack 500 win 4096
15:07:24.974112 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 8246 win 14023
15:07:24.984112 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 8246:8694(448) ack
500 win 4096
15:07:24.984112 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 10486 win 11839
15:07:24.984112 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 10486 win 13399
```

```
15:07:26.014129 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 10486:10934(448)
ack 500 win 4096
15:07:26.014129 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 12278 win 13087
15:07:26.014129 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 12278 win 14335
15:07:26.034129 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 12278:12726(448)
ack 500 win 4096
15:07:26.034129 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 12726 win 14023
15:07:26.034129 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 12726:13174(448)
ack 500 win 4096
15:07:26.034129 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 13174 win 14023
15:07:26.074130 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 13174:13622(448)
ack 500 win 4096
15:07:26.074130 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 13622 win 14023
15:07:26.074130 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 13622:14070(448)
ack 500 win 4096
15:07:26.074130 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 14070 win 14023
15:07:26.084130 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 14070:14518(448)
ack 500 win 4096
15:07:26.084130 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 14518 win 14023
15:07:26.114131 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 14518:14966(448)
ack 500 win 4096
15:07:26.114131 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 14966 win 14023
15:07:26.134131 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 14966:15414(448)
ack 500 win 4096
15:07:26.134131 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 15414 win 14023
15:07:26.144131 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 15414:15862(448)
ack 500 win 4096
15:07:26.144131 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 15862 win 14023
15:07:26.144131 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 15862:16310(448)
ack 500 win 4096
15:07:26.144131 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 16310 win 14023
15:07:26.144131 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 16310:16758(448)
ack 500 win 4096
15:07:26.144131 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 16758 win 14023
15:07:26.154131 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 16758:17206(448)
ack 500 win 4096
15:07:26.154131 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 17206 win 14023
15:07:26.164132 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 17206:17654(448)
ack 500 win 4096
15:07:26.164132 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 17654 win 14023
15:07:26.164132 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 17654:18102(448)
ack 500 win 4096
15:07:26.164132 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 18102 win 14023
15:07:26.164132 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 18102:18550(448)
ack 500 win 4096
15:07:26.164132 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 18550 win 14023
15:07:26.164132 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 18550:18998(448)
ack 500 win 4096
15:07:26.164132 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 18998 win 14023
15:07:26.174132 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 18998:19446(448)
ack 500 win 4096
15:07:26.174132 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 19446 win 14023
15:07:26.184132 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 19446:19894(448)
ack 500 win 4096
15:07:26.184132 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 19894 win 14023
15:07:26.194132 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 19894:20342(448)
ack 500 win 4096
15:07:26.194132 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 20342 win 14023
15:07:26.194132 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 20342:20790(448)
```

```
ack 500 win 4096
15:07:26.194132 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 20790 win 14023
15:07:26.204132 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 20790:21238(448)
ack 500 win 4096
15:07:26.204132 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 21238 win 14023
15:07:26.234133 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 21238:21686(448)
ack 500 win 4096
15:07:26.234133 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 21686 win 14023
15:07:26.234133 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 21686:22134(448)
ack 500 win 4096
15:07:26.234133 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 22134 win 14023
15:07:26.234133 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 22134:22582(448)
ack 500 win 4096
15:07:26.234133 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 22582 win 14023
15:07:26.234133 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 22582:23030(448)
ack 500 win 4096
15:07:26.234133 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 23030 win 14023
15:07:26.234133 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 23030:23478(448)
ack 500 win 4096
15:07:26.234133 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 23478 win 14023
15:07:26.244133 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 23478:23926(448)
ack 500 win 4096
15:07:26.244133 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 23926 win 14023
15:07:26.284134 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 24374:24822(448)
ack 500 win 4096
15:07:26.284134 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 23926 win 14023
15:07:26.284134 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 24822:25270(448)
ack 500 win 4096
15:07:26.284134 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 23926 win 14023
15:07:26.284134 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: FP 25270:25653(383)
ack 500 win 4096
15:07:26.284134 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 23926 win 14023
15:07:27.044146 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . 23926:24374(448)
ack 500 win 4096
15:07:27.044146 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 25654 win 13119
15:07:27.044146 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: . ack 25654 win 14335
15:07:27.064147 lrcpc8.epfl.ch.1071 > ezinfo.ethz.ch.www: F 500:500(0) ack
25654 win 14335
15:07:27.104147 ezinfo.ethz.ch.www > lrcpc8.epfl.ch.1071: . ack 501 win 409
```

Exercise 1.9 *1. Consider the transparency ”Nagle’s Algorithm: Example”. Assume that the packet at line 4 is lost in the network. Give a possible continuation of the message chart.*

2. Assume Nagle’s algorithm is disabled for a given connection. Is it possible that some data written by the application is still delayed ? Prove your answer.

Exercise 1.10 *Quiz*

1. *true* □ *false* □ *When a multiport repeater has some bits to send on a half-duplex Ethernet interface, it should first wait until the channel is idle.*
2. *true* □ *false* □ *When a bridge sends a packet towards the final destination over a full duplex Ethernet interface, it should put as destination MAC address the MAC address of the next hop.*

3. *true □ false □ When a bridge has a packet ready to send on a full-duplex Ethernet port, it listens to the medium and waits until the medium is idle.*
4. *true □ false □ Bridges are said to be "multiprotocol" because a bridged network works independently of network layer protocols such as IPv4 or IPv6.*
5. *true □ false □ A bridge is an intermediate system for layer 2.*
6. *true □ false □ Assume host A sends an IP packet to host B via bridge X, and assume all three systems are on the same bridged network. Then the destination MAC address in the packet sent by A is the MAC address of X.*
7. *true □ false □ On a full duplex Ethernet link, there is no CSMA/CD protocol.*
8. *true □ false □ With an Ethernet switch, there is one collision domain per port.*
9. *true □ false □ A multiport repeater separates collision domains.*
10. *true □ false □ When a bridge has a packet ready to send on a half-duplex Ethernet port, it listens to the medium and waits until the medium is idle.*
11. *true □ false □ In a bridged LAN with more than one bridge and with redundant paths, packet sequence is not guaranteed.*
12. *true □ false □ Assume hosts A and B are on the same bridged LAN, with one bridge X. When host A sends a packet to host B, the source MAC address is that of A, and the destination MAC address is that of the bridge*
13. *true □ false □ A router is an intermediate system for layer 3.*
14. *true □ false □ Ethernet bridges do not use IP addresses when deciding where to send a packet.*
15. *true □ false □ If an IP host A receives an IP packet with TTL=255, then A can conclude that the source of the packet is on-link.*
16. *true □ false □ If host A at EPFL wants to send an IP packet to host B at ETHZ, and if A's ARP cache is empty, then A sends an ARP request in order to determine the IP address of the next hop router.*
17. *true □ false □ Assume A and B are two IPv4 hosts, and that the hosts are on Ethernet. If A and B have the same network mask and the same network prefix, then when A sends a packet to B, the packet still contains an IP destination address, equal to the IP address of B.*
18. *true □ false □ When an IP router between two Ethernet segments forwards an IP packet, it does not modify the destination IP address.*
19. *true □ false □ Assume that host A has an IP packet to send to host B, and that the two hosts are on two Ethernet segments separated by a bridge BR. Assume the ARP table at A is empty. Host A will send an ARP packet in order to find the MAC address of the bridge BR.*
20. *true □ false □ Assume A and B are two IPv4 hosts, and that the hosts are on Ethernet. If A and B have the same network mask and the same network prefix; if A has no entry in its ARP, then before sending a packet to B, A sends an ARP request with target IP address = IP address of B.*
21. *true □ false □ The route indicated by* `traceroute` *may not be the real one because parallel paths may exist in the Internet.*
22. *true □ false □ In an intranet with more than one router, packet sequence is guaranteed by means of the TTL field.*
23. *true □ false □ When an IP router between two Ethernet segments forwards an IP packet, it does not modify the destination MAC address.*
24. *true □ false □ Assume A and B are two IPv4 hosts on the EPFL network. Assume that host A is configured by error with a network mask equal to 255.255.0.0. When A sends a packet to another EPFL host B, if the ARP cache at A is empty, then A will send an ARP packet in order to find the MAC address of B.*
25. *true □ false □ If there are some errors in the routing tables at some routers, then, with IPv4,*

it is possible that a packet loops for ever.

26. *true □ false □ When a router sends a packet towards the final destination over a full duplex Ethernet interface, it should put as destination MAC address the MAC address of the next hop.*
27. *true □ false □ The subnet mask is used by a host or a router in order to know whether it belongs to the same subnet as a machine identified by some IP address.*
28. *true □ false □ When an application receives a block of data from TCP, the application knows that the data was sent as one message by the source.*
29. *true □ false □ Assume host A sends data to host B using TCP. In some cases, it may happen that two blocks of data generated by the application at A are grouped by TCP into one single IP datagram.*
30. *true □ false □ Assume host A sends data to host B using a TCP socket. If A writes three blocks of data into the TCP socket, then there will be three packets sent to B.*
31. *true □ false □ It is possible for a UDP source A to send data to a destination process P_1 on host B_1, using source port a and destination port b, and at the same time send (different) data to another destination process P_2 on a different host B_2, still using the same source port a and destination port b.*
32. *true □ false □ With TCP, the goal of silly window syndrome avoidance is to avoid that out of sequence data is delivered to the application.*
33. *true □ false □ When an application receives data from UDP, the application knows that the data was sent as one message by the source.*
34. *true □ false □ Assume host A sends data to host B using UDP. In some cases, it may happen that two blocks of data generated by the application at A are grouped by UDP into one single IP datagram.*
35. *true □ false □ With a sliding window protocol and for a constant round trip time, increasing the window size increases the throughput if there is no loss, up to a certain limit.*
36. *true □ false □ With a sliding window protocol, the window size is the maximum amount of unacknowledged data that can be sent by the source.*
37. *true □ false □ Assume host A sends one block of data to host B using UDP. In some cases, it may happen that the blocks of data generated by the application at A is fragmented by the IP layer at A into several IP packets.*

2 Module 2: Dynamic Routing

Exercise 2.1 1. *Why do bridges have to build a spanning tree whereas routers do not ?*

2. *What happens to packets if there is a routing loop with bridges ? with routers ?*
3. *Is it possible for a link-state algorithm to use the Bellman-Ford algorithm ? Why or why not ?*

⋆ **Exercise 2.2** *Consider the network in Figure 4. R1 to R6 are routers. Each of these routers has 3 (external) IP interfaces:*

- *two interfaces, called* backbone interfaces, *connect the router to neighbouring routers; the prefix length for these interfaces is 28 bits.*

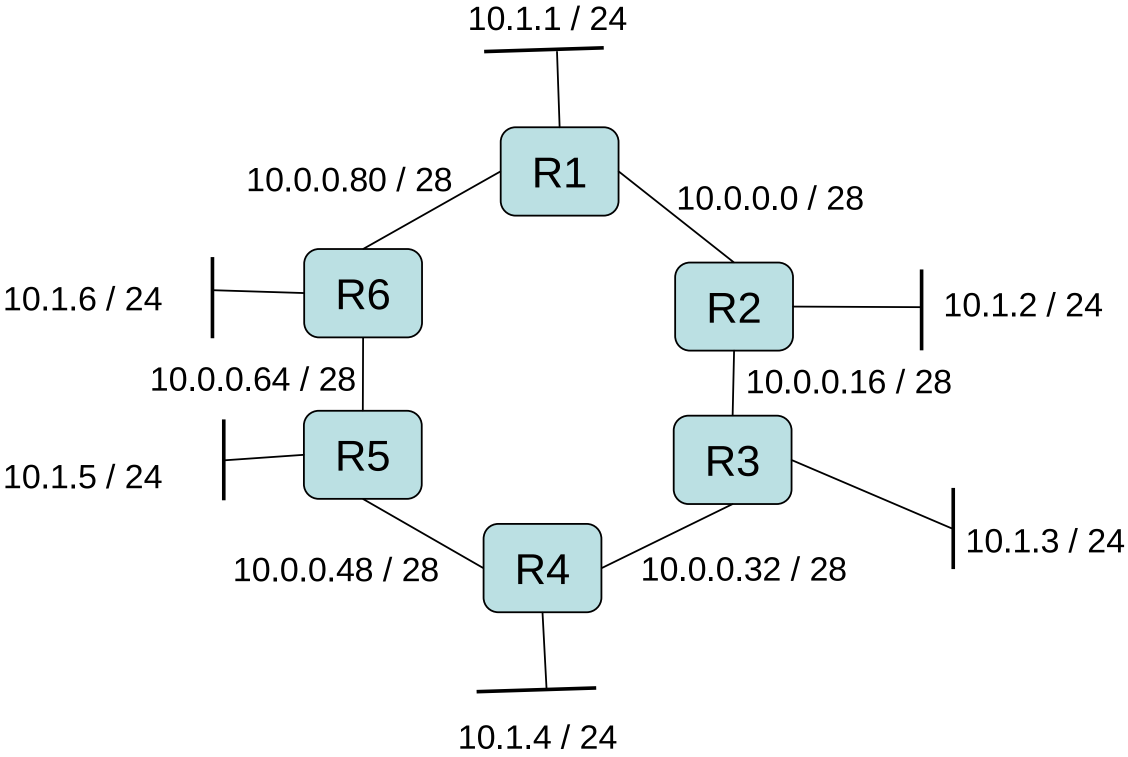

Figure 4: The network in Exercise 2.2.

- *one of them, called* edge interface*, is an interface to a set of hosts; the prefix length for this interface is 24 bits.*

All routers run a distance vector routing protocol such as RIP. The costs of a link between two adjacent routers is equal to 1. The cost from a router to a directly connected network is also equal to 1.

1. *What is the subnet mask at each of the router interfaces shown on the picture ? (give the answer in dotted decimal notation)*
2. *Give the routing table at R1, assuming the routing protocol has converged. Also assume that there is no other network connected to these routers than shown on the picture.*
3. *Assume there exists a host M with IP address 10.0.0.24 and a host A with IP address* $10.1.1.23$*. What are the possible default routers for M and A ? For each combination of cases, what is the path (=sequence of routers) followed by a packet from M to A ?*
4. *Assume now that, on router R2, the edge interface with network prefix 10.1.2/24 is brought down and replaced by a new edge interface, which has now prefix 10.1.7/24. Explain by which mechanisms the other routers will become aware of the change.*
5. *Assume just after this change of configuration, router R2 receives a distance vector from R1, which is based on the values before the change. Explain what will happen, assuming the routing protocol does not implement split horizon. What would happen if the routing protocol* would *implement split horizon ?*
6. *Assume the network has converged after the changes in the previous questions. Assume we do the same manipulation on router R5, with the* same *new prefix (i.e. the edge interface with network prefix 10.1.5/24 is brought down and replaced by a new edge interface, which has now prefix 10.1.7/24, the same prefix as on router R2).*

Normally, this should not be done, since in principle different LANs should have different prefixes. However, this was done by the network managers, maybe by mistake.
Explain the actions that the routing protocol will take, and give the routing table entries at all routers that, after convergence, have changed.

7. *Assume there is a host B2 connected to router R2's edge interface, with IP address 10.1.7.2 and a host B5 connected to router R5's edge interface, with IP address 10.1.7.5. Assume host A has a packet to send to B2 and a packet to send to B5. What is the path followed by each of these packets ? What happens at the last router on the path, in both cases ?*

⋆ **Exercise 2.3** *Consider the network in Figure 5. It represents a small corporate network. The IP addresses are shown explicitly; M1 to M15 mean MAC addresses. B1, B2 and B3 are bridges; R1, R2 and R3 are routers.*

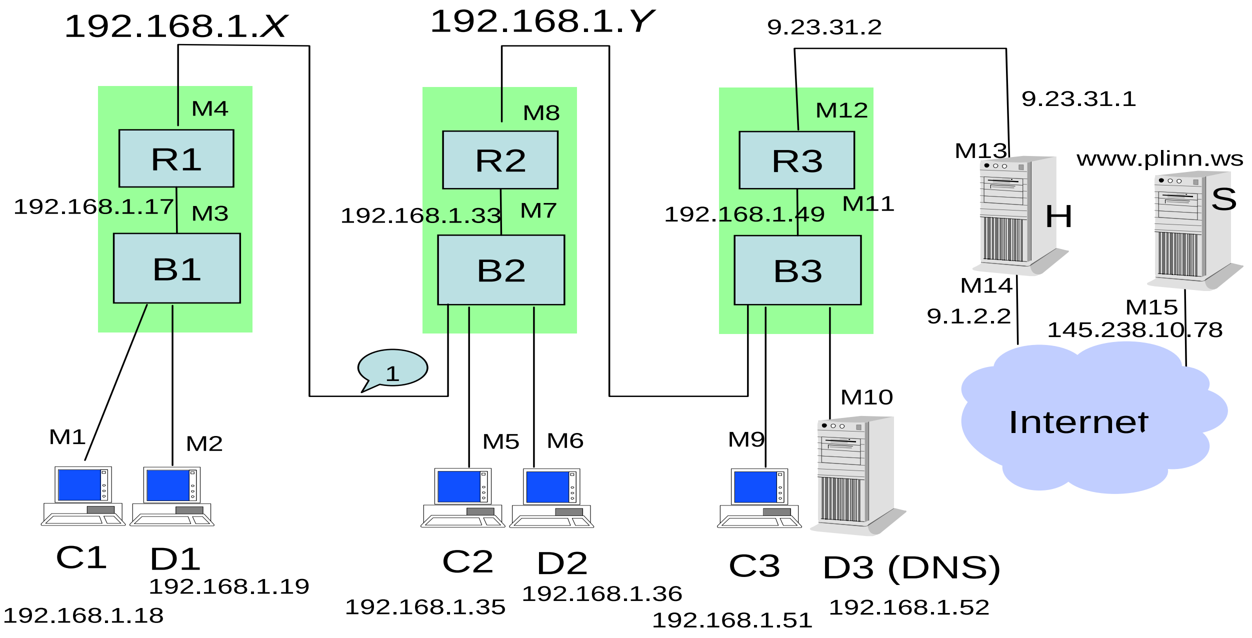

Figure 5: A small corporate network (exercise 2).

D3 is the DNS server for this network. The machines C1, D1, C2, D2, and C3 are configured with DNS server address = 192.168.1.52.

The network is connected to the Internet only by means of a web proxy (the machine H is an application layer gateway).

All interfaces that have IP addresses of the form 192.168.x.y are configured with netmask = 255.255.255.240.

The default gateway are configured as follows

- *at C1 and D1: 192.168.1.17*
- *at C2 and D2: 192.168.1.33*
- *at C3 and D3: 192.168.1.49*

1. *Give a possible value for the X in the IP address of the interface M4 of router R1 (i.e. give a possible value for the address marked* 192.168.1.X *on the figure). Justify your answer. Same question for the*

Y in the IP address of the interface M8 of router R2.

2. *We assume that R1, R2 and R3 are manually configured, i.e. they do not run any routing protocol. Put in the table below the routing table entries that need to be written in these three routers. Give only the entries for destination prefixes that are* not *on-link with this router.*

(Manual Configuration)	*Destination prefix*	*Destination mask*	*Next hop*
R1			
R2			
R3			

3. *The user at host C1 uses a web browser to connect to the server www.plinn.ws, which is on the machine marked S on the figure. As a result, the web browser at C1 sends a DNS query to determine the IP address that corresponds to the DNS name www.plinn.ws. A packet sniffer placed at the location labelled 1 on the figure reads the DNS query and its answer. In the table below, mark the values of the fields that are read in these two packets.*

	MAC header		*IP header*			*Transport Protocol header*	
Packet	*Source MAC address*	*Destination MAC address*	*Source IP address*	*Destination IP address*	*Protocol*	*Source Port*	*Destination Port*
Query from C1 to DNS server							
Response from DNS server to C1							

4. *The web browser at C1 has now received the response from the DNS server and sends an HTTP query. Same question as before for the packets that contain the HTTP query sent by C1 and for the resulting response.*

	MAC header		*IP header*			*Transport Protocol header*	
Packet	*Source MAC address*	*Destination MAC address*	*Source IP address*	*Destination IP address*	*Protocol*	*Source Port*	*Destination Port*
HTTP Request from from C1							
Response to C1							

5. *Assume that we change (by mistake) the netmask for the interface M1 of host C1. The new mask value is 255.255.255.0. Will C1 continue to work normally ? Justify your answer.*

6. *Instead of manual configuration as in question 2, routers R1 R2 and R3 use now RIP. After RIP has converged, what are the routing tables at each router ? Give only the entries for destination prefixes that are* not *on-link with this router.*

(RIP, Figure 5)	*Destination prefix*	*Destination mask*	*Next hop*
R1			
R2			
R3			

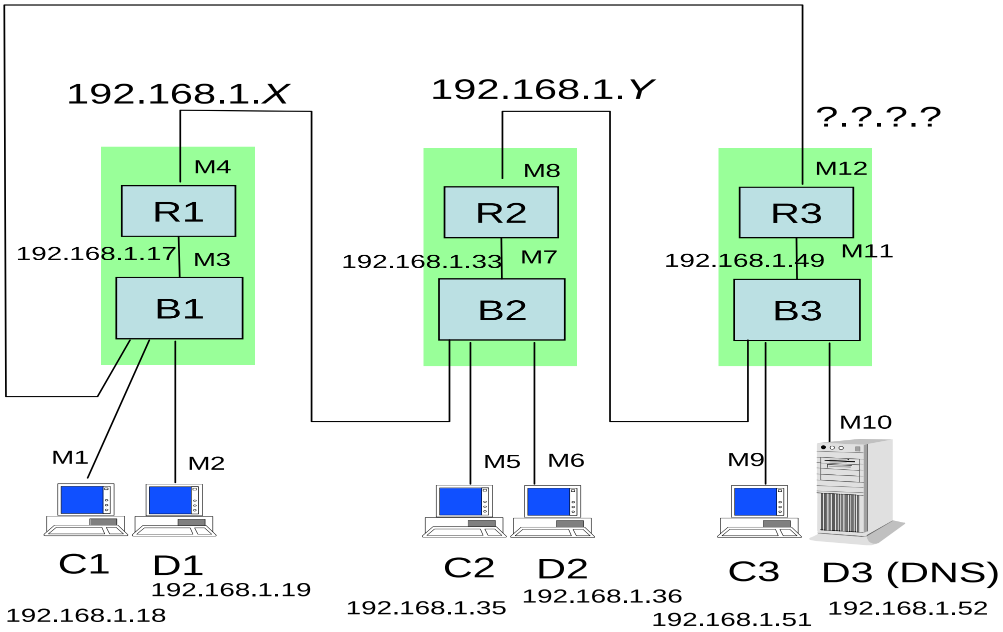

Figure 6: The second network (exercise 2).

7. *We now pull the wire between M12 and M13; then we change the IP address of the interface at M12 and connect M12 to bridge B1; the resulting new configuration is in Figure 6. What IP address and netmask should we give to M12 ?*

Explain what RIP does immediately after the re-connection ?

In the following table, write the routing tables after RIP has stabilized. (As before, give only the *entries for destination prefixes that are* not *on-link with this router.)*

(RIP, Figure 6)	*Destination prefix*	*Destination mask*	*Next hop*
R1			
R2			
R3			

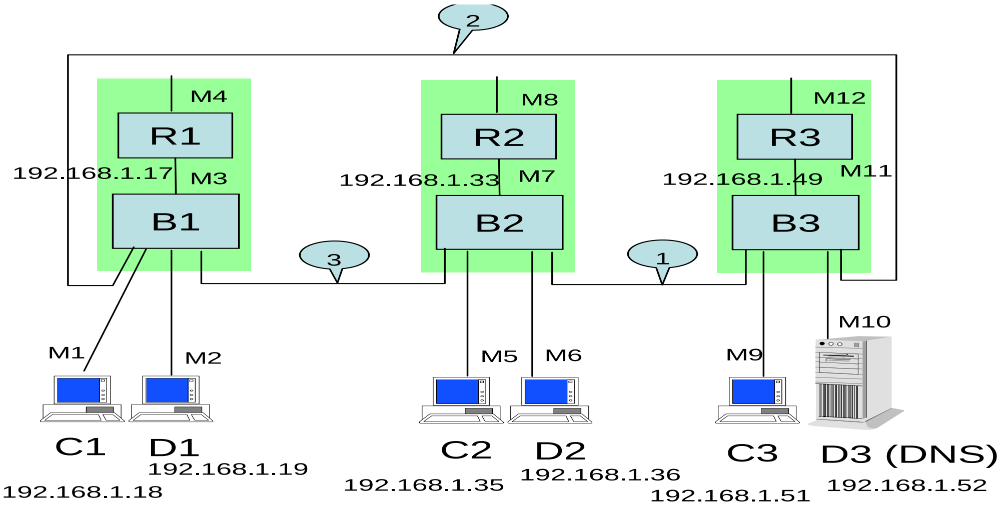

Figure 7: The third network (exercise 2).

8. *We reconfigure the network as shown in Figure 7. The interfaces at M4, M8 and M12 are not used. We change the network mask to 255.255.255.0 on all systems, the IP addresses remain the same. We do a ping from C1 to C2, C2 to C3 and C3 to C1. Packet sniffers are placed at locations labeled 1, 2 and 3 on the figure. In the table below, mark the values of the fields that are read in the ping packets corresponding to each of the ping exchanges if the packet is visible at this location. Consider only the ping packets themselves, not the replies.*

		MAC header		*IP header*		
Sniffing Location	*Ping Packet*	*Source MAC address*	*Destination MAC address*	*Source IP address*	*Destination IP address*	*Protocol*
1	*C1 → C2*					
	C2 → C3					
	C3 → C1					
2	*C1 → C2*					
	C2 → C3					
	C3 → C1					
3	*C1 → C2*					
	C2 → C3					
	C3 → C1					

⋆ **Exercise 2.4** *Consider the network in Figure 8. Domain A [resp. B] is a service provider to domains Z and T [resp. X and Y]. A and B peer over the link shown in the figure. A1, A2, A3 and A4 run BGP and OSPF. B1, B2, B3 and B4 run BGP and RIP. Some of the routers (like A3 and B3) run BGP but do not have external BGP connections. Both domains A and B never redistribute BGP into their IGP. The link costs for OSPF or RIP are all equal to 1.*

The decision process inside all BGP routers in domain A is such that the route selected is, by order of decreasing priority

(1) *the route that has the smallest MED*
(2) *the route that has the shortest IGP distance from this node to the NEXT-HOP of the route*

The decision process inside all BGP routers in domain B uses the reverse order, i.e., it is such that the route selected is, by order of decreasing priority

(1) *the route that has the shortest IGP distance from this node to the NEXT-HOP of the route*
(2) *the route that has the smallest MED*

1. *Here is the list of all BGP announcements that are made over E-BGP:*

```
A1 to B1
     133.29/16 AS-PATH = B  MED = 10
```

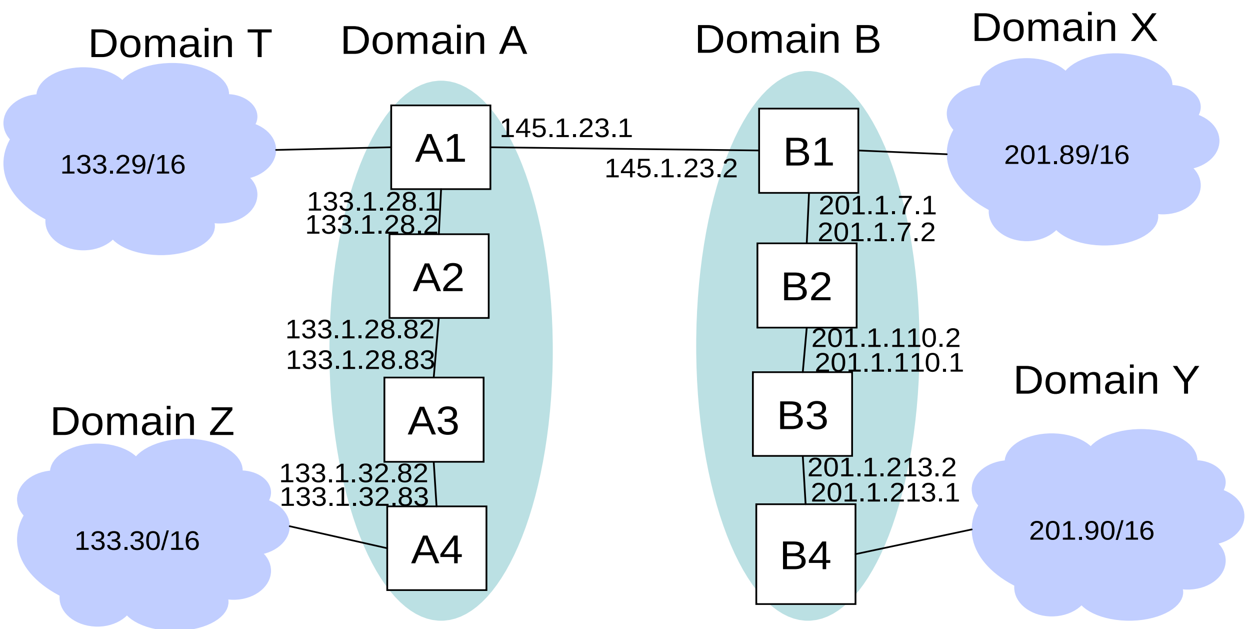

Figure 8: Network for exercise 2.4, step 1.

```
      133.30/16 AS-PATH = B  MED = 40
B1 to A1
      201.89/16 AS-PATH = B  MED = 10
      201.90/16 AS-PATH = B  MED = 40
```

All BGP routers that receive announcements made over E-BGP accept them, and store them in their RIB-in with the corresponding NEXT-HOP attribute. For example, the RIB-in at A1 for the link to B1 is

```
RIB-in at A1
  from B1
     201.89/16 AS-PATH = B  MED = 10 NEXT-HOP = 145.1.23.2
     201.90/16 AS-PATH = B  MED = 40 NEXT-HOP = 145.1.23.2
```

Since A1 receives only one route to 201.89/16 [resp. 201.90/16] we assume that its decision process accepts this route. We assume a similar behavior at B1.

(a) *Will A4 also learn the route* `201.89/16 AS-PATH = B MED = 10` *? If so, by which protocol ?*

(b) *What is the routing table entry (in the forwarding table) at router A3 for destination prefix* `201.89/16`*? at B3 for destination prefix* `133.29/16`*?*

2. *Assume now that an external link is opened between A4 and B4, as shown on Figure 9. The BGP announcements that are made over E-BGP are the same as previously plus the following ones:*

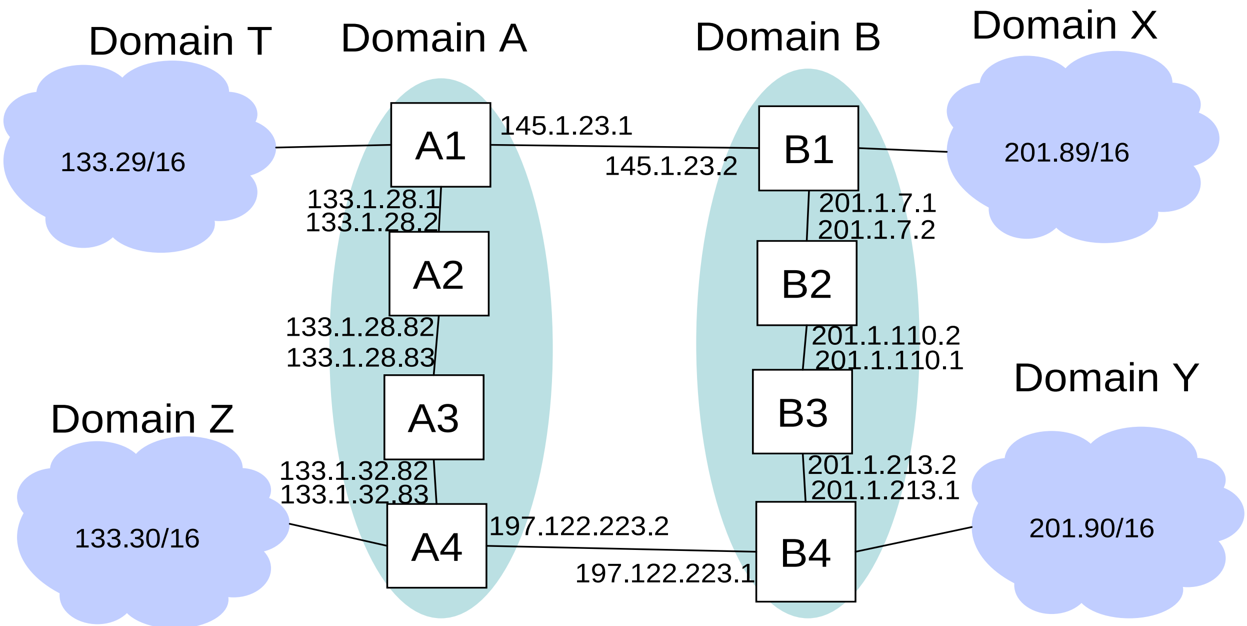

Figure 9: Network for exercise 2.4, step 2.

```
A4 to B4
      133.29/16 AS-PATH = B  MED = 50
      133.30/16 AS-PATH = B  MED = 10
B4 to A4
      201.89/16 AS-PATH = B  MED = 50
      201.90/16 AS-PATH = B  MED = 10
```

(a) *Which route does the decision process at A4 now select toward the networks 201.89/16 and 201.90/16? at B4 toward 133.29/16 and 133.30/16 ? Give the answers in the form*

```
destination prefix / AS-path / next-hop
```

(b) *Does this change the outcome of the decision process at A1 ? at B1 ? (justify your answer)*

(c) *After BGP converges, assume router A3 has a packet with destination address 201.89.1.1 . To which directly connected neighbor does A3 send the packet? By which mechanisms did A3 learn where to send this packet? Same questions for B3 with a packet to send to 133.29.1.1 .*

3. *Assume now that the link between A1 and B1 fails. Will connectivity between domains T and X be affected and how ? Will the BGP routers be able to find new routes from T to X ? Explain what will happen, i.e. what will trigger the events, which messages (information) will be exchanged and between which routers and by which protocol, which entries will change and in which tables and at which routers?*

3 Module 3: Congestion Control

⋆ **Exercise 3.1** *Consider the network illustrated on Figure 10. Source 1 uses links 1 and 2; source 2 uses links 2 and 3: source 3 uses link 3. Each of the links has the same capacity c.*

1. *Assume that the rates x_1, x_2 and x_3 of the three sources are distributed according to max-min fairness. Compute their values.*
2. *Same question assuming that the rates are distributed according to proportional fairness.*

⋆ **Exercise 3.2** 1. *Assume that a TCP sender, called S, does not implement fast retransmit, but does implement slow start and congestion avoidance. Assume that*

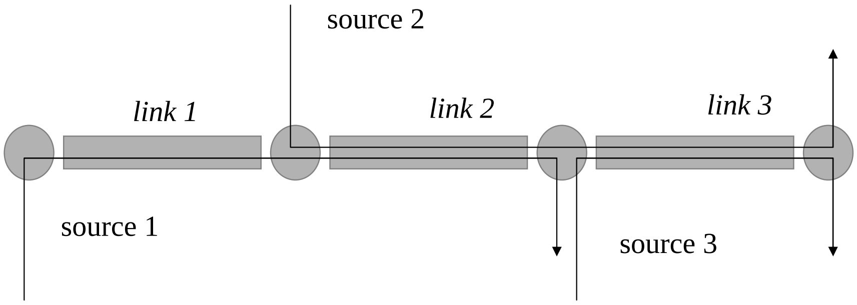

Figure 10: A network with 3 sources.

- *segments $n, n+1, n+2, \ldots, n+10$ are transmitted at times $0, 1, 2, \ldots, 10$ deciseconds (ds)*
- *transmission time per segment is 1 ds*
- *round trip time (2-way propagation plus packet transmission, ack processing and transmission) is fixed and equal to 10 ds*
- *segment n is lost*
- *no other segment or ack is lost*
- *there is no misordering of segments or acks by the network*
- *the retransmission timer for segment n is set to 60 ds, starting at the end of the transmission*
- `cwnd` *=* `twnd` *= 64 segments at time 0*
- `offeredWindow` *= 70 segments all along the exercise*

By which means and at what time will the loss of segment n be detected ?

2. *Immediately after re-transmitting segment n (due to a timeout), S has 3 more segments ready to send (we call them segments $n+11$ to $n+13$). At what time is the ack for segment $n+3$ received (assuming again that there are no losses other than that of segment n) ? For segment $n+13$?*

3. *Same question as before assuming now that fast retransmit and fast recovery are implemented and segments $n+11$ to $n+13$ are available for transmission at time 60.*

$\star$ **Exercise 3.3** *Consider the network illustrated in Figure 11. There are three links and six flows. The flows that use one link are called "short", the other flows are called "long". The short flow number 1 [resp. 2,3] uses link 1, and its rate is called x_1 [resp. x_2, x_3]. The long flow number 1 [resp. 2,3] uses links 1 and 2 [resp. 2 and 3, 3 and 1]. All links have the same capacity c.*

Case 1 *Assume that the rates of the six flows are distributed according to max-min fairness. Compute their values.*

Case 2 *Same question assuming that the rates are distributed according to proportional fairness.*

Case 3 *We assume now that each flow is a TCP connection. Assume that the round trip times of each short flow is equal to the same number T, while the round trip time of each long flow is exactly $2T$ (i.e. most of the round trip time is spent queuing and processing at the link buffers; time spent at destinations is negligible). Assume that:*

- *The loss ratio at the entrance to link i is q, the same for all links, and the same for all flows that use a link. Thus the loss ratio for a flow that uses one link is q*

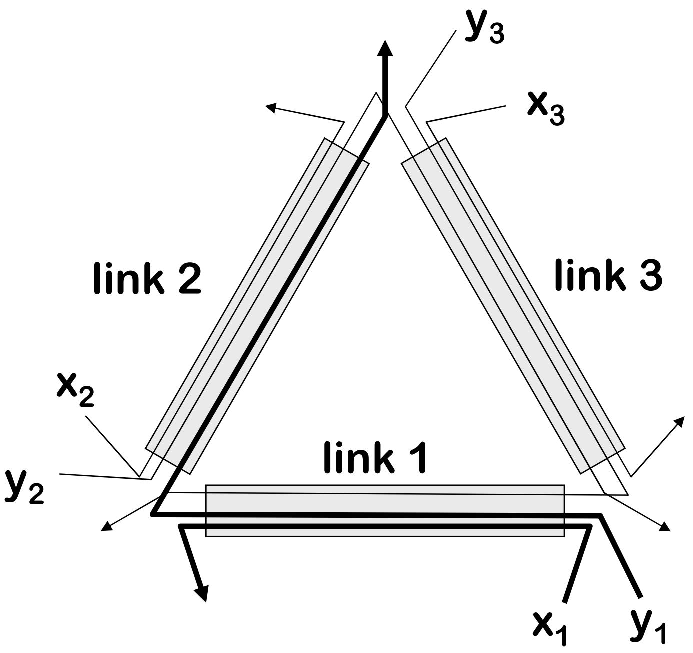

Figure 11: A network with 6 flows. The short flow 1 and the long flow 1 are shown with thick lines.

- *Loss rates are small, so that the loss ratio for a flow that uses two links is considered to be equal to $2q$*
- *All flows have the same TCP segment size*
- *Flows are limited only by the link capacities and overhead is negligible, so we assume that the sum of the rates of flows is equal to link capacity. For example:*

$$x_1 + y_1 + y_3 = c \tag{1}$$

- *TCP acknowledgements are never lost.*

Compute the rates achieved by all flows under these assumptions.

Case 4 *Same question assuming that the round trip times of all six flows are approximately the same (i.e. most of the round trip time is due to processing at the servers, instead of queuing and processing at the link buffers). Compare to the rates found in cases 1 and 2 and explain what you find. Also compare to the rates found in case 3 and explain what you find.*

⋆ **Exercise 3.4** *Consider the scenario in Figure 12.*

A and B each transfer a very large file over a TCP connection with server S. The link rates are:

- *between A and R1: 54 Mb/s in each direction*
- *between B and R1: 100 Mb/s in each direction*
- *between R1 and R2: 6 Mb/s in each direction*
- *between R2 and S: 1000 Mb/s in each direction*

Assume there is no other traffic than these two TCP connections. The RTT for A is 1000 ms, for B it is 200 ms.

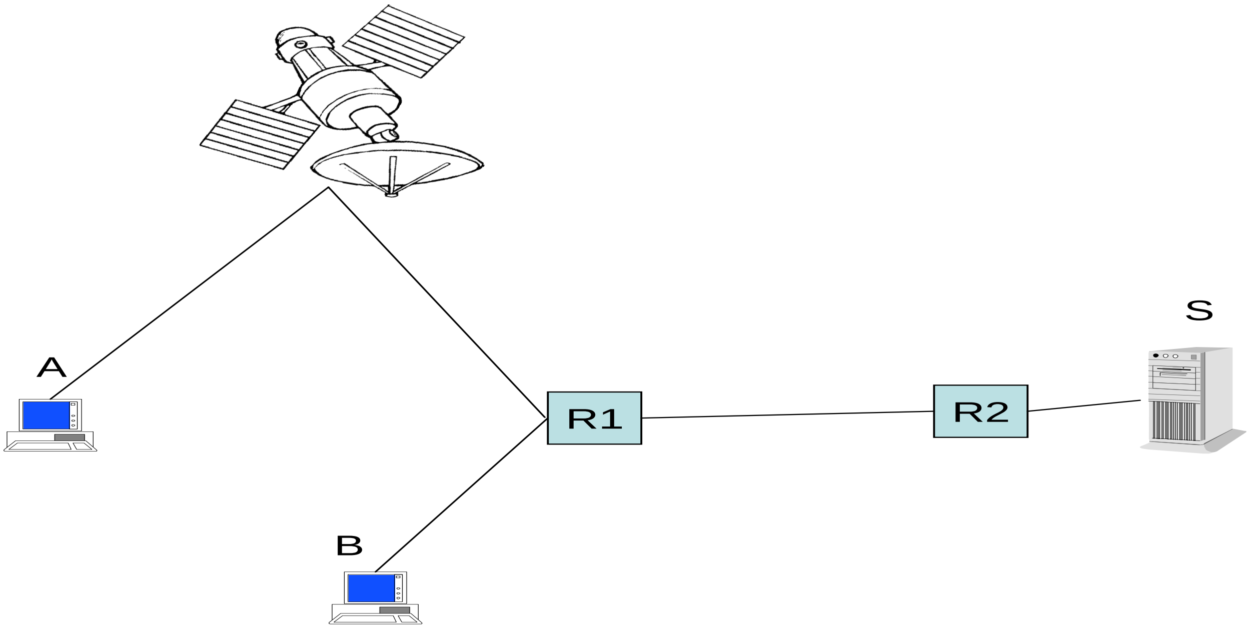

Figure 12: Network for exercise 3.4, step 1.

1. *What would be the throughputs of A and B if they were max-min fair ? if they were proportionally fair ?*

2. *Assume we can neglect all losses except on the link R1-R2. Also assume that the two TCP connections can, together, fully utilize the link R1-R2 (i.e. the sum of their throughputs is 6 Mb/s). What are the throughputs achieved by A and B ?*

3. *Assume now that the satellite link has a high loss ratio, due to transmission errors (FEC is disabled). The loss rate for the connection from A is now assumed to be fixed and equal to 0.01. We assume the RTTs stay the same, and, as before, that the two TCP connections can, together, fully utilize the link R1-R2 (i.e. the sum of their throughputs is 6 Mb/s). What are now the throughputs achieved by A and B ?*

4. *Assume we change the configuration and introduce a transport layer gateway H, as shown on Figure 13.*

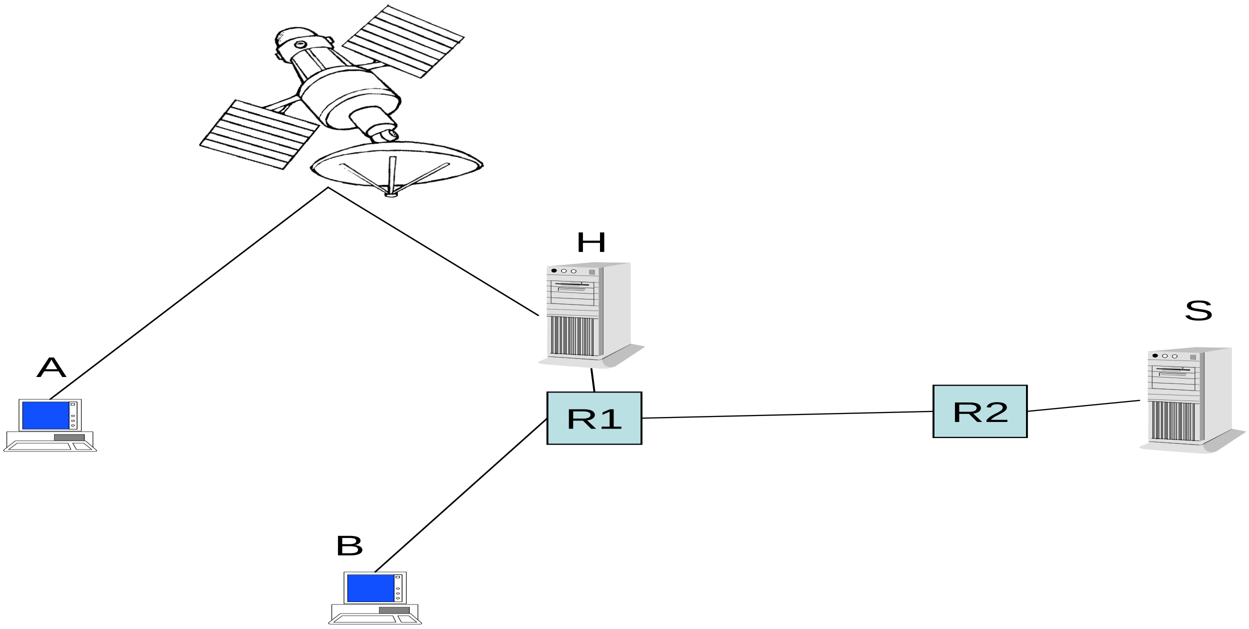

Figure 13: Network for exercise 3.4, step 2.

There is now one TCP connection from A to H, a second one from H to S (and one from B to S as *before). The RTT between H and S is 200 ms, between A and H is 800 ms and between B and S 200 ms as before. We assume FEC is enabled so that there is no loss between A and H. What is the end to end throughput of A and B ? Same question if FEC is disabled, so that the loss rate between A and H is again 0.01.*

Is it useful to use a transport layer gateway in this scenario ?

⋆ **Exercise 3.5** *Assume that we change the code of TCP by modifiying the congestion avoidance phase as follows. For every acknowledgement received, we increase the window size as in slow start. Give the pros and cons of this modification.*

⋆ **Exercise 3.6** *Assume two hosts have a very large file to transfer, and decide to open n parallel TCP connections for the transfer.*

1. *Find one advantage and one drawback, for the user of the two hosts, of having $n > 1$. Justify your answer.*
2. *Find one advantage or one drawback, for the rest of the network, of having $n > 1$. Justify your answer.*

4 Module 4

Exercise 4.1 *Host redirect*

1. *Supposing that if a host receives (due to an error) a packet destined for an IP address that is not that of the host, that it will retransmit this packet towards the destination ("good citizen rule"). Now suppose that a computer with IP address a, following a hardware error, thinks that its Ethernet address is A= "xFF:FF:FF:FF:FF:FF". What happens when computers with IP addresses b,c,... (and corresponding MAC addresses B,C,...) try to send IP datagrams to address a? We assume that all the computers are on the same LAN.*
2. *Among the following rules, which ones seem to you to remove this problem:*
   - *ICMP messages are never sent as a response to datagrams received on a link layer broadcast address;*
   - *IP datagrams received on a layer 2 broadcast address are never forwarded;*
   - *ARP table entries should never map to the link layer broadcast for non-multicast and non broadcast addresses.*

Exercise 4.2 *1. The Total Length field in the IP header is the length of the IP datagram (true/false)*

2. *Replay the "IP Fragmentation (2)" transparency, assuming that the MTU of the network on the right is 250 (instead of 1500).*
3. *Supposing that the lifetime of a packet in transit is TTL seconds at the most and that the reassembly timer is at 120 seconds at the most. What is the maximum number of IP packets per second that a system can send so as to avoid two different packets having the same Identification field, for TTL = 60 and TTL = 255?*
4. *In the fragmentation algorithm, what happens if* `new(fragmentList, P0.(identification, source address), f1))` *returns* `false`*? How can the performance be improved, avoiding creating a list if a fragment has already been refused because it was not possible to create the list at that moment.*
5. *What is the maximum amount of free memory that it is necessary to allocate at the creation of a fragment list?*
6. *Show that the reassembly method shown in the course transparency can lead to a deadlock. How can this be avoided?*
7. *Can the following functions be performed (yes or no):*

```
                 Router   Host
   Segmentation  -------- ------
   Reassembly    -------- ------
```

Exercise 4.3 *1. The first byte of an option of an IP header contains "Option Code"; if it is 1, the options field must be copied into each of the fragments at fragmentation. Modify the algorithm in the "Fragmentation Algorithm" transparency to take this field into account.*

2. *In the modified algorithm, is there a case where the occurrence of copy bit = 1 in an option should be considered as an error?*

Exercise 4.4 *1. A low quality modem link has a bit error rate of 0.0001. We suppose the errors to be independent and equally distributed. What is the probability that a packet of length L is received correctly ? Assume acks are never lost.*

2. *The modem uses SRP error recovery. We assume a very large window, so that the source always has something to send. We also assume that we have data units of long length to send, segmented into IP packets of size $L + K$, where K is the overhead. What is the line utilisation for: $K = 48$ octets (as for TCP over IP over PPP without compression) $L = 296$ octets, $L = 1500$ octets.*

Exercise 4.5 *A set of n web servers is connected to a single router. The n web servers are all replications of the same site. We want to give only one IP address, shared by the set of n web servers, so that remote web clients have to know only one address. Find two different solutions for this problem. Describe each of your two solutions in 5 to 25 lines. Discuss the advantages and drawbacks of each of your solutions in 5 to 25 lines.*

Exercise 4.6 *Explain in 1 to 5 lines each of the concepts below.*

1. *Automatic Tunnel*
2. *Sollicited Node Multicast Address*
3. *DHCP*
4. *Flow Label*
5. *Transition from IPv4 to IPv6*
6. *IPv4 compatible IPv6 address*
7. *IPv4 mapped IPv6 address*

Exercise 4.7

1. *Explain why some of the logical interfaces on a dual stack host attached to Ethernet have an MTU of 1480 instead of 1500*
2. *How does a dual stack host know wether it should talk IPv4 or IPv6 with a distant host ?*
3. *How does a dual stack host know to which router and over which mecanism it should send an IP packet ?*

Exercise 4.8

1. *Explain in at most 15 lines the principles, benefits and limitations of an IPv4-IPv6 interworking unit operating at the network layer.*
2. *Explain in at most 15 lines the principles of an IPv4-IPv6 interworking unit operating at a layer other than the network layer.*

Exercise 4.9

1. *On slide "B: Some Examples: Email", explain which function is in a router and which one is in a host.*
2. *What is a mail exchanger ? Is it a router ?*
3. *Do mail exploders use multicast IP addresses ?*

⋆ **Exercise 4.10** 1. *Consider the intranet illustrated on Figure 14. There are three Ethernet segments at 10 Mb/s, each corresponding to a* `net:subnet` *prefix noted* n_1, n_2 *and* n_3. *Every Ethernet segment is connected to two routers as indicated on the figure. There is no external connection to this intranet. Each Ethernet segment has a number of hosts directly attached to it. The Ethernet segments are shared media, there is no Ethernet switching equipment.*

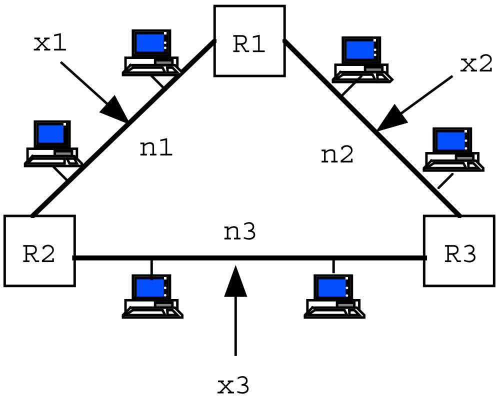

Figure 14: The first network for Problem 4.10

(a) *We assume that the IP routing tables in R1, R2 and R3 are setup in such a way that traffic from subnet* n_i *to a subnet* n_j, *with* $i \neq j$ *goes through exactly one router. How can the hosts and routers be configured to achieve this ? Show all the relevant routing tables. If necessary, introduce additional notation (for addresses and interfaces).*

(b) *We call* x_i *the total traffic generated by all hosts directly attached to segment* i. *We neglect the effect of collisions on one Ethernet and thus assume that the maximum amount of traffic possible on every Ethernet segment is 10 Mb/s.*
We further assume that the destination of traffic originating from subnet i *is uniformly distributed among the three subnets. Thus, for example, the amount of traffic originating from subnet* 1 *which has a destination in subnet* 2 *is* $\frac{x_1}{3}$.
What is the maximum value of the total traffic $x_1 + x_2 + x_3$ *which is possible with these assumptions ?*

2. *We consider the intranet illustrated on Figure 15. There are three Ethernet segments at 10 Mb/s, interconnected by means of three bridges B1, B2 and B3. There is no router and no external connection in this intranet. Each Ethernet segment has a number of hosts directly attached to it. The Ethernet segments are shared media, there is no Ethernet switching equipment apart from the three bridges B1, B2 and B3.*

(a) *How many subnets are there in principle in this intranet ?*

(b) *We call* x_i *the total traffic generated by all hosts directly attached to segment* i. *We neglect the effect of collisions on one Ethernet and thus assume that the maximum amount of traffic possible on every Ethernet segment is 10 Mb/s.*
We further assume that the destination of traffic originating from subnet i *is uniformly distributed among the three subnets. Thus, for example, the amount of traffic originating from subnet* 1 *which has a* final *destination in subnet* 2 *is* $\frac{x_1}{3}$. *We further assume that the bridges have had enough time to learn and build their forwarding tables.*
What is the maximum value of the total traffic $x_1 + x_2 + x_3$ *which is possible with these assumptions ?*

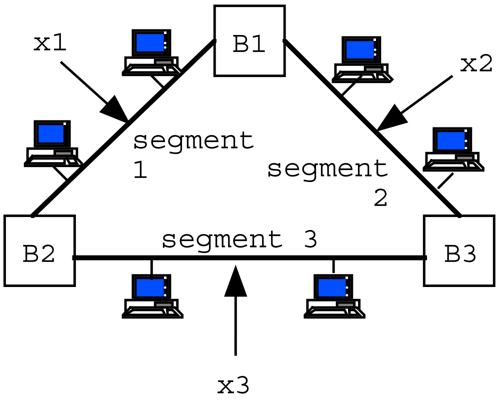

Figure 15: The second network for Problem 4.10

Exercise 4.11 1. *Based on what addresses (MAC addresses, IP addresses, UDP/TCP port numbers) do bridges forward packets ?*

2. *For the rest of this problem, we define a new bridging method (called "structured bridging"). Unlike the bridging method we have seen in the lecture, this method is not transparent. It uses a supplementary field in the Ethernet packet, called "Routing Information" field. The Routing Information field already exists in some versions of Ethernet; it is used by source routing bridges (which we did not see in the lecture). We redefine it here for the purpose of supporting structured bridging. Therefore, structured bridging is incompatible with source routing and with transparent bridging.*

   Consider a collection of Ethernet segments, interconnected at layer 2 by bridges (one Ethernet segment is one collision domain). We assume that the bridges implement structured bridging, defined as follows.

   - *every LAN segment is allocated a 16 bit number, called the 'LAN Segment Identifier' (LSI)*
   - *consider a station A, connected to a LAN segment with LSI equal to L1, and a station B connected to a LAN segment with LSI equal to L2. When A sends a MAC frame to B, it puts in the MAC frame: source address = A, destination address = B, Routing Information : (destination LSI=L2, source LSI=L1).*
   - *bridges have forwarding tables based on LSIs. Namely, for every destination LSI, a bridge knows on which port it has to send the MAC frame. Unless specified otherwise, we assume that the forwarding tables are written manually.*

   Assume station A knows the IP address of station B, but not its own LSI nor B's LSI, nor B's MAC address. Assume also that a host knows its own MAC address but does not know its LSI at initialization. Explain (in at most 15 lines) how station A can communicate with station B over the bridged network using structured bridging (without using any routers).

3. *Propose the principles of a method for the bridges to build the LSI forwarding tables automatically (at most 15 lines)*

Exercise 4.12 *Quiz*

1. *true* □ *false* □ *When an IPv4 system A sends an ICMP packet to the IPv4 system B, it sends it as a UDP packet with destination port number = 1 (the port reserved for ICMP).*

2. *true* □ *false* □ *A filtering router may be configured to discard all IP packets that have protocol type = UDP and source port number = 234.*