Exercises
TCP/IP Networking
With Solutions

Jean-Yves Le Boudec

Fall 2009

1 Module 1: TCP/IP Architecture

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.*

Solution: In the hypothesis that the override timer is bigger than the retransmission timer, we have **repacketization**: at the retransmission timeout the lost bits are retransmitted, together with the bits arrived in the meantime. In this case Nagle's algorithm does not come into play.

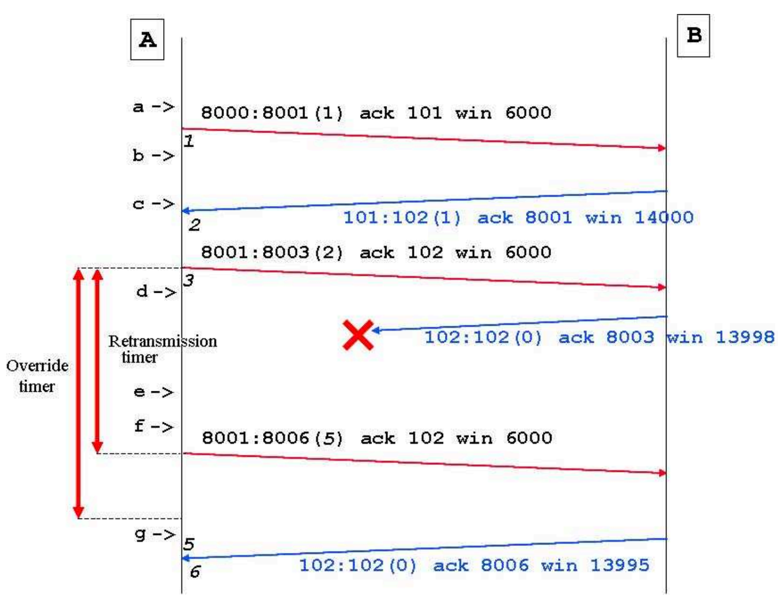

1

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.*

[1] file: 22

Solution: Due to the sliding window mechanism, whenever the available window does not allow the application to send all the data that it produces, all the remaining data is buffered and therefore delayed.

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.*

   Solution: False. The repeater simply repeats bits, even if this causes a collision.

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.*

   Solution: False; the bridge does not modify MAC addresses.

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.*

   Solution: False, there is no CSMA/CD over full duplex Ethernet.

4. *true* ☐ *false* ☐ *Bridges are said to be "multiprotocol" because a bridged network works independently of network layer protocols such as IPv4 or IPv6.*

   Solution: True.

5. *true* ☐ *false* ☐ *A bridge is an intermediate system for layer 2.*

   Solution: True.

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.*

   Solution: False. Bridges on Ethernet are transparent.

7. *true* ☐ *false* ☐ *On a full duplex Ethernet link, there is no CSMA/CD protocol.*

   Solution: True. A full duplex Ethernet link uses Ethernet physical layer but is not a shared medium link.

8. *true* ☐ *false* ☐ *With an Ethernet switch, there is one collision domain per port.*

   Solution: True.

9. *true* ☐ *false* ☐ *A multiport repeater separates collision domains.*

   Solution: False.

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.*

Solution: True, the bridge executes CSMA/CD on all half-duplex ports.

11. *true* ☐ *false* ☐ *In a bridged LAN with more than one bridge and with redundant paths, packet sequence is not guaranteed.*

    Solution: False. Packet sequence is guaranteed by the spanning tree algorithm, which reduces the active topology to a tree.

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*

    Solution: False; Bridges are transparent. The destination address is normally that of B; it may also be the broadcast address, or a multicast address.

13. *true* ☐ *false* ☐ *A router is an intermediate system for layer 3.*

    Solution: True.

14. *true* ☐ *false* ☐ *Ethernet bridges do not use IP addresses when deciding where to send a packet.*

    Solution: True. Bridges do not look at layer 3 information and are therefore said to be multiprotocol

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.*

    Solution: True.

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.*

    Solution: False. The ARP request is to find the MAC 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.*

    Solution: True.

18. *true* ☐ *false* ☐ *When an IP router between two Ethernet segments forwards an IP packet, it does not modify the destination IP address.*

    Solution: true.

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.*

    Solution: False. The bridge is not visible to A. The ARP is to find the MAC address of B.

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.

Solution: True. Comment: if proxy ARP is used, a proxy ARP server may respond with another MAC address than that of B

21. *true* □ *false* □ *The route indicated by* `traceroute` *may not be the real one because parallel paths may exist in the Internet.*

    Solution: True.

22. *true* □ *false* □ *In an intranet with more than one router, packet sequence is guaranteed by means of the TTL field.*

    Solution: False. Packet sequence is not guaranteed with IP.

23. *true* □ *false* □ *When an IP router between two Ethernet segments forwards an IP packet, it does not modify the destination MAC address.*

    Solution: false.

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.*

    Solution: True. This is not the normal configuration, but it will works because in such cases the default router for A will use proxy ARP and respond with its own MAC address

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.*

    Solution: False. The packet is discarded when the TTL fields becomes 1.

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.*

    Solution: True.

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.*

    Solution: True.

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.*

    Solution: False.

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.*

    Solution: True. TCP does its own packetization.

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.*

    Solution: False. TCP does its own packetization. There may be any number of packets, depending on how much data is written by 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.*

    Solution: True.

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.*

    Solution: False.

33. *true* ☐ *false* ☐ *When an application receives data from UDP, the application knows that the data was sent as one message by the source.*

    Solution: True.

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.*

    Solution: False.

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.*

    Solution: True.

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.*

    Solution: True.

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.*

    Solution: True.

2 Module 2: Dynamic Routing

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

Solution: Bridges have to build a *spanning* tree because they forward packets according to MAC addresses which are not structured and they do not detect frames that loop. Routers do not have to build a spanning tree since they forward packets according to IP addresses which are structured and eventually discard packets that loop.

2. *What happens to packets if there is a routing loop with bridges ? with routers ?*

   Solution: Packets loop indefinitely if there is a routing loop with bridges. Packets will eventually be discarded if there is a routing loop with routers because of the TTL field.

3. *Is it possible for a link-state algorithm to use the Bellman-Ford algorithm ? Why or why not ?*

   Solution: The link-state algorithm can use the Bellman-Ford algorithm (static version) for computing the shortest path to all other nodes since the Bellman-Ford algorithm requires only a partial view of the network and the link-state algorithm provides a complete topology view of the network.

⋆ **Exercise 2.3** *Consider the network in Figure 1. 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.*

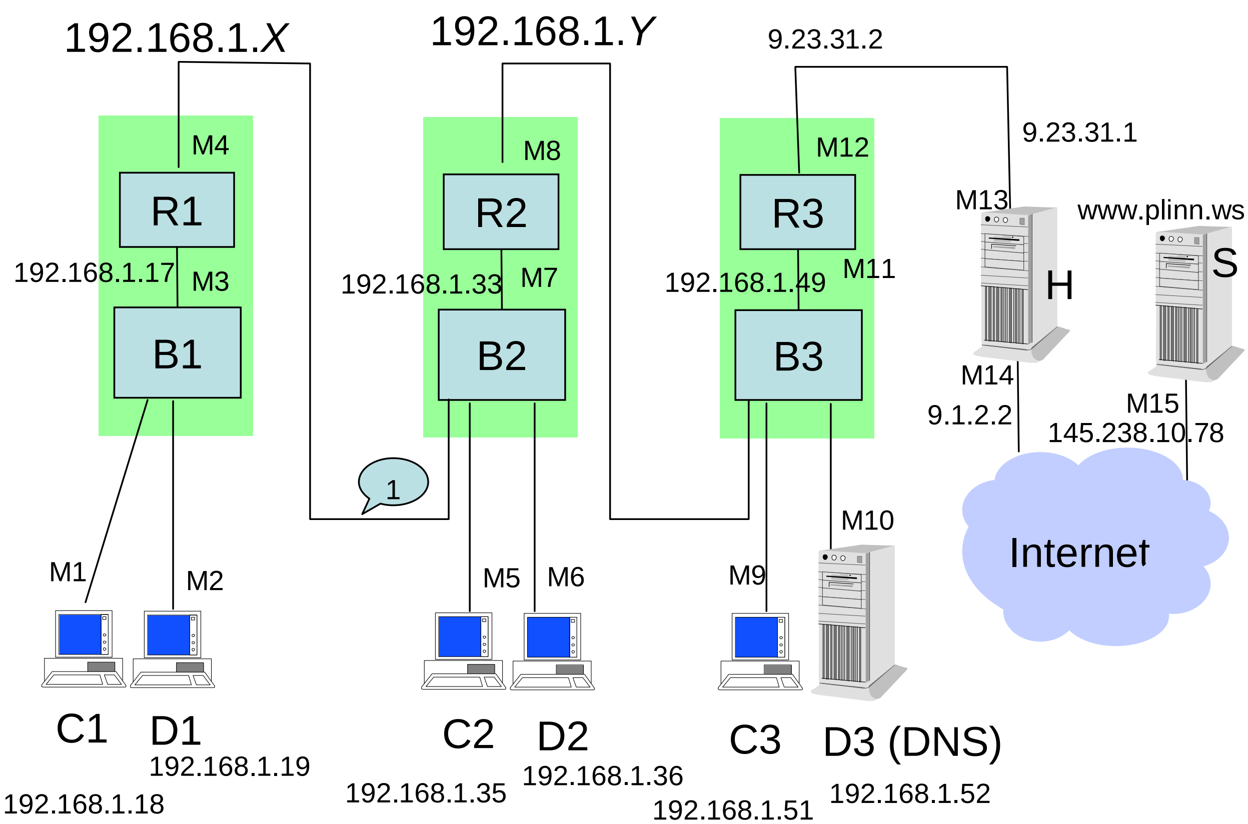

Figure 1: A small corporate network (exercise 2.3).

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 gateways 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.*

   Solution: The M4 interface must belong to the same subnet as C2 and D2. Since the mask is over 28 bits, X must lie in the interval $[33, 46]$ (the host parts $b0000$ [$X = 32$] and $b1111$ [$X = 47$] are not possible). The value must also not be already allocated. A possible value is 34. Similarly, Y must lie in the interval $[49, 62]$; a possible value is 50.

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.*

   Solution:

(Manual Configuration)	*Destination prefix*	*Destination mask*	*Next hop*
R1	*0.0.0.0*	*0.0.0.0*	*192.168.1.33*
R2	*192.168.1.16* *9.23.31.0*	*255.255.255.240* *255.255.255.240*	*192.168.1.34* *192.168.1.49*
R3	*0.0.0.0*	*0.0.0.0*	*192.168.1.50*

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.*

   Solution: The port p needs to be one of the non-reserved ports.

Packet	MAC header: Source MAC address	MAC header: Destination MAC address	IP header: Source IP address	IP header: Destination IP address	IP header: Protocol	Transport Protocol header: Source Port	Transport Protocol header: Destination Port
Query from C1 to DNS server	*M4*	*M7*	*192.168.1.18*	*192.168.1.52*	*UDP*	*p>1024*	*53*
Response from DNS server to C1	*M7*	*M4*	*192.168.1.52*	*192.168.1.18*	*UDP*	*53*	*p*

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	*M4*	*M7*	*192.168.1.18*	*9.23.31.1*	*TCP*	$p>1024$	*80*
Response to C1	*M7*	*M4*	*9.23.31.1*	*192.168.1.18*	*TCP*	*80*	p

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.*

Solution: C1 will continue to work normally if router R1 functions also as an ARP proxy. Otherwise not, because no one will answer to C1's ARP requests for the machines in 192.168.1/24, which C1 sees as machines on its local subnet.

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.*

   Solution:

(RIP, Figure 1)	*Destination prefix*	*Destination mask*	*Next hop*
R1	*192.168.1.48*	*255.255.255.240*	*192.168.1.33*
	9.23.31.0	*255.255.255.240*	*192.168.1.33*
R2	*192.168.1.16*	*255.255.255.240*	*192.168.1.34*
	9.23.31.1	*255.255.255.240*	*192.168.1.49*
R3	*192.168.1.32*	*255.255.255.240*	*192.168.1.50*
	192.168.1.16	*255.255.255.240*	*192.168.1.50*

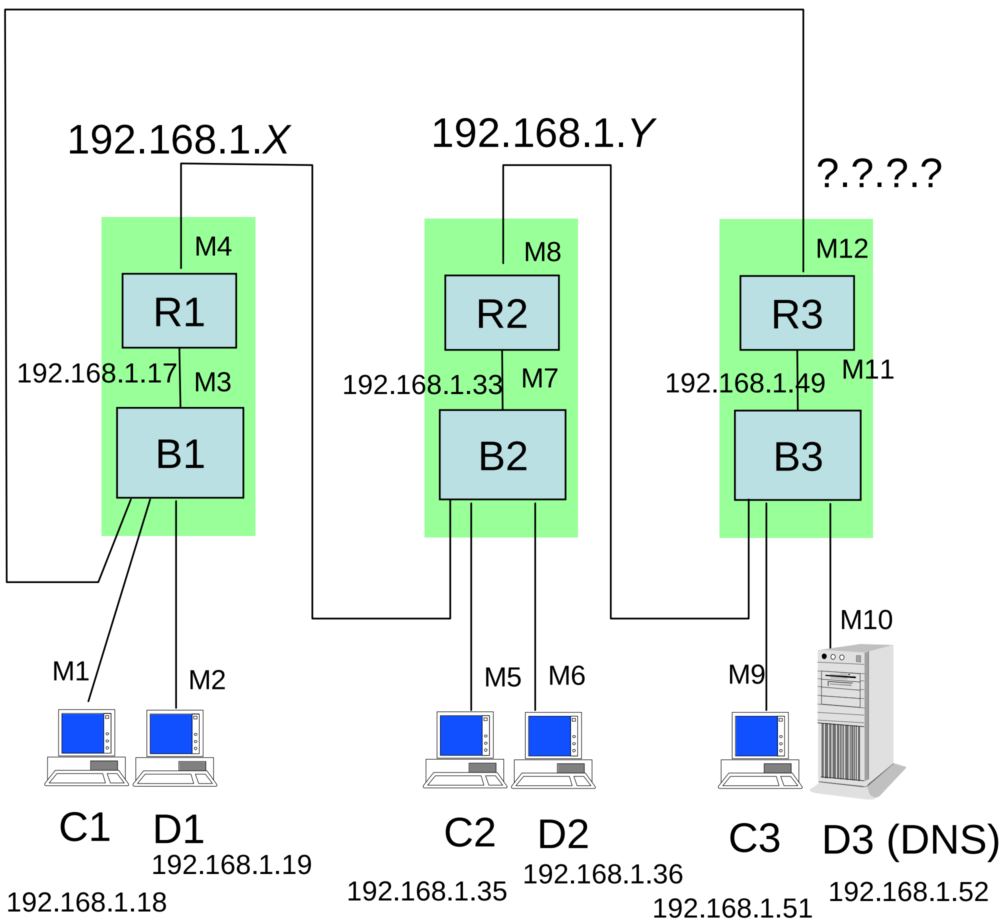

Figure 2: The second network (exercise 2.3).

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 2. What IP address and netmask should we give to M12 ?*

   Solution: An address in subnet 192.168.1.16/28, for example 192.168.1.20. The netmask should be 255.255.255.240 (i.e. the prefix is 28 bits)

   Explain what RIP does immediately after the re-connection ?

   Solution: After discovering that R3 is a neighbor, R1 gets from R3 a route to 192.168.1.48/28 but it is not better than the existing one so there is no change. Symmetrically, R3 gets from R1 a route to 192.168.1.32/28, but it is not better than the existing one, so again there is no change in the routing table. Since R3 is now on-link with 192.168.1.16/28, R3 sends to R2 a route to 192.168.1.16/28, but again it is not better than the one that R2 already has. Finally, the entries in R1 and R2 to 9.23.31.x will timeout but this normally does not cause a message to be sent.

   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 2)	*Destination prefix*	*Destination mask*	*Next hop*
R1	*192.168.1.48*	*255.255.255.240*	*192.168.1.33*
R2	*192.168.1.16*	*255.255.255.240*	*192.168.1.34*
R3	*192.168.1.32*	*255.255.255.240*	*192.168.1.50*

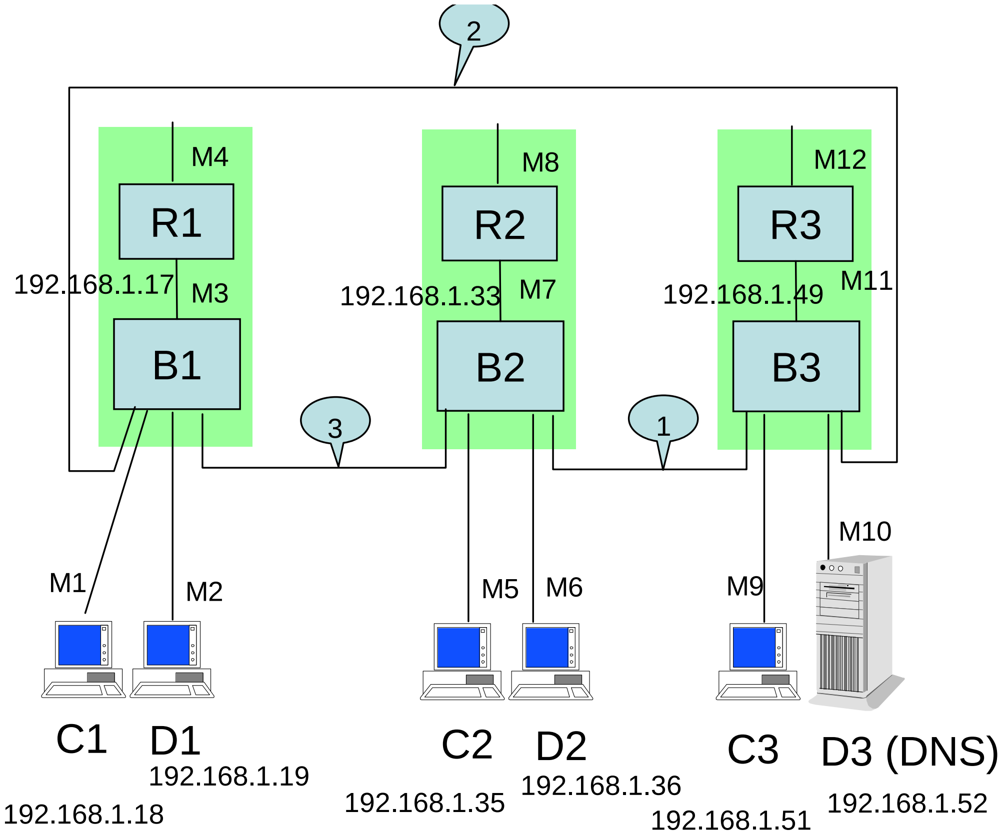

Figure 3: The third network (exercise 2.3).

8. *We reconfigure the network as shown in Figure 3. 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.*

Sniffing Location	*Ping Packet*	*MAC header*		*IP header*		
		Source MAC address	*Destination MAC address*	*Source IP address*	*Destination IP address*	*Protocol*
1	*C1 → C2*					
	C2 → C3	*M5*	*M9*	*192.168.1.35*	*192.168.1.51*	*ICMP*
	C3 → C1	*M9*	*M1*	*192.168.1.51*	*192.168.1.18*	*ICMP*
2	*C1 → C2*					
	C2 → C3					
	C3 → C1					
3	*C1 → C2*	*M1*	*M5*	*192.168.1.18*	*192.168.1.35*	*ICMP*
	C2 → C3					
	C3 → C1	*M9*	*M1*	*192.168.1.51*	*192.168.1.18*	*ICMP*

Solution: All the hosts are on the same network: 192.168.1.0/24. One of the links must be disabled by the spanning tree, so at one of the observation points we should see nothing. We assume it is the link between B1 and B3, so nothing is observed at point 2.