Mid-term Exam
TCP/IP Networking
Duration: 2 hours
With Solutions

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.

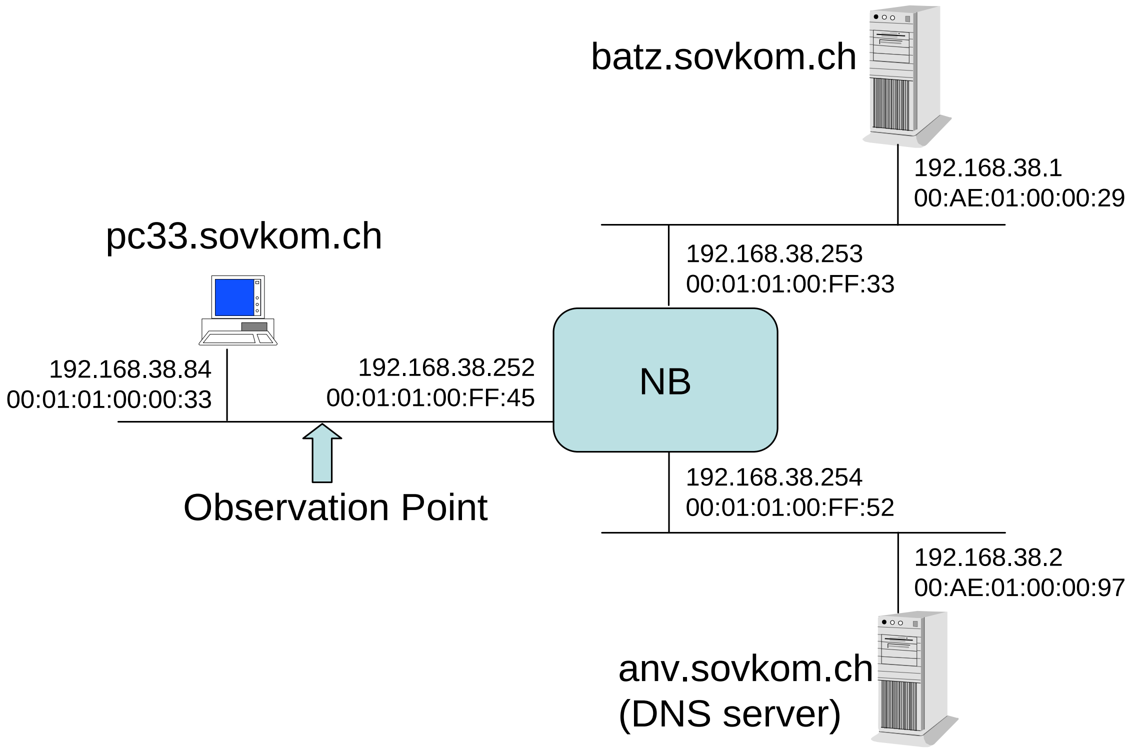

Figure 1: The network of Problem 1, with NB configured as bridge

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

- 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.

Solution.

(a) All the systems in the network belong to the same ethernet network: as all the IP addresses in the network are correct and as they belong to the same subnet, they all are correct.

(b) Nb needs at least one IP address, as it runs a web server.

(c) Let's use the following labels:

Device Address	Symbol
PC33 (IP)	a1
PC33 (MAC)	a2
DNS server (IP)	b1
DNS server (MAC)	b2
Batz (IP)	c1
Batz (MAC)	c2
NB (IP)	n1
NB (MAC)	n2

For NB we consider only the network interface on the same ethernet segment as pc33.

Packet	IP src	IP dest	MAC src	MAC dst	protocol type-dest port
pc33⟶ all	(a1)	(b1)	a2	(broadcast)	ARP
DNS ⟶ pc33	(b1)	(a1)	b2	a2	ARP
pc33⟶ DNS	a1	b1	a2	b2	UDP (port 53)-dns query
DNS⟶ pc33	b1	a1	b2	a2	UDP (local port)-dns answer
pc33⟶ all	(a1)	(c1)	a2	(broadcast)	ARP
batz⟶ pc33	(c1)	(a1)	c2	a2	ARP
pc33⟶ batz	a1	c1	a2	c2	TCP (port 13)- syn
batz⟶ pc33	c1	a1	c2	a2	TCP (local port)- syn ack
pc33⟶ batz	a1	c1	a2	c2	TCP (port 13)- ack
batz⟶ pc33	c1	a1	c2	a2	TCP (local port)- (data, fin)
pc33⟶ batz	a1	c1	a2	c2	TCP (port 13)- ack
pc33⟶ batz	a1	c1	a2	c2	TCP (port 13)- fin
batz⟶ pc33	c1	a1	c2	a2	TCP (local port)- ack

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.

Solution.

1. *From the picture, the two servers and the PC33 each belong to a different ethernet network, connected each to a separate network interface of the router. Then each IP address in each ethernet network belong to the same subnet, and all are correct: then all the IP address assigned in the picture are acceptable.*
2. *In order to route packets, NB needs IP addresses for all of its three network interfaces.*
3. *We use the same labels as before:*

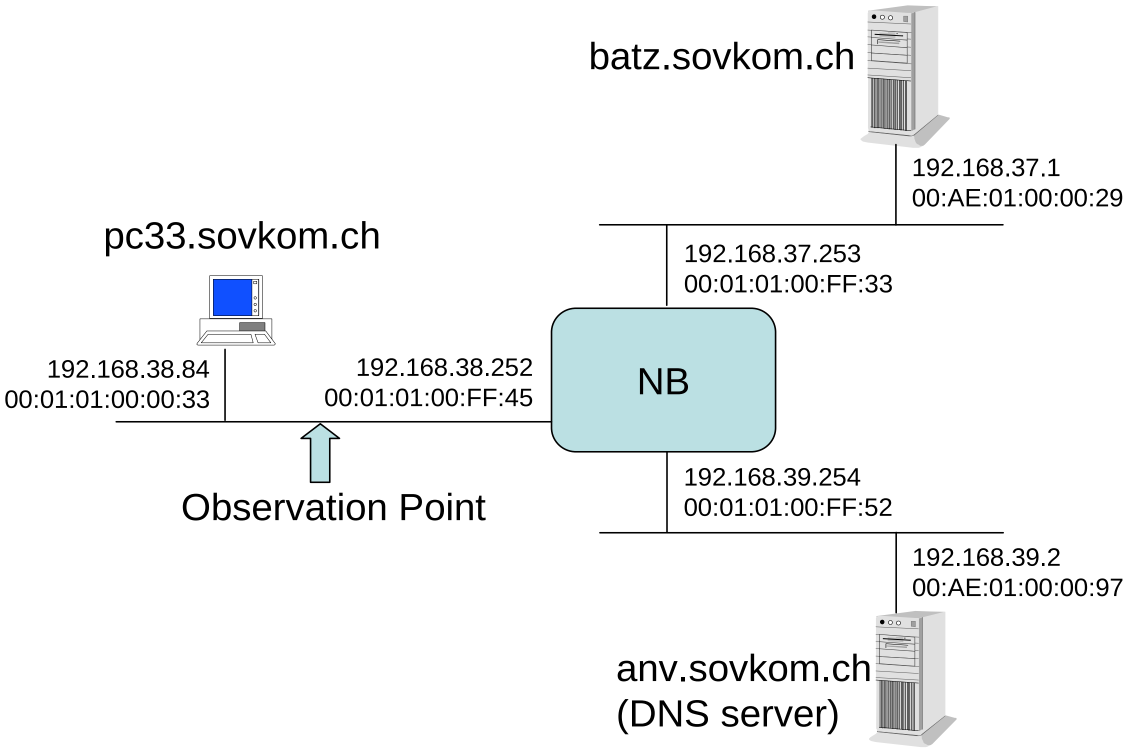

Figure 2: The network of Problem 1, with NB configured as router

Packet	*IP src*	*IP dest*	*MAC src*	*MAC dst*	*protocol type-dest port*
pc33⟶ all	*(a1)*	*(n1)*	*a2*	*(broadcast)*	*ARP*
NB ⟶ pc33	*(n1)*	*(a1)*	*n2*	*a2*	*ARP*
pc33⟶ DNS	*a1*	*b1*	*a2*	*n2*	*UDP (port 53)-dns query*
DNS⟶ pc33	*b1*	*a1*	*n2*	*a2*	*UDP (local port)-dns answer*
pc33⟶ batz	*a1*	*c1*	*a2*	*n2*	*TCP (port 13)- syn*
batz⟶ pc33	*c1*	*a1*	*n2*	*a2*	*TCP (local port)- syn ack*
pc33⟶ batz	*a1*	*c1*	*a2*	*n2*	*TCP (port 13)- ack*
batz⟶ pc33	*c1*	*a1*	*n2*	*a2*	*TCP (local port)- (data, fin)*
pc33⟶ batz	*a1*	*c1*	*a2*	*n2*	*TCP (port 13)- ack*
pc33⟶ batz	*a1*	*c1*	*a2*	*n2*	*TCP (port 13)- fin*
batz⟶ pc33	*c1*	*a1*	*n2*	*a2*	*TCP (local port)- ack*

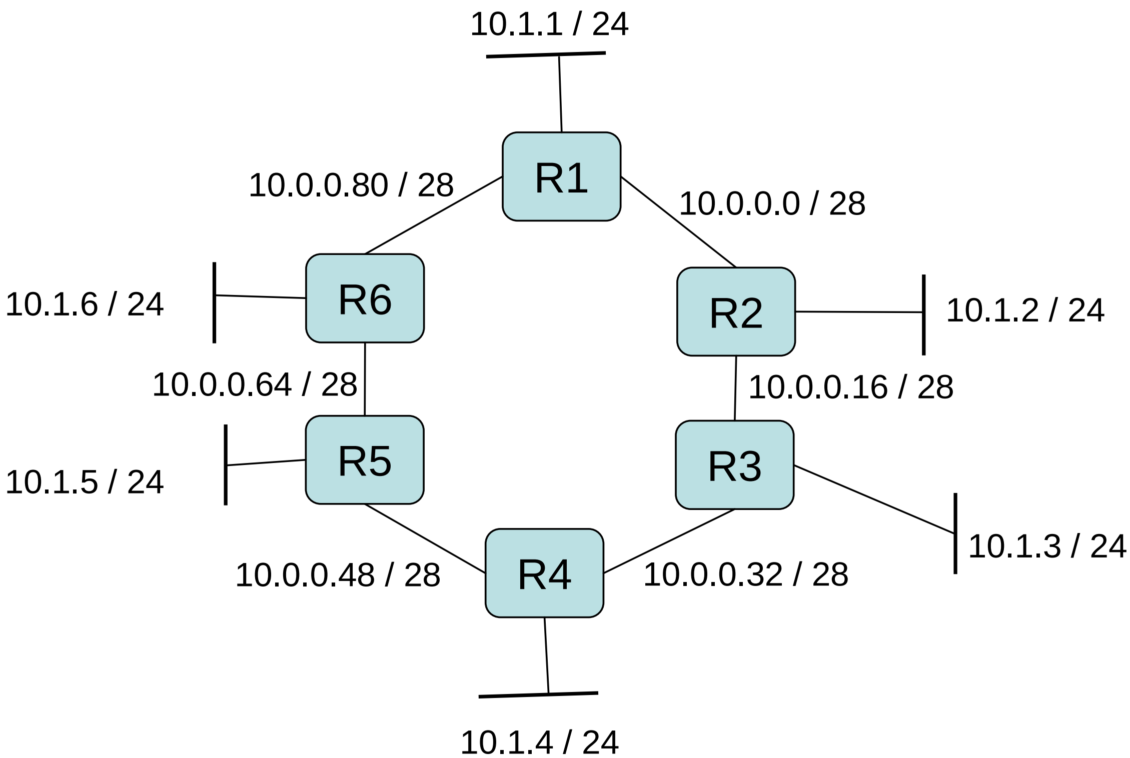

Figure 3: The network of Problem 2

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)

   Solution. *Backbone: 255.255.255.240. Edge: 255.255.255.0.*

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.

   Solution. *Let's assume that the last two bits of the clock-wise side backbone interfaces are equal to* 01 *and the other side backbone interfaces have the last two bits equal to* 10*. Lets also assume that the edge interfaces are called eth0, clock-wise side backbone interfaces are called eth1 and the other side backbone interfaces are called eth2. Then the routing table is:*

Destination Network	*Next Hop*	*Interface*	*Distance*
10.1.1.0/24	*on-link*	*eth0*	*1*
10.0.0.0/28	*on-link*	*eth1*	*1*
10.1.2.0/24	*10.0.0.2*	*eth1*	*2*
10.0.0.16/28	$-\|\|-$	$-\|\|-$	*2*
10.1.3.0/24	$-\|\|-$	$-\|\|-$	*3*
10.0.0.32/28	$-\|\|-$	$-\|\|-$	*3*
10.1.4.0/24	$-\|\|-$	$-\|\|-$	*4*
10.0.0.48/28	*10.0.0.81*	*eth2*	*3*
10.1.5.0/24	$-\|\|-$	$-\|\|-$	*3*
10.0.0.64/28	$-\|\|-$	$-\|\|-$	*2*
10.1.6.0/24	$-\|\|-$	$-\|\|-$	*2*
10.0.0.80/28	*on-link*	$-\|\|-$	*1*

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 ?

   Solution. *The possible default routers for M: $R2$ and $R3$. For A: $R1$.*
   Case $(R2, R1)$: the path is $M - R2 - R1 - A$.
   Case $(R3, R1)$: for the first packet the path is $M - R3 - R2 - R1 - A$. Then, the redirection will happen and the path $M - R2 - R1 - A$ will be used for the subsequent packets.

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.

   Solution. *In the next RIP update message (sent every 30s) R2 will announce toward its neighbors the network 10.1.7 with metric 1 and will not announce the network 10.1.2 (which is equivalent to the metric ∞). Bellman-Ford algorithm will cause the updates of routing tables or R1 and R3. R1 and R3 will propagate changes of their routing tables in their next update messages sent towards their neighbors. The change propagates further in the same way.*

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 ?

   Solution. *If split horizon is not used there will be "count to infinity" scenario: R2 will increment it's metric toward 10.1.2/24 to 3 and tell this route to R1 and R3. R1 will increment its metric to 4 and tell to R2 and R6, and so on. In this way there is count to infinity between R1 and R2.*
   If split horizon is used, R1 will not send to R2 an update about the network 10.1.2/24.

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.

   Solution. *The actions: RIP update messages exchange and routing table updates according to*

Bellman-Ford algorithm. Only the entries toward 10.1.7/24 will change at some of the routers:

Router	*Destination Network*	*Next Hop*	*Distance*	*changed?*
R1	*10.1.7.0/24*	*R2*	*2*	*no*
R2	*10.1.7.0/24*	*on-link*	*1*	*no*
R3	*10.1.7.0/24*	*R2*	*2*	*no*
R4	*10.1.7.0/24*	*R5*	*2*	*yes*
R5	*10.1.7.0/24*	*on-link*	*1*	*yes*
R6	*10.1.7.0/24*	*R5*	*2*	*yes*

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 ?

   Solution. *A→B2:* $A - R1 - R2 - B2$*, success.*
   A→B5: $A - R1 - R2$*, then ARP will fail, and the ICMP "destination host unreachable" message will be sent by R2 to A, and the packet will be dropped by R2.*

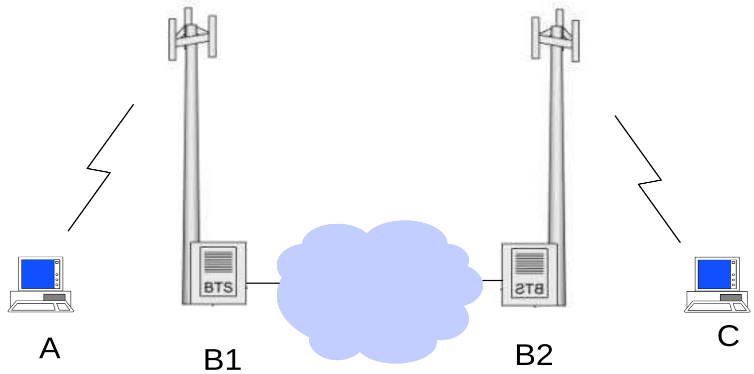

Figure 4: The network of Problem 3, question 2

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 $W \geq nL$.

   Solution.

   (a) In this case we can send only one packet before receiving an acknowledgement. Assuming the transmission time for the acknowledgements is negligible (that is, they are small enough), the minimum time to transmit the whole file and receive all necessary acknowledgements is: $10 \times (RTT + L/R)$*, where* RTT *is the roundtrip time.*

   (b) In this case we can send all packets without waiting for any acknowledgements. So, the minimum time is $10L/R + RTT$*, since we need to wait for the last acknowledgement.*

   The transmission time for one packet is $L/R = 10ms$ *and the roundtrip time to New Zealand is approximately 300ms.*

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) ?

   Solution.

Option 1 *The end to end packet delivery probability is equal to the product of the delivery probabilities of each link. The delivery probability of each wireless link (A-B1 and C-B2) is (1− loss rate) equal to 0.75. Therefore, the end to end packet delivery probability is* $0.75^2 = 56.25\%$*; and the end to end capacity is* $1Mb/s \times 0.5625 = 0.5625Mb/s$.

Option 2 *The end to end capacity is the minimum of the link capacities. The capacity of each wireless link is* $0.75Mb/s$ *so the end to end capacity is 0.75 Mb/s.*

PROBLEM 4

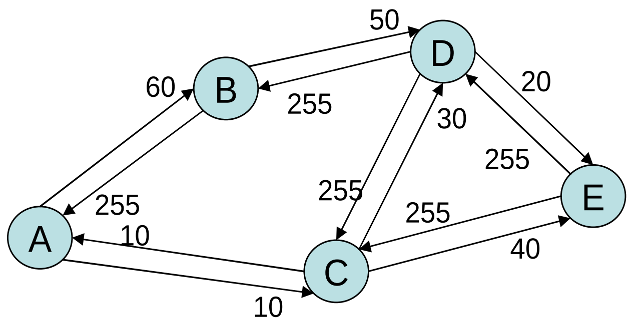

Consider a new routing method, that would work as follows. Assume we have a network with routers and links, all of equal bit rate. The cost of a link is an integer between 0 and 255 that represents the congestion status, with a high cost meaning a congested link (it is called *congestion cost*). Links 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 ?

   Solution. *$A \to C \to D \to E$, congestion cost is 30*

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 $n = 1,2,3,4,...$). Does the algorithm converge ? If so, in how many iterations ?

Solution.

n	A	B	C	D	E
0	$+\infty$	$+\infty$	$+\infty$	$+\infty$	0
1	$+\infty$	$+\infty$	40	20	0
2	40	50	30	20	0
3	30	50	30	20	0
4	30	50	30	20	0

The algorithm stops at $n = 4$ (i.e. $q^3 = q^4$).

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)

   Solution. *Yes, it is true. We prove by induction the statement "$q^n(i, k)$ is the congestion cost from i to k in at most k hops".*

   (a) *$n = 1$: the congestion cost from i to k in 0 or 1 hop is*
   - *0 if $i = k$*
   - *else if k is a neighbour of i: $c(i, k)$*
   - *else $+\infty$*

   This is precisely what the algorithm gives for $n = 1$.

   (b) *Assume the statement is true up for $n - 1$; we show that it is true for n. Let c be the congestion cost c from i to k in at most n hops. We want to show that $c = q^n(i, k)$.*
   Now c is the cost of one path from i to k that has at most n hops. Pick one such path, and let j be the next hop starting from i on this path. The rest of this path leads from j to k in at most $n - 1$ hops, thus its congestion cost c' staisfies, by the induction hypothesis: $c' \geq q^{n-1}(j, k)$. Now the congestion cost of this path (which is equal to c) is

   $$c = c(i, j) \vee c' \geq c(i, j) \vee q^{n-1}(j, k)$$

   thus $c \geq q^n(i, k)$.
   Conversely, by definition of the algorithm, $q^n(i, j)$ is equal to $c(i, j_0) \vee q^{n-1}(j_0, k)$ for some *j_0, thus it is the cost of* some *path from i to k in at most n hops. Thus $q^n(i, k) \geq c$. Combining the two shows that $q^n(i, k) = c$*

4. Does the algorithm provide the same result if we change the initial conditions ? (justify your answer)

   Solution. *no, for example if we start with $q^0(i, E) = 0$ the algorithm stops at a wrong value:*

n	A	B	C	D	E
0	0	0	0	0	0
1	10	50	10	20	0
2	10	50	10	20	0