MID-TERM EXAM
TCP/IP NETWORKING
Duration: 2 hours

Jean-Yves Le Boudec

2005 December 8

Do not forget to put your names on *all* sheets of your solution.

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

Manage your time: All problems have equal weight in the final appreciation; consider this before spending too much time on one question.

The exam is open book. You can use all written documents, but no electronic equipment.

You can write your solution in English, French or German.

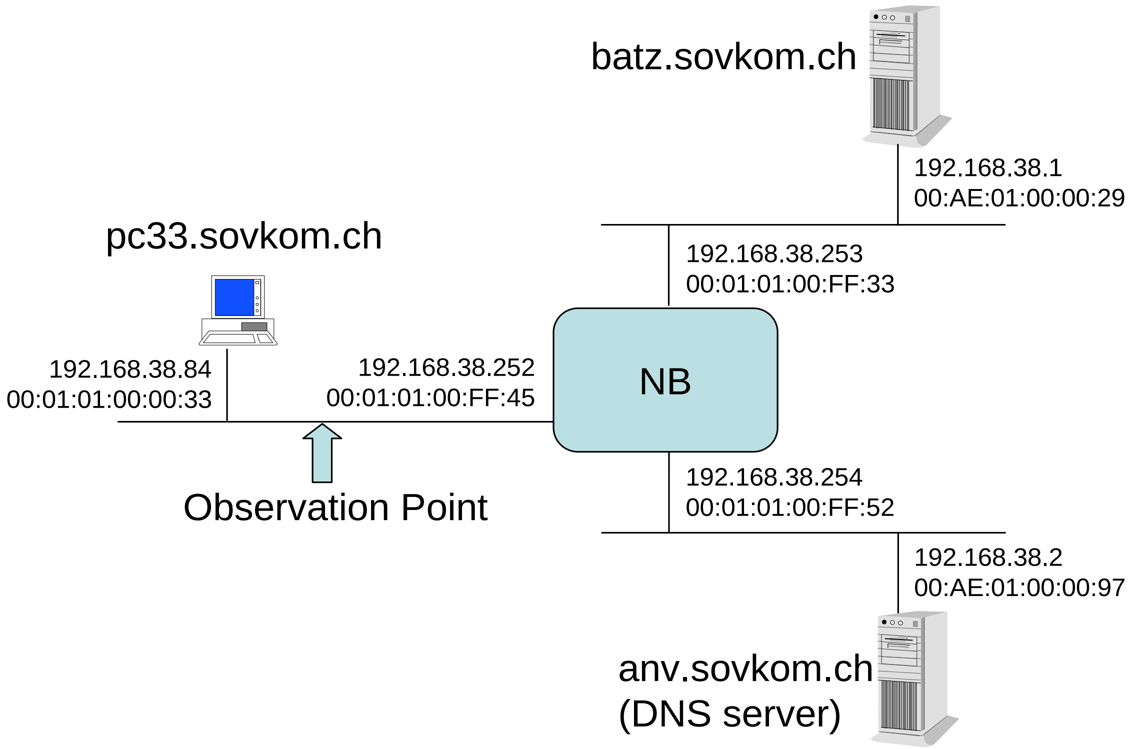

Figure 1: The network of Problem 1, with NB configured as bridge

PROBLEM 1

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

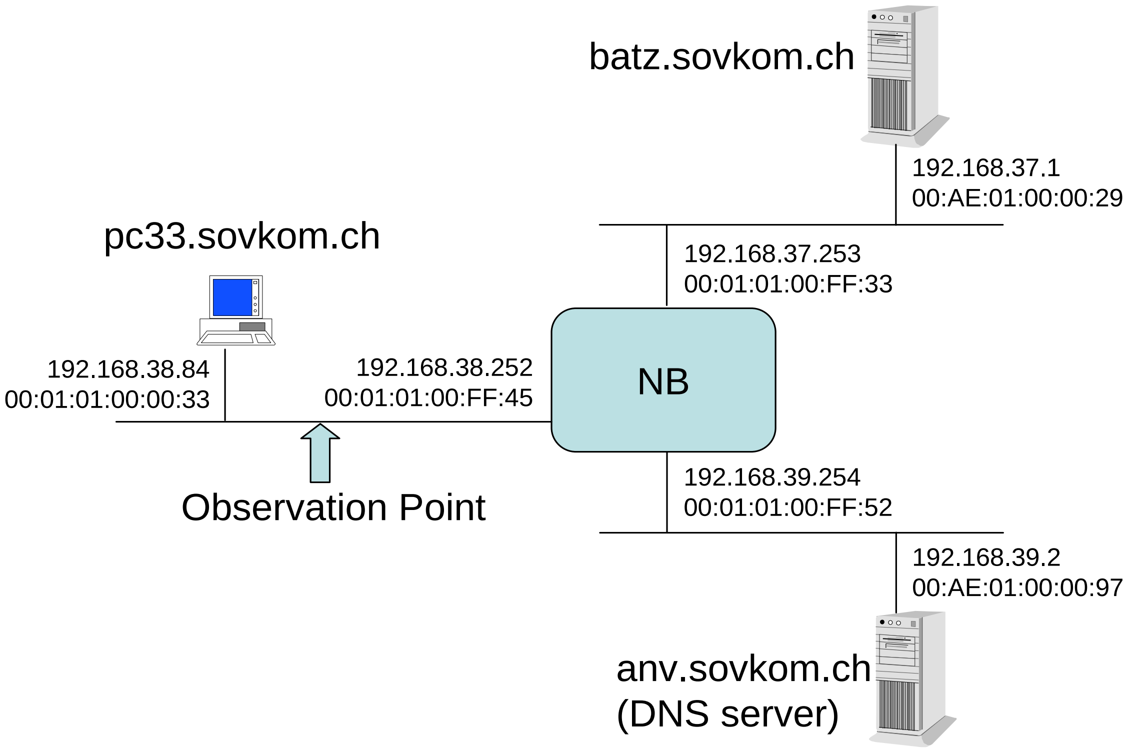

Figure 2: The network of Problem 1, with NB configured as router

- IP source address
- 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.

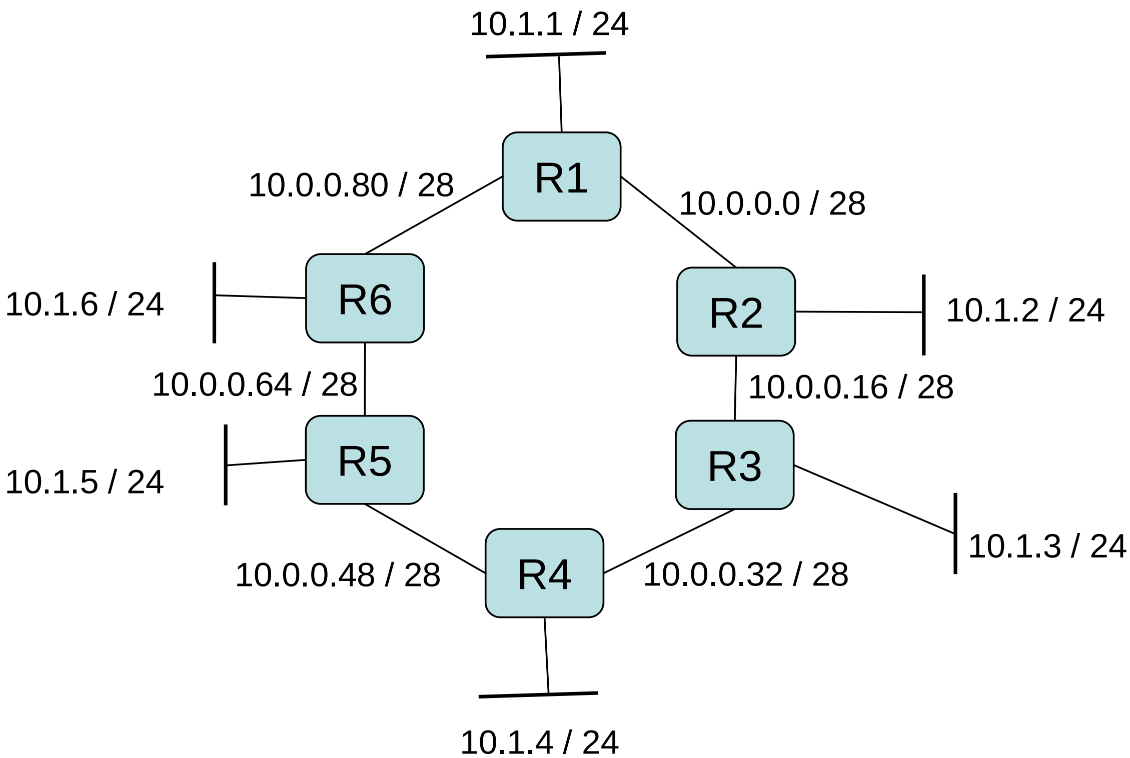

Figure 3: The network of Problem 2

PROBLEM 2

Consider the network in Figure 3. 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.
- 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 ?

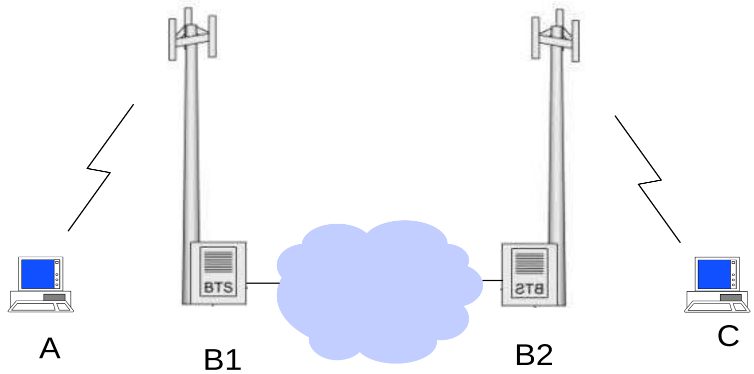

Figure 4: The network of Problem 3, question 2

PROBLEM 3

In this problem you need to make some assumptions. Please describe them explicitly.

1. Assume we send a file with a sliding window protocol from EPFL to a host in New Zealand. We do not know exactly all the details of the sliding protocol, but we do know the following.
   - The file is such that it takes $n = 10$ packets of size $L = 10^4$ bits to transmit it (the Path MTU is equal to L).
   - The bit rate available for transmission is $R = 10^6$b/s.
   - The destination sends one ack for every packet received

   (a) Assume we use a window size $W = 10^4$ bits. What is the minimum time it takes to transmit the file and receive all necessary acknowledgements ?

   (b) Same question with a window size $W \geq nL$.

2. Assume we want to transmit a very large file from A to C using HTTP on the system of Figure 4. We consider two options for the base stations B1 and B2.

   Option 1 B1 and B2 act as a router.
   Option 2 B1 and B2 act as an application layer gateways.

   A and B1 [resp. C and B2] are connected by a wireless link at 1 Mb/s. B1 and B2 are connected by a wired infrastructure at a very high bit rate (much larger than 1 Mb/s). The residual packet loss rates on links A-B1 and C-B2 (after all recovery at the MAC layer) are 25%. The packet loss rate between B1 and B2 is negligible.
   For each of the two options, what is the throughput achievable (i.e. capacity of the end to end path) ?

PROBLEM 4

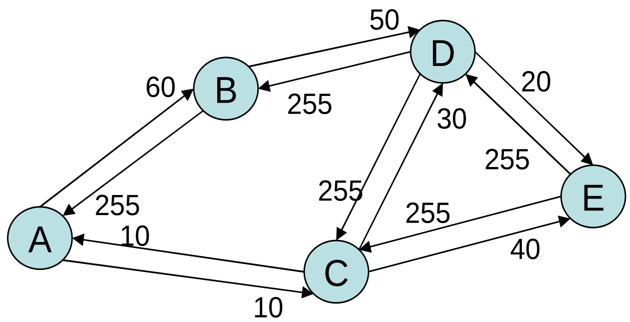

Consider a new routing method, that would work as follows. Assume we have a network with routers and links, all of equal bit rate. The cost of a link is an integer between 0 and 255 that represents the congestion status, with a high cost meaning a congested link (it is called *congestion cost*). Links are unidirectional, the directions may have different costs. The cost of a *path* is, by definition, the *maximum* of the costs of the constituent links (it is also called the *congestion cost* of the path).

Consider for example the picture above, with the congestion costs of the links as shown: the congestion cost of the link $A \to B$ is 60, of the link $B \to A$ is 255. The congestion cost of the path $A \to B \to D$ is

$$60 \vee 50 = 60$$

where the notation $a \vee b$ means the maximum of a and b. We see on this example that the congestion cost of a path is the congestion cost of the most congested link on the path.

We say that a path between two nodes X and Y is "least congested" if it has no cycle and its congestion cost is minimum, among all paths from X to Y.

1. Give, for the picture above, a least congested path between A and E ? What is its congestion cost ?

Assume we have some procedure to determine, every 5 minutes, all links costs. Assume these link costs are fed to a central network management station, which uses these costs to compute the congestion costs from any router to any router. The following centralized algorithm is used.

For all destination routers k do
    For all routers i do $q^0(i,k) = +\infty$
    $q^0(k,k) = 0$
    For $n = 1$ to $n_{\max}$ do
        For all routers $i \neq k$ do
            $q^n(i,k) = \min_{j:\text{ neighbour of } i, j \neq i} \left(c(i,j) \vee q^{n-1}(j,k) \right)$
        $q^n(k,k) = 0$
        If $q^n = q^{n-1}$ leave

where $c(i,j)$ is the congestion cost of the link from router i to router j, and $n_{\max}$ is some predefined large integer, much larger than the network diameter. The condition $q^n = q^{n-1}$ means that $q^n(i,k) = q^{n-1}(i,k)$ for all i, k. Thus, the algorithm stops either because $n_{\max}$ is reached, or because q^n does not change. In the latter case we say that the algorithm converges.

2. Apply the algorithm to the example on the figure for $k = E$ (i.e. show in a table the values of $q^n(i,E)$ for $i \in \{A, B, C, D, E\}$ and for $n = 1, 2, 3, 4, \ldots$). Does the algorithm converge ? If so, in how many iterations ?

3. Someone claims that, in this algorithm, $q^n(i,j)$ is the congestion cost of all paths from i to j that have n hops or less. Is this true ? (justify your answer)
4. Does the algorithm provide the same result if we change the initial conditions ? (justify your answer)