First name: ________________ Family name: ________________

MID TERM EXAM
TCP/IP NETWORKING
Duration: 90 min.
With Solutions

Jean-Yves Le Boudec, Patrick Thiran

2009 November 5

Write your solution into this document and return it to us at the end. You may use additional sheets if needed. Do not forget to put your name on this document and *all* additional sheets of your solution.

If you need to make assumptions in order to solve some questions, write them down explicitly.

Documents or electronic equipment are *not* allowed.

You can write your solution in English or French.

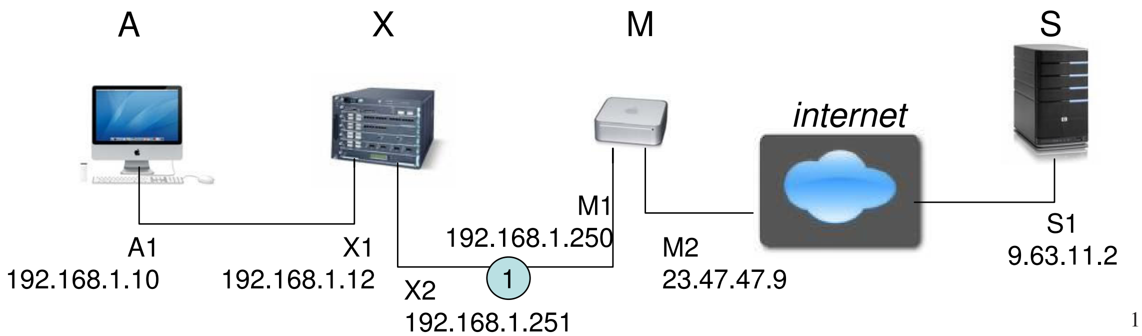

Figure 1: The configuration used in Question 1.

QUESTION 1

Consider the configuration in Figure 1. A is a personal computer, X is a box that can be configured either as a bridge or a router. M is a computer used as web proxy; it does not work as a router. A, X and M are in Joe's apartment; S is a web server on the internet. Only S and M are connected to the internet.

The connections are as indicated on the figure; there is no other connection. M and S are both connected to the internet. The IP addresses are shown on the figure. The MAC addresses are represented as $A1$, $X1$, etc. The circle labeled "1" is an observation point (where we put a packet sniffer).

Unless otherwise specified, all machines are configured to set TTL $= 64$ in the IP packets for which they are the source.

1. In this question, X is configured as a router.
   (a) What should the IP configuration at the interfaces $A1$ and $M1$ be for this setting to work well ? (by configuration we mean: netmask and default gateway only).
   The netmask should allow a separation of $A1, X1$ on one hand, $X2, M1$ on the other hand into two subnetworks. For example: netmask = `255.255.255.192`.
   The default gateway at $A1$ should be 192.168.1.12; at $M1$ it should be 192.168.1.251.
   (b) Assume that all ARP caches are empty. Then A sends a ping packet to `192.168.1.11`; there is no system with this address. We observe all packets resulting from this activity with a packet sniffer at observation point 1. For each of these packets, give: the MAC source and destination addresses; for the packets that are IP packets, give: IP source and destination address and TTL.
   Since the destination IP address in on-link with A, A sends a broadcast ARP request for the non existent address `192.168.1.11`. The router does not respond to or forward the ARP request, so nothing is visible at 1.
   (c) Assume now that the ARP caches and the DNS caches are populated by the correct values. A downloads a file from S. We observe all packets resulting from this activity with a packet sniffer at observation point 1. We observe only the packets flowing from A to S (not in the reverse direction). For each of these packets, give: the MAC source and destination addresses; for the packets that are IP packets, give: IP source and destination address and TTL.
   First notice that there are some packets to be observed, even though this is the reverse direction (TCP acknowledgements). The packets are sent to M, not to S, since M is a proxy and there is no other way for A to access the internet. In other words, M acts as an application layer intermediate system.
   All packets have the same fields:

```
Source MAC address      : X2
Destination MAC address : M1
```

```
Source IP address        : 192.168.1.10
Destination IP address   : 192.168.1.250
TTL                      : 63
```

(d) A runs a media server program that sends broadcast packets with destination IP address 255.255.255.255 and TTL $= 1$. We observe all IP packets resulting from this activity with a packet sniffer at observation point 1. For each of these packets, give: the MAC source and destination addresses; the IP source and destination address; the TTL.
The router does not forward this packet, so nothing is visible at 1.

2. In this question, X is configured as a bridge.

(a) What should the IP configuration at the interfaces $A1$ and $M1$ be for this setting to work well ? (by configuration we mean: netmask and default gateway only).
The netmask should make sure that$A1$, $X1$, $X2$, $M1$ are all on the same subnetwork. For example: netmask = `255.255.255.0`.
The default gateway is absent, since there is no router on this network.

(b) Assume that all ARP caches are empty. Then A sends a ping packet to `192.168.1.11`; there is no system with this address. We observe all packets resulting from this activity with a packet sniffer at observation point 1. For each of these packets, give: the MAC source and destination addresses; for the packets that are IP packets, give: IP source and destination address and TTL.
Since the destination IP address in on-link with A, the ping packets are sent directly to the non existent address `192.168.1.11`. A sends an ARP request to obtain the MAC address of `192.168.1.11`, therefore we see an ARP packet (note that this is not an IP packet) with:

```
Source MAC address      : A1
Destination MAC address : FF:FF:FF:FF:FF:FF
```

There is no answer to this request.

(c) Assume now that the ARP caches are populated by the correct values. A downloads a file from S. We observe all packets resulting from this activity with a packet sniffer at observation point 1. We observe only the packets flowing from A to S (not in the reverse direction). For each of these packets, give: the MAC source and destination addresses; for the packets that are IP packets, give: IP source and destination address and TTL.
All packets have the same fields:

```
Source MAC address       : A1
Destination MAC address  : M1
Source IP address        : 192.168.1.10
Destination IP address   : 192.168.1.250
TTL                      : 64
```

(d) A runs a media server program that sends broadcast packets with destination IP address 255.255.255.255 and TTL $= 1$. We observe all IP packets resulting from this activity with a packet sniffer at observation point 1. For each of these packets, give: the MAC source and destination addresses; the IP source and destination address; the TTL.
The bridge forwards the packet (it is transparent), therefore we see it unchanged at 1. The destination MAC address is the broadcast address. There is no ARP packet resulting from this activity:

```
Source MAC address       : A1
Destination MAC address  : FF:FF:FF:FF:FF:FF
Source IP address        : 192.168.1.10
Destination IP address   : 255.255.255.255
```

TTL : 1

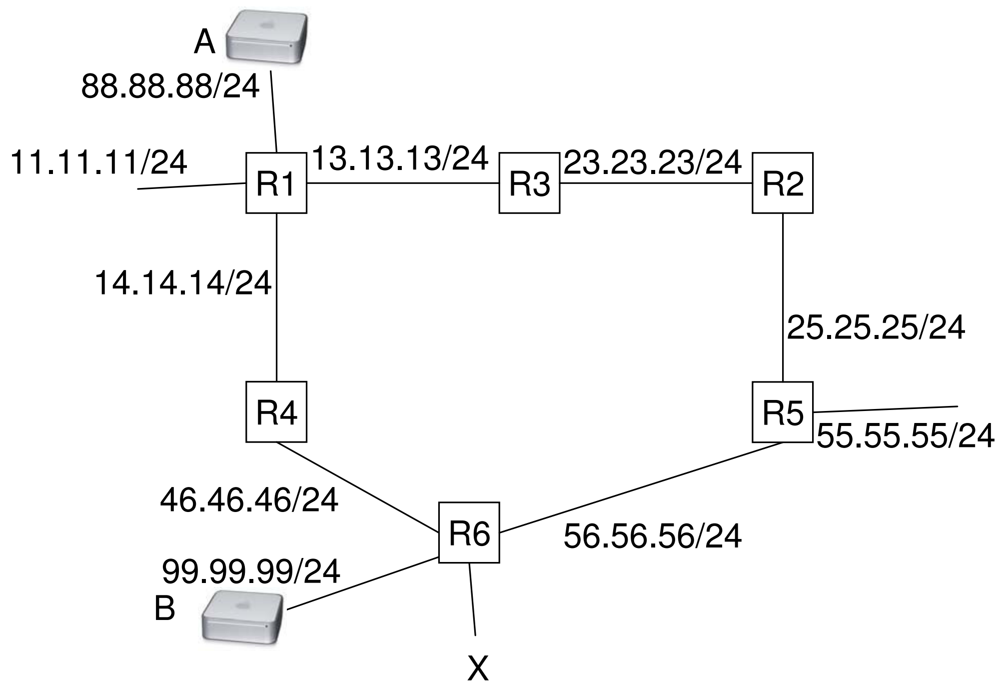

Figure 2: The network used in Question 2.

QUESTION 2

Consider the network in Figure 2. Boxes $R1$ to $R6$ are routers. A and B are hosts. Each router has 2 or more interfaces, as shown on the figure. All prefixes are 24 bits. All routers run a distance vector routing protocol such as RIP. Unless otherwise specified, the cost of a link between two routers is 1. The cost from a router to a directly attached network is 0. There is no other network than shown in the figure.

1. The interface at X is disabled, for the moment. Give a possible value of the routing table at router $R4$, at a time t_1 such that the routing protocol has converged. Give the values in the table below (do not give the value of the "interface" field).

At $R4$		
Destination Network	Next-Hop	Distance
14.14.14/24	directly connected	0
46.46.46/24	directly connected	0
11.11.11/24	14.14.14.1	1
13.13.13/24	14.14.14.1	1
88.88.88/24	14.14.14.1	1
99.99.99/24	46.46.46.1	1
56.56.56/24	46.46.46.1	1
23.23.23/24	14.14.14.1	2
55.55.55/24	46.46.46.1	2
25.25.25/24	46.46.46.1	2

2. At time $t_2 > t_1$, the interface at X is now enabled, but it is wrongly configured: router $R6$ thinks that the prefix associated with X is $55.55.55/24$. $R6$ does its job and sends a routing update to all its

neighbours. Assume that $t_2 - t_1$ is short enough, so that no other message was generated in the time interval $[t_1, t_2]$.
Explain which computation $R4$ will do upon receiving the routing update from $R6$.
$R4$ receives the message from $R6$: `prefix = 55.55.55/24, cost = 0`. $R6$ is the next-hop to that prefix, therefore $R4$ adopts the new value and sets its cost to 55.55.55/24 to

$$c(R4, R6) + 0 = 1$$

and keeps $R6$ as next hop to 55.55.55/24. Then $R4$ sends an update to $R1$: `prefix = 55.55.55/24, cost = 1`.
At a time $t_3 > t_2$ such that the routing protocol has converged again, give the values of the routing table at $R4$ (in the table below; do not give the value of the "interface" field).

At $R4$		
Destination Network	Next-Hop	Distance
14.14.14/24	directly connected	0
46.46.46/24	directly connected	0
11.11.11/24	14.14.14.1	1
13.13.13/24	14.14.14.1	1
88.88.88/24	14.14.14.1	1
99.99.99/24	46.46.46.1	1
56.56.56/24	46.46.46.1	1
23.23.23/24	14.14.14.1	2
55.55.55/24	46.46.46.1	1
25.25.25/24	46.46.46.1	2

3. At time $t_4 > t_3$, the configuration at router $R6$ is changed as follows. The interface at X continues to be enabled and wrongly configured with the prefix 55.55.55/24. But from now on, $R6$ sees subnet X as a subnet with distance 10 (instead of 0 as before). $R6$ does its job and sends a routing update to all its neighbours. Assume that $t_4 - t_3$ is short enough, so that no other message was generated in the time interval $[t_3, t_4]$.
   Explain which computation $R4$ will do upon receiving the routing update from $R6$.
   $R6$ now has two entries in the RIP table: `prefix = 55.55.55/24, cost = 1` and `prefix = 55.55.55/24, cost = 10`. $R6$ will promote the route with shorter distance to its routing table and advertise it to the neighbors (respecting the split horizon rule). This means that $R4$ will receive the following message from $R6$: `prefix = 55.55.55/24, cost = 1`. $R6$ is the next-hop to that prefix, therefore $R4$ adopts the new value and sets its cost to 55.55.55/24 to

$$c(R4, R6) + 1 = 2$$

   and keeps $R6$ as next hop to 55.55.55/24. Then $R4$ sends an update to $R1$: `prefix = 55.55.55/24, cost = 2`.
4. Let $t_5 > t_4$ be the time at which the routing protocol converges again. Give a possible sequence of messages received by $R1$ during the time interval $[t_4, t_5]$, with the computations that $R1$ performs upon receiving these messages.
   $R1$ will receive updates with the same distance from both $R3$ and $R4$:
   `prefix = 55.55.55/24, cost = 2, next-hop = R3`
   `prefix = 55.55.55/24, cost = 2, next-hop = R4`
   Since the routing protocol used in this network is RIP, depending on configuration, $R1$ will add one, or both of these updates to the routing table (RIP2 allows equal cost load balancing with up to six routes).

$R1$ will send the route(s) added to the routing table message to the neighbors respecting the split horizon rule.

5. At time $t_6 > t_5$, you start a program at A that continuously pings B. Then, at time $t_7 > t_6$, the router $R4$ goes down. Explain what will happen with RIP (explain in at most 15 lines a possible sequence of events, until the routing protocol converges). Explain what A will observe regarding all sent pings. Ping messages from A won't be able to reach B after the failure. $R6$ and $R1$ detect the failure after the timeout. $R6$ will promote, from the RIP table to the routing table, alternative routes to prefixes (`11.11.11/24, 13.13.13/24, 14.14.14/24, 88.88.88.88/24`) that had $R4$ as next hop. They will be replaced with the routes that use $R5$ as the next hop. Similarly, $R1$ will promote, from the RIP table to the routing table, alternative routes to prefixes (`46.46.46/24, 56.56.56/24, 99.99.99/24`) that had $R4$ as next hop. They will be replaced with the routes that use $R3$ as the next hop. Similar recalculation is going to happen in the routers $R3$ and $R5$. The ping request and reply will start working again.

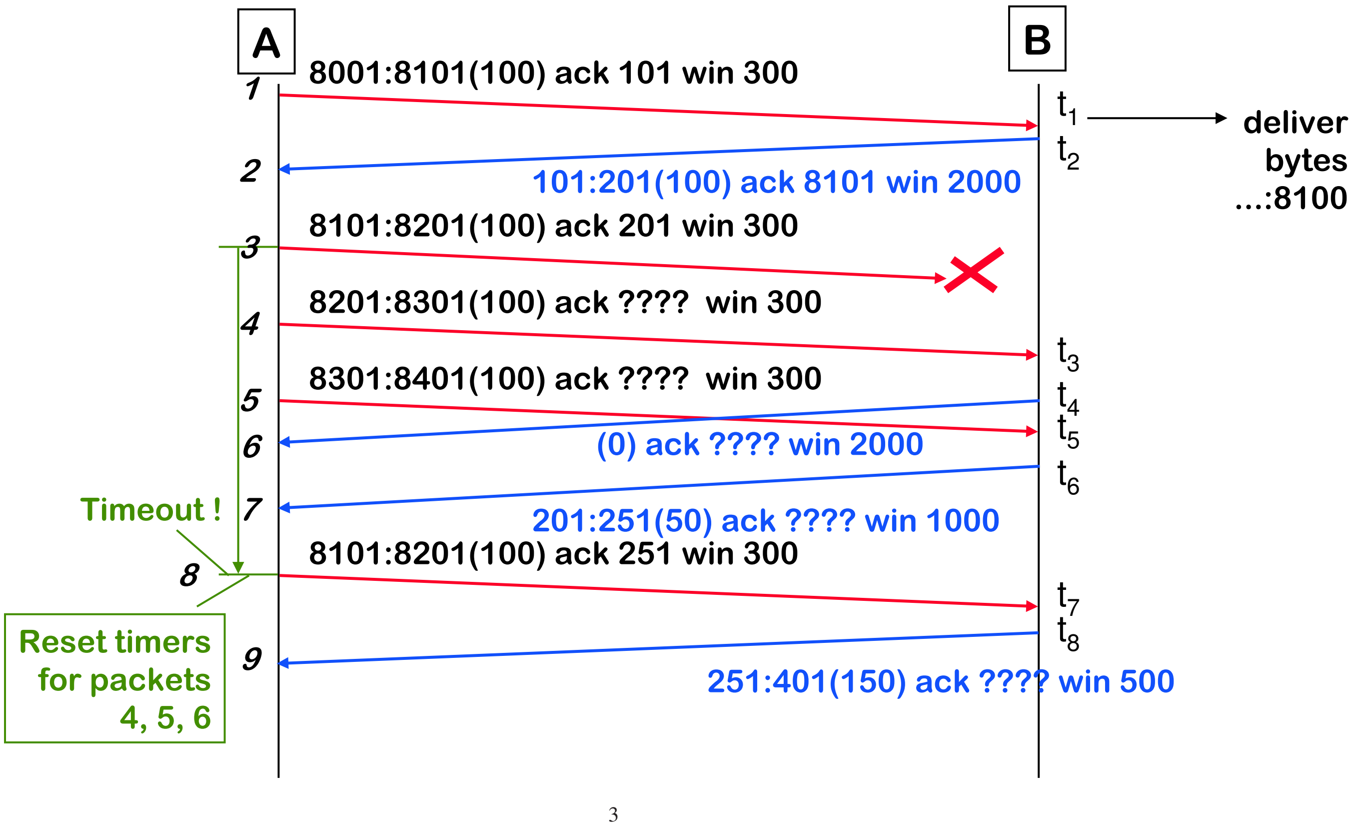

Figure 3: Message sequence diagram used in Question 3.

QUESTION 3

1. Figure 3 shows one possible message sequence for two hosts A and B that are transferring data using one TCP connection. The times t_1 to t_8 are the instants at which a TCP segment is either received or sent by B. The figure shows the values of sequence numbers and acknowledgements, but some acknowledgement numbers are missing.
The figure shows some of the fields relative to the TCP header. For example, the first packet is shown with

   `8001:8101(100) ack 101 win 1000`

   which means the following:

   - the segment contains the bytes with sequence numbers 8001 to 8100; there are 100 data bytes in the packet
   - the ACK number of 101
   - the advertised window size is 1000

   The figure also shows that, at time t_1, B delivers bytes numbered up to 8100 to the application layer.

   (a) Give the values of the missing acknowledgement numbers. Use the table below.

Packet at line ...	Acknowledgement number
4	201
5	201
6	8101
7	8101
9	8401

(b) At what time are bytes 8301 to 8400 delivered to the application by B ?

At t_7.

(c) After receiving the packet at line 9, A has 1000 bytes of data ready to be sent. The maximum segment size at A is 100 bytes. How many TCP segments can A immediately send without waiting for acknowledgement ?

The window size (advertised by B) is 500. Since A has no non-acknowledged packet, A may use the full window size and will thus send 5 packets consecutively.

2. Using the network diagnostic tools (`ndt.switch.ch`), we transfer a large file from a host S to a host C. C is in Switzerland while S is in some unknown geographical location. The file is transferred using one TCP connection, with a maximum window size of[4] 10^5 bits. We find that the throughput achieved on the connection is 5 Mb/s. Is it possible that S is in Australia ?

The maximum throughput θ achievable with a window size W is $\frac{W}{\tau}$ where τ is the round trip time, thus

$$\tau \leq \frac{W}{\theta} = \frac{10^5}{5 \times 10^6} = 0.02 \text{ sec} = 20 \text{ msec}$$

The round trip time to Australia is at least given by the propagation time, which is close to 0.2 sec. Therefore it is impossible. In fact, S must be very close, probably in Switzerland.

[4] The window size is usually given in bytes, but we give it in bits to avoid numerical complexity