First name: _____ Family name: _____

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

Jean-Yves Le Boudec

2007 November 8

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.

All printed documents are allowed. Electronic equipment is not allowed except calculators.

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

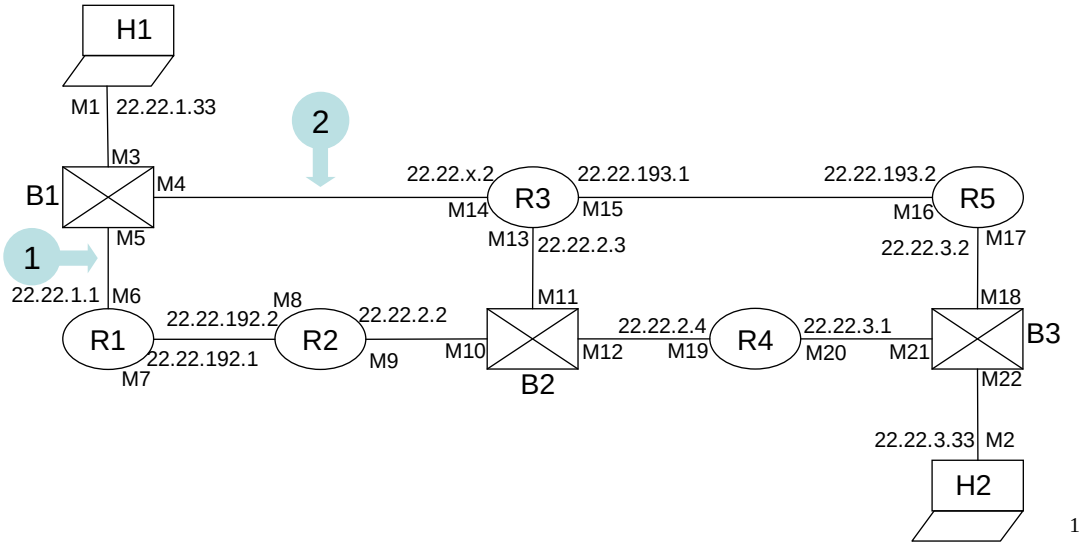

Figure 1: The network used in Question 1.

QUESTION 1

Consider the network in Figure 1. All links are Ethernet cables. B1 to B3 are bridges. R1 to R5 are routers. H1 and H2 are hosts. There is no other network than shown on the figure.

The figure shows the IP and MAC addresses. The MAC addresses are written symbolically as M1 to M22.

All interfaces are configured with network mask 255.255.255.0. All routers run a distance vector routing protocol.

Since B2 is a bridge, routers R2 and R3 are neighbours for the distance vector routing protocol (and similarly, R2 and R4 are neighbours, R3 and R4 are neighbours, and R1 and R3 are neighbours.

The link costs are the same in both directions. They are configured as follows.

R1-R3: 10
R1-R2: 1
R2-R3: 1
R2-R4: 1
R3-R4: 1
R3-R5: 10
R4-R5: 1

Thus assume that the costs given above are the costs of the specific networks. For directly connected networks this then means that e.g. the cost from R5 to `22.22.3/24` is 1, the cost from R1 to 22.22.1/24 is 10.

Also assume that the initial TTL value is 64 for all machines.

1. Give a possible value of the routing table entry at each router for destination network = `22.22.3/24`, assuming the routing protocol has converged. Give the values in the table below (do not give the value of the“interface” field).

Destination Network = 22.22.3/24		
At Router	*Next-Hop*	*Distance*
R1	*22.22.192.2*	*3*
R2	*22.22.2.4*	*2*
R3	*22.22.2.4*	*2*
R4	*on link*	*1*
R5	*on link*	*1*

2. Host H1 is configured with default gateway = 22.22.1.1. We observe all packets with packet sniffers at points 1 and 2. All caches are empty at H1. Then H1 decides to send one single IP packet with IP destination address H2. Explain which packets, resulting from this activity, will be seen at observation point 1. Also explain which packets will be seen at observation point 2. For each of the packets, give the value of the following fields:
   - MAC source address, MAC destination address
   - For those packets that are IP packets (and only those!) give: source IP address, IP destination address, TTL.

@pt1:

type	*MAC src*	*MAC dest*	*IP src*	*IP dest*	*TTL*
ARP request	*M1*	*FF:FF:FF:FF:FF:FF*	*-*	*-*	*-*
ARP reply	*M6*	*M1*	*-*	*-*	*-*
IP packet	*M1*	*M6*	*22.22.1.33*	*22.22.3.33*	*64*

@pt2:

type	*MAC src*	*MAC dest*	*IP src*	*IP dest*	*TTL*
ARP request	*M1*	*FF:FF:FF:FF:FF:FF*	*-*	*-*	*-*

3. What should the value x in the IP address at R3 be ? Justify your answer.

x=1 otherwise R3 would not be part of the same subnet. B1 is a bridge, it only separates collision domains.

Re-do the previous question (question 1.2) assuming the default gateway at H1 is now `22.22.`x`.2`, where x is the value you have found.

@pt1:

type	*MAC src*	*MAC dest*	*IP src*	*IP dest*	*TTL*
ARP request	*M1*	*FF:FF:FF:FF:FF:FF*	-	-	-

@pt2:

type	*MAC src*	*MAC dest*	*IP src*	*IP dest*	*TTL*
ARP request	*M1*	*FF:FF:FF:FF:FF:FF*	-	-	-
ARP reply	*M14*	*M1*	-	-	-
IP packet	*M1*	*M14*	*22.22.1.33*	*22.22.3.33*	*64*

4. The link between R2 and B2 fails. Explain what happens in all routers, in particular respective to the destination network `22.22.3/24`, by giving a possible sequence of events (max 15 lines). Your sequence of events needs to be detailed and specific (e.g. it is not enough to just say that the protocol converges but you need to give a possible sequence of packets exchanged and decisions made at the routers).

   - *R2, R3, R4 detect failure after timeout*
   - *they set the cost of routes to networks, whose next hop is one of the neighbors who timed out, to infinity. R2 for example had R4 as next hop to reach network `22.22.3/24` with a cost of 2, it will change this to infinity. R3 and R4 will not change anything with respect to `22.22.3/24` as their routes were not affected by the failure. R3 and R4 will however set e.g. costs to `22.22.192/24` to infinity as the next hop was R2.*
   - *R2 sends its updated routing information to R1*
   - *R1 adopts the entry `22.22.3/24`, cost infinity, because it had R2 as next hop of this route.*
   - *some time R1 will get an update from R3, saying that `22.22.3/24` is reachable from R3 at a cost of 2. This plus the cost of the link R1-R3 equals 12 which is smaller than the current cost which is infinity so it will be adopted.*
   - *R1 sends its updated routing information to R2, announcing among others that it has a route to `22.22.3/24` with a cost of 12.*
   - *This plus the cost of the link R1-R2 equals 13 which is smaller than the current cost which is infinity, so R2 will adopt it.*

   Give the new value of the routing table entry at each router for destination network = `22.22.3/24`, after the routing protocol has converged again. Give the values in the table below (do not give the value of the“interface” field).

Destination Network = `22.22.3/24`		
At Router	*Next-Hop*	*Distance*
R1	*22.22.1.2*	*12*
R2	*22.22.192.1*	*13*
R3	*22.22.2.4*	*2*
R4	*on link*	*1*
R5	*on link*	*1*

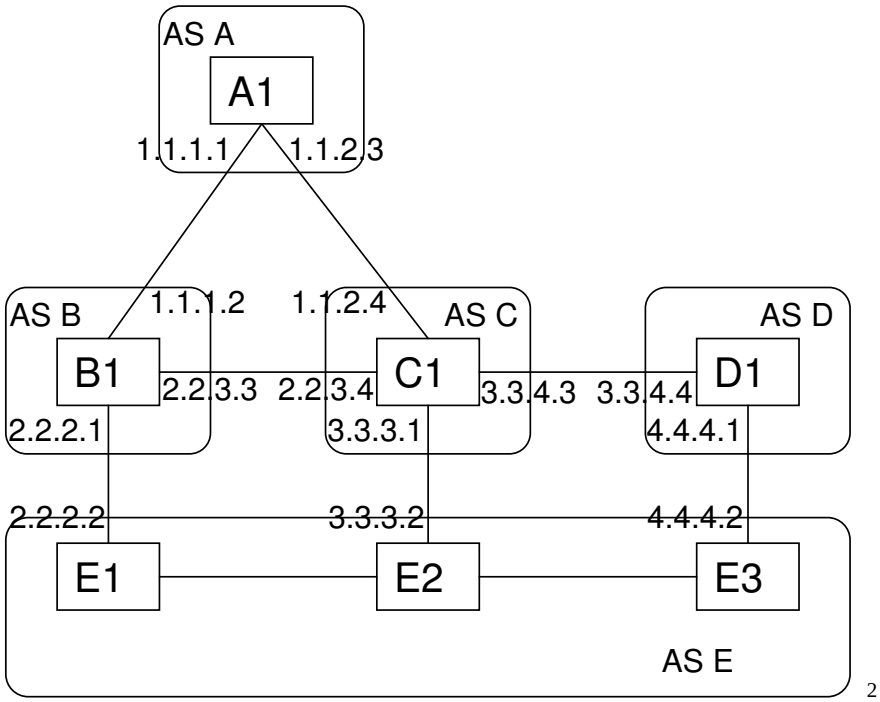

Figure 2: The network used in Question 2.

QUESTION 2

Consider the network in Figure 2. Boxes A1 to E3 are routers. There are five autonomous routing domains, each shown by a rounded box on the figure. The AS numbers are A to E.

The physical connections are shown by the lines in the figure. There is no other router or prefix than shown on the figure or introduced in the text.

All routers run RIP inside their domain. *All* routers run BGP. There is *no* redistribution of BGP into RIP.

There are external BGP sessions between routers as shown by the lines across domain boundaries (for example: A1-B1). There are internal BGP sessions between BGP routers, as many as required.

The decision process at all routers inside each AS is such that the route selected is, by order of decreasing priority

(1) the route that has the shortest sequence of ASs
(2) the route that has the shortest IGP distance from this router to the NEXT-HOP of the route
(3) the route that has the smallest IP address value in NEXT-HOP (here we consider an IP address as an integer)

There is no filtering, all routers accept all announcements.

1. Routers in E send the following announcements:

   E1 to B1: `66.66.0/17, AS path =E, NEXT-HOP=2.2.2.2`
   E2 to C1: `66.66.0/17, AS path =E, NEXT-HOP=3.3.3.2`
   E3 to D1: `66.66.0/17, AS path =E, NEXT-HOP=4.4.4.2`

After BGP has converged, what route to `66.66.0/17` will the decision process at C1 choose ? (give a detailed explanation of the steps involved in this). Same question for A1.

C1 will eventually also receive an announcement for `66.66.0/17` from B1, but its decision will not select it as its AS path is longer. Same for a possible announcement received from D1 or A1. So the decision process at C1 will continue to select the announcement sent by E2, namely `66.66.0/17, AS path =E, NEXT-HOP=3.3.3.2`.
A1 will receive two competing announcements from B1 and C1, with equal AS path lengths and IGP distances to next-hop. Therefore, the third criterion applies and A1 selects the announcement with smallest NEXT-HOP address, namely, `66.66.0/17, AS path =B E, NEXT-HOP=1.1.1.2`

2. The link between C1 and E2 fails. Explain what happens, say in particular which route to `66.66.0/17` router C1 will now select.

   (a) C1 and E2 will notice failure because of missing keepalive messages.
   (b) for C1 the route to `66.66.0/17` in loc-RIB is the one through E2, so it will delete it from loc-RIB
   (c) C1 thus has to rerun decision process. Because of criterion 3 the route through B1 will be promoted from RIB-in associated with B1 to loc-RIB, thus C1 will now select `66.66.0/17, AS path =B E, NEXT-HOP=2.2.3.3`.
   (d) as loc-RIB has been updated, C1 will propagate this new information to its BGP neighbors.

3. From now on we assume the link between C1 and E2 was repaired and BGP has converged again. Router A1 sends the following announcements.

   A1 to B1: `55.55/16, AS path =A, NEXT-HOP=1.1.1.1`

   A1 to C1: `55.55/16, AS path =A, NEXT-HOP=1.1.2.3`

   Explain which route to `55.55/16` the decision process at E1 will eventually choose. Same question at E2 and E3. Give a detailed explanation.

   (a) B1 and C1 will choose the announcement from A1 over all other announcements they might get later because it has the shortest AS path (criterion 1) and they will propagate this information to their BGP neighbours (B1 to E1 and C1; C1 to E2, B1 and D1).

   (b) D1 will get an announcement from C1 and choose this one again over all other announcements it might possibly get (criterion 1). It will also propagate this information to its neighbors (to E3).

   (c) After convergence, E1 applies the decision process to the following announcements:

   i. *from B1 via E-BGP: 55.55/16, AS path = B A, NEXT-HOP=2.2.2.1*
   ii. *from E2 via I-BGP: 55.55/16, AS path = C A, NEXT-HOP=3.3.3.1*
   iii. *from E3 nothing as the best route at E3 was learnt via I-BGP and is thus not repeated inside the AS.*

   Both announcements have same length of AS path (criterion 1), so the IGP distance (criterion 2) decides and the first one (via B1) is chosen.

   (d) After convergence, E2 applies the decision process to the following announcements:

   i. *from C1 via E-BGP: 55.55/16, AS path = C A, NEXT-HOP=3.3.3.1*
   ii. *from E1 via I-BGP: 55.55/16, AS path = B A, NEXT-HOP=2.2.2.1*
   iii. *from E3 nothing as the best route at E3 was learnt via I-BGP and is thus not repeated inside the AS.*

   Both announcements have same length of AS path (criterion 1), so the IGP distance (criterion 2) decides and the first one (via C1) is chosen.

   (e) After convergence, E3 applies the decision process to the following announcements:

   i. *from D1 via E-BGP: 55.55/16, AS path = D C A, NEXT-HOP=4.4.4.1*
   ii. *from E1 via I-BGP: 55.55/16, AS path = B A, NEXT-HOP=2.2.2.1*
   iii. *from E2 via I-BGP: 55.55/16, AS path = C A, NEXT-HOP=3.3.3.1*

   The first one is not chosen because its AS path is longer (criterion 1), from the remaining two the one with the shortest IGP distance (criterion 2) decides, and thus the last one (via E2) is chosen.

4. Router A1 now sends the following announcements.

   A1 to B1: `66.66/16, AS path =A, NEXT-HOP=1.1.1.1`

   A1 to C1: `66.66/16, AS path =A, NEXT-HOP=1.1.2.3`

   Explain which route to `66.66/16` the decision process at D1 will eventually choose.

   First of all note that `66.66/16` and `66.66.0/17` are different prefixes, so there will be one entry for each of them. D1 will choose the route through C1 (3.3.4.3) because this is the only one announced by C1 to D1 and the one announced by E3 to D1 must have a longer AS path.

   Assume an IP packet with destination address `66.66.254.1` is forwarded by router D1. To which neighbour does D1 send this packet ? Justify your answer. Same question for an IP packet with destination address `66.66.1.1`.

   for the fist one longest prefix match gives destination network `66.66/16` and it is thus sent to 3.3.4.3. for the second one, longest prefix match yields network `66.66.0/17` and it is thus sent to 4.4.4.2.