CSE 473S – Introduction to Computer Networks *Roch Guérin*

Final Exam

Your Name: *12/20/2017*

Six (6) Problems for a total of 120 points and a max of 100 points

(No points for answers without explanations/justifications)

1) **[25 points]** Consider an Ethernet network that consists of 9 switches, A to I, as shown below.

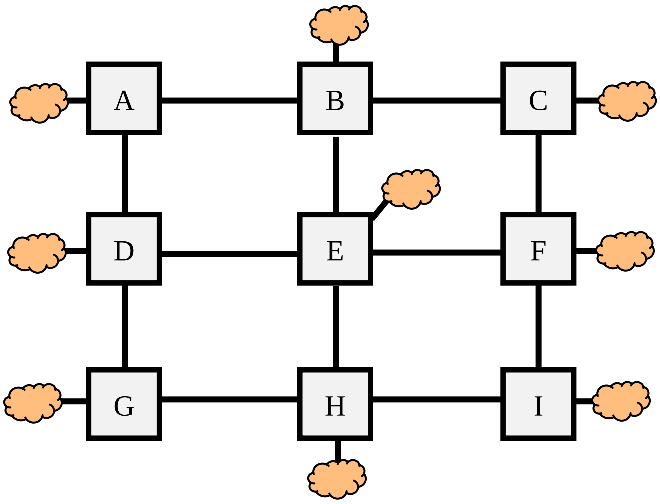

The little "clouds" correspond to local LAN segments, *i.e.*, LAN segments to which hosts are attached. Local LAN segments run at 10Gbps, while links between switches are 1Gbps links.

a. **[10 points]** Assume that switch priority is in decreasing alphabetical order, *i.e.*, switch A has the highest priority and switch I the lowest. Draw the resulting spanning tree, and for each switch indicate its root port (RP), designated port(s) (DP) and blocked port(s) (BP), if any.

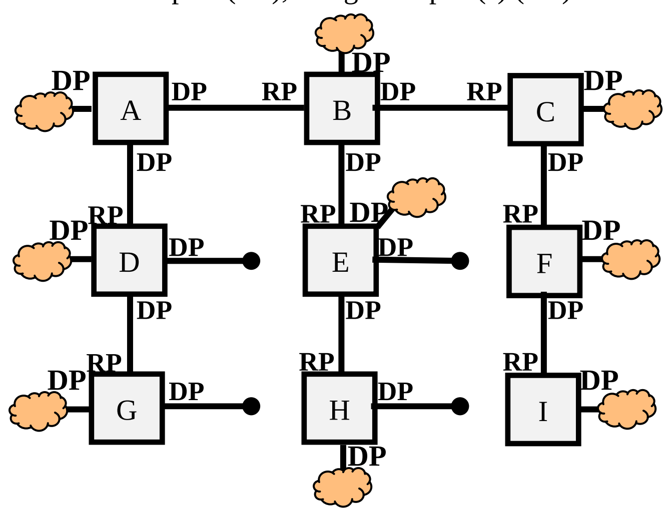

The spanning tree and port types at each bridge are identified in the above diagram, with switch A as the root bridge.

b. **[5 points]** Under the same switch priority assignments as in the previous question and assuming that switches transmit BPDUs every 100ms, how long would it take for the spanning tree to be constructed if all the switches were powered on at the same time? Justify your answer.

In each 100ms cycle, a switch learns about switches that are one hop further away. Since all the links have the same bandwidth, the shortest path between any two switches is the minimum hop path. The maximum hop count of the shortest path between any switch and switch A, the eventual root bridge, is 4, i.e., between switch A and I. This implies that the spanning tree will be finalized after 400ms or 4 cycles, as this is how long it takes for switch I to discover the root bridge A and its shortest path to it. All other switches take less time.

c. **[5 points]** Assume that servers are connected to local LAN segments attached to switches C, F, and I, with hosts on local LAN segment attached to switches A, D, G having heavy traffic to/from the server attached to switches C, F, and I, respectively, *e.g.*, hosts attached to switch A have heavy traffic to the server attached to switch C. Identify an assignment of priorities to switches that would ensure that traffic between hosts and their respective servers can flow on the most direct paths on the resulting spanning tree. **JUSTIFY YOUR ANSWER**, *e.g.*, identify the corresponding spanning tree and switch priorities, and why it results in the most direct paths.

The most direct paths between hosts and servers are A-B-C; D-E-F; G-H-I. In order to ensure that traffic flows on those paths, they need to be on the spanning tree. One (out of many) possible option is a spanning tree of the form shown below

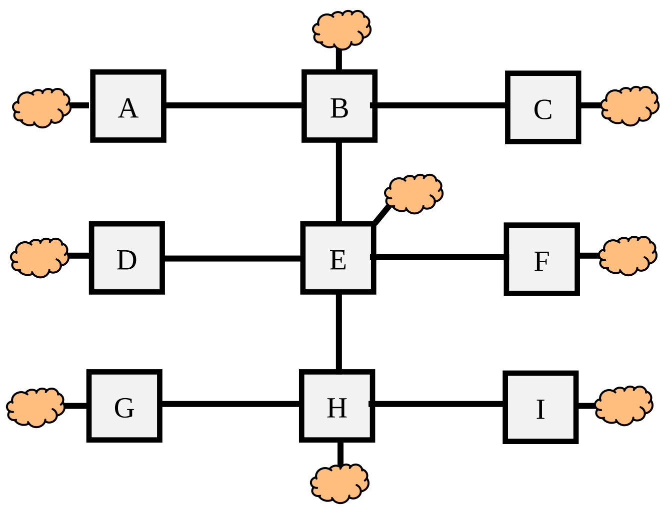

This can be realized by assigning the following priorities to switches (smaller is higher):
E: 1; B: 2; H: 3; and the other priorities are irrelevant, as long as they are lower than 3.

d. **[5 points]** Considering again the Ethernet network of question *a*, what is the maximum total traffic volume that this network could ever carry, and what traffic pattern does it correspond to, *i.e.*, identify the locations and traffic intensities of traffic source and sink hosts for a configuration that achieves this maximum traffic volume?

The maximum possible traffic volume is 8Gbps, since the spanning tree has a total of 8 switch links, each with a capacity of 1Gbps. This can be realized by having hosts attached to the local segment of a given switch send 1Gbps worth of traffic to hosts on the local LAN segment of a neighbor switch on the spanning tree. Note that because local LAN segments run at 10Gbps, they can source/sink traffic to/from multiple switch links. Assuming the spanning tree of question 1, a possible realization is as follows:

A->B; A->D; D->G; B->E; E->H; C->B; C->F; F->I.

2) **[10 points]** The network administrator of an office building has been allocated the following seven (7) subnets by its service provider:

(a) 192.166.10.0/25	(b) 192.166.10.129/25	(c) 192.166.11.0/24
(d) 192.166.12.0/22	(e) 192.166.8.0/24	(f) 192.166.9.0/24
(g) 192.166.16.0/21		

What is the <u>smallest</u> number of prefixes that the network administrator needs to advertise to the outside world to properly announce reachability to all its users, but no other? Explicitly identify the required prefixes in support of your answer.

The bit patterns for the last two bytes of the seven subnets that have been allocated to the network administrator are as follows:

*(a). 00001010.0*******,*

*(b). 00001010.1******* (with a mask size of 25, 129 is equivalent to 128)*

*(c). 00001011.*********

*(d). 000011**.*********

*(e). 00001000.*********

*(f). 00001001.*********

*(g). 00010***.*********

(a) and (b) can be combined into 192.166.10.0/24, which can then be combined with (c) to yield 192.166.10.0/23. Similarly, (e) and (f) can be combined to produce 192.166.8.0/23. The two /23 prefixes obtained from these aggregation steps can be in turn combined into 192.166.8.0/22, which together with (d) produces the supernet 192.166.8.0/21. This supernet can unfortunately not be combined with the remaining supernet 192.166.16.0/21, so that our network administrator needs to advertise the two /21 prefixes to ensure that it properly announces reachability to all its users.

3) **[25 points]** Consider three traffic sources, A, B, and C (see below), that are feeding a common multiplexer M connected to a shared 100 Mbps link. All three sources generate fixed size, 10,000 bits packets according to the patterns described in Table 1, with the first packet arriving at M at t=0. There is no packet processing delay in M, so that any delay that packets experience is caused by queueing delays waiting for the transmission of other packets.

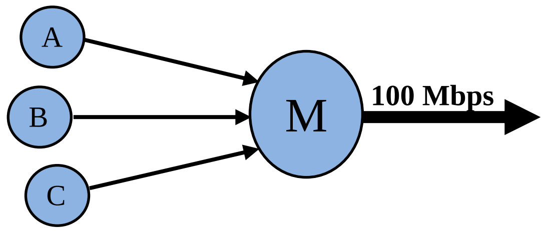

Source	Packet generation times (in units of 100µsecs)
A	A1=0, A2=1, A3=2, A4=4, A5=5
B	B1=0, B2=2, B3=3, B4=4, B5=5, B6=6, B7=7, B8=8
C	C1=2.5, C2=5.5, C3=8.5

Table 1: Packet arrival times

a. **[10 points]** Assume that the queueing discipline for the Ethernet link at M is FIFO, and that packets <u>that arrive at the same time</u> are inserted in the FIFO queue in the alphabetical order of their respective source, *e.g,* if packets from sources A, B, and C arrive together, A will be inserted first in the FIFO queue, followed by B, and then C. What is the maximum <u>delay</u> (time between the packet arrival and when it starts transmission on the link) experienced by any packet across all three sources? Justify your answer by showing how you compute this maximum delay value.

The 100Mbps link is capable of transmitting one 10,000 bits packet every 100µsecs. If we track packet arrivals and departures at M after time 0, we see that the link remains always busy until the last one of the packets listed in Table 1 has been transmitted. This is because at any point in time during that period, more bits have arrived than can have been sent. Specifically, we have the following progression in terms of number of packets received, corresponding backlog, and associated delay:

Time	*0*	*1*	*2*	*2.5*	*3*	*4*	*5*	*5.5*	*6*	*7*	*8*	*8.5*
Arrivals	*2 A1,B1*	*1 A2*	*2 A3,B2*	*1 C1*	*1 B3*	*2 A4,B4*	*2 A5,B5*	*1 C2*	*1 B6*	*1 B7*	*1 B8*	*1 C3*
Backlog	*0*	*1*	*1*	*2.5*	*3*	*3*	*4*	*5.5*	*6*	*6*	*6*	*6.5*
Delay	*1*	*1*	*2*	*2.5*	*3*	*4*	*5*	*5.5*	*6*	*6*	*6*	*6.5*

Given this, we see that the highest delay is equal to 0.65ms and is experienced by packet C3 (it sees the highest backlog upon arrival).

b. **[10 points]** The multiplexer now implements a Round-Robin (RR) scheduling policy with separate queues for A, B, and C, which are served in alphabetical order one packet at the time. If multiple packets arrive simultaneously and the multiplexer is idle, it starts serving queues in alphabetical order. Under this assumption, what is the delay experienced by packet C3? Justify your answer by identifying the sequence of packets transmitted up to packet C3.

Using the table from the previous question, we readily see that the order of transmission is as follows: A1(0), B1(1), A2(2), B2(3), C1(4), A3(5), B3(6), C2(7), A4(8), B4(9), C3(10), A5(11), B5(12), B6(13), B7(14), B8(15), where the number in parenthesis next to each packet is the time when the packet transmission starts. Based on this information, packet C3 arrives at t=8.5, and starts service at t=10, for a difference of 1.5, so that its delay is 150μsecs.

c. **[5 points]** Under the RR scheduler, when does the last of the packets specified in Table 1 finish transmission on the link?

Because the RR scheduler is work-conserving, the last packet starts transmission at the same time as in the case of a FIFO scheduler, and only the order of packet transmissions is different.

*The multiplexer receives a total of 16 packets, all 10,000 bits long. Because the link is always busy after time t=0, it takes 16*100μsecs to transmit those packets. As a result, the last packet finishes transmission at t=1.6ms.*

4) **[20 points]** Consider the OSPF network below that consists of 8 routers, R1 to R8 and one Ethernet switch T, which serves as a transit network inter-connecting routers R2 to R7. Router R*i*, *i*=1,…,8, has IP address 10.0.*i*.1 and is also connected to local subnets 10.0.*i*.0/24 and 10.*i*.0.0/16. Transit network T maps to subnet 10.128.0.0/24. Link costs are as shown on the diagram below, and the cost to local subnets are set to 1 on all routers.

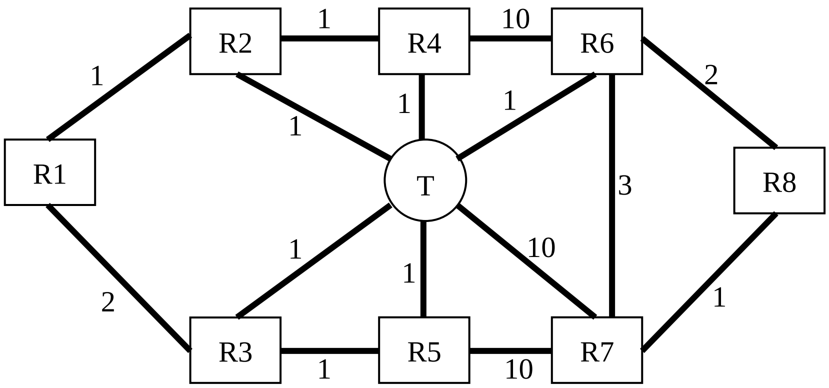

a. **[15 points]** Compute the routing table at router R2. Rows in the routing table should have the following format: <subnet, next_hop(s), cost>.

   If and when one of the next hops is to be reached through transit network T, identify it as T->R*i*. Justify your answer by identifying all intermediate shortest path computation steps involved in computing the routing table. You can assume that Flooding of LSAs has completed, so that router R2 has access to a full "map" of the network.

Computation of the routing table proceeds in two steps. Router R2 first relies on the Dijkstra algorithm to compute shortest paths to all other routers and transit networks in the network. Dijkstra proceeds by moving one node from the set of candidate nodes to the set of labeled nodes (nodes with a shortest path) at each step, as follows:

- *L={<R2,loc,0>};C={<R1,R1,1>,<T,T,1>,<R4,R4,1>}*
- *L={<R2,loc,0>,<T,T,1>};C={<R1,R1,1.,<R3,T->R3,1>,<R4,(T,R4),1>,*
  <R5,T-R5,1>,<R6,T->R6,1>,<R7,T->R7,1>}
- *L={<R2,loc,0>,<R1,R1,1>,<T,T,1>};C={<R3,T->R3,1>,*
  <R4,(T->R4,R4),1>,<R5,T->R5,1>,<R6,T->R6,1>, <R7,T->R7,1>}
- *L={<R2,loc,0>,<R1,R1,1>,<T,T,1>,<R3,T->R3,1>};*
  C={<R4,(T->R4,R4),1>,<R5,T->R5,1>,<R6,T->R6,1>, <R7,T->R7,1>}
- *L={<R2,loc,0>,<R1,R1,1>,<T,T,1>,<R3,T->R3,1>,<R4,(T->R4,R4),1>};*
  C={<R5,T->R5,1>,<R6,T->R6,1>, <R7,T->R7,1>}
- *L={<R2,loc,0>,<R1,R1,1>,<T,T,1>,<R3,T->R3,1>,<R4,(T->R4,R4),1>,*
  <R5,T->R5,1>};C={<R6,T->R6,1>, <R7,T->R7,1>}
- *L={<R2,loc,0>,<R1,R1,1>,<T,T,1>,<R3,T->R3,1>,<R4,(T->R4,R4),1>,<R5,T-R5,1>,*
  <R6,T->R6,1>}; C={<R8,T-R6,3>,<R7,T->R7,1>}

- *L={<R2,loc,0>,<R1,R1,1>,<T,T,1>,<R3,T->R3,1>,<R4,(T->R4,R4),1>,<R5,T->R5,1>, <R6,T->R6,1>,<R7,T->R7,1>};C={<R8,T->R7,2>}*
- *L={<R2,loc,0>,<R1,R1,1>,<T,T,1>,<R3,T->R3,1>,<R4,(T->R4,R4),1>,<R5,T->R5,1>, <R6,T->R6,1>,<R7,T->R7,1>, <R8,T->R7,2>}*

Where entries in the labeled and candidate node sets are of the form <destination, next hop(s), cost>, and we have used the fact that in cases of ties, transit networks always need to be moved first to the set of labeled nodes (recall that this is because their outgoing link cost is ***zero****).*

Once shortest paths are known to all, subnets are added as a post-processing step. The final routing table at R2 is of the following form:

Subnet	*Next hop(s)*	*cost*
10.0.1.0/24	*R1*	*2*
10.1.0.0/16	*R1*	*2*
10.0.2.0/24	*Local*	*1*
10.2.0.0/16	*Local*	*1*
10.0.3.0/24	*T->R3*	*2*
10.3.0.0/16	*T->R3*	*2*
10.0.4.0/24	*R4,T->R4*	*2*
10.4.0.0/16	*R4,T->R4*	*2*
10.0.5.0/24	*T->R5*	*2*
10.5.0.0/16	*T->R5*	*2*
10.0.6.0/24	*T->R6*	*2*
10.6.0.0/16	*T->R6*	*2*
10.0.7.0/24	*T->R7*	*2*
10.7.0.0/16	*T->R7*	*2*
10.0.8.0/24	*T->R7*	*3*
10.8.0.0/16	*T->R7*	*3*
10.128.0.0/24	*T*	*1*

b. **[5 points]** The link R6-R7 goes down. How is this detected, if at all, by router R2, and how does it affect its routing table? Justify your answer.

Router R2 is notified of the failure by receiving new Router_LSAs originated by both routers R6 and R7 that indicate that their link to the other router is now gone. Routers R6 and R7 would typically detect that the link went down through their Hello protocol.

Upon receiving the updated Router_LSAs from R6 and R7, router R2 proceeds to re-run its Dijkstra algorithm. However, because the failed link was not part of any shortest paths, the new routing table is identical to the previous one.

5) **[20 points]** Consider the corporate network below with three routers, R1, R2, and R3, which are interconnected using an Ethernet networks consisting of three VLANs numbered 1, 2 and 3, with links labeled according to the VLAN(s) they belong to. VLAN 1 is associated with subnet 11.1.0.0/16, VLAN 2 with subnet 11.2.0.0/16, and VLAN 3 with subnet 11.3.0.0/16. In addition, router R2 provides connectivity to the rest of the Internet and consequently advertises a default route 0.0.0.0/0 that it itself receives from its ISP. Finally, router R1 provides connectivity to the rest of the corporate network, and therefore also advertises connectivity to 11.0.0.0/8.

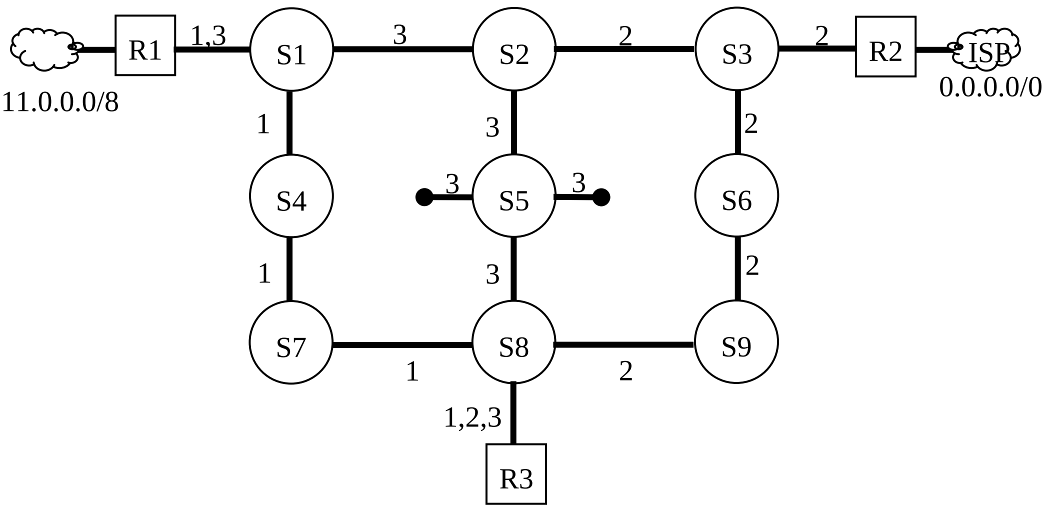

a. **[5 points]** Consider host 11.3.25.12 that has just booted and that issues a DNS query to DNS server 53.125.45.87. Identifies which Ethernet switches in the corporate network should contain a forwarding entry for the MAC address of host 11.3.25.12 and that of the DNS server as a result of its DNS query. Justify your answer.

Because the DNS server is not in the subnet of host 11.3.25.12, it will need to send the query to its default gateway, which is router R3. As a result, the host will issue an ARP query for R3 on VLAN 3. This will result in the MAC address of host 11.3.25.12 being present in the forwarding tables of all the switches that have a link in VLAN 3, i.e., S1, S2, S5 and S8. Furthermore, since the DNS server is not part of any subnet of the corporate network, no Ethernet switches will contain an entry for its MAC address.

b. **[5 points]** Assume that the three routers are running a standard IGP such as OSPF or EIGRP. What routes are present in the routing table of each router?

All three routers have the same set of routes in their respective routing tables, namely:

0.0.0.0/0, 11.0.0.0/8, 11.1.0.0/16, 11.2.0.0/16, and 11.3.0.0/16.

c. **[10 points]** Host 11.23.34.56 starts an `ftp` connection to host 150.12.58.61. Identify the set of switches and routers from the above network that the packets will traverse.

Host 11.23.34.56 is in the segment of the corporate network connected through router R1, so that the packet will arrive at R1. The best match for IP address 150.12.58.61 is the default route 0.0.0.0/0 advertised by router R2. However, R1 is not directly connected to R2 (they don't have any subnets in common), so that it needs to first forward the packet to R3, which will in turn deliver it to R2. This means that the full path followed by the packets through the corporate network is as follows:

R1-S1-S4-S7-S8-R3-S8-S9-S6-S3-R2

Note that because R1 and R3 are on both VLAN 1 and VLAN 3, an alternate path would be

R1-S1-S2-S5-S8-R3-S8-S9-S6-S3-R2

6) **[20 points]** Consider AS 1234 comprised of 4 routers that run iBGP among themselves and that each connect to other routers in other ASes using eBGP. A logical representation of the AS is shown below with the numbers next to each link indicating the shortest paths between pairs of routers, as also shown in the table on the left. Router R*i* has BGP identifier 10.0.0.*i*.

	R1	R2	R3	R4
R1	0	10	8	10
R2	10	0	11	5
R3	8	11	0	7
R4	10	5	7	0

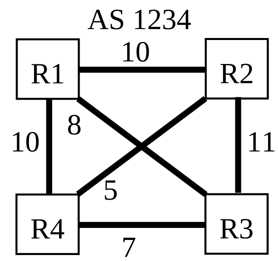

The four routers advertise to their iBGP peers the following routes that are all learned from their eBGP peers.

R1	0.0.0.0/0, AS_PATH: <32>, LOCAL_PREF=200 153.0.0.0/8, AS_PATH: <32-1235-6578>, LOCAL_PREF=50 153.10.0.0/16, AS_PATH: <32-1235-6578>, LOCAL_PREF=50 153.10.10.0/24, AS_PATH: <32-1235-6578>, LOCAL_PREF=50 153.10.11.0/24, AS_PATH: <32-1235-6578>, LOCAL_PREF=50 72.12.34.0/24, AS_PATH: <32-4444-3498-1278>, LOCAL_PREF=50
R2	153.0.0.0/8, AS_PATH: <22-5432-4333-6578>, LOCAL_PREF=100 153.10.0.0/16, AS_PATH: <22-1235-4333-6578>, LOCAL_PREF=50 72.12.34.0/24, AS_PATH: <22-1278>, LOCAL_PREF=50 72.12.0.0/16, AS_PATH: <22-3498-1278>, LOCAL_PREF=50 53.40.51.0/24, AS_PATH: <22-235-3343-5278>, LOCAL_PREF=100
R3	0.0.0.0/0, AS_PATH: <1>, LOCAL_PREF=100 153.0.0.0/8, AS_PATH: <1-6578>, LOCAL_PREF=50 153.10.11.0/24, AS_PATH: <1-6578>, LOCAL_PREF=50 72.12.34.0/24, AS_PATH: <1-3498-1278>, LOCAL_PREF=50 14.34.0.0/16, AS_PATH: <1-784>, LOCAL_PREF=50 14.34.10.0/24, AS_PATH: <1-784>, LOCAL_PREF=50
R4	153.0.0.0/8, AS_PATH: <11-459-6578>, LOCAL_PREF=50 153.10.11.0/24, AS_PATH: <11-459-6578>, LOCAL_PREF=60 72.12.34.0/24, AS_PATH: <11-3498-1278>, LOCAL_PREF=50 14.34.0.0/16, AS_PATH: <11-784>, LOCAL_PREF=50 14.34.10.0/24, AS_PATH: <11-784>, LOCAL_PREF=50

Assume that none of the routes include a MED attribute, and that all have the same ORIGIN value.

a) **[15 points]** Identify the routing tables at each of the four routers. For each route entry, identify the next hop router (use "*local*" when it is the router itself), and the reason for selecting the route. **No points will be awarded to answers without justifications.** For conciseness, give the routing table at R1, and for other routers only identify differences between their routing table and that of R1.

Recall from the BGP selection process that LOCAL_PREF has the highest precedence, followed by AS_PATH count, followed by eBGP over iBGP, followed by IGP cost, and finally followed by BGP lowest identifier. Based on this, the routing tables at the four routers are identical in the entries they contain, but differ in the NEXT_HOP they select as follows:

R1	*0.0.0.0/0, R1*	*chosen due to higher LOCAL_PREF of 200 for R1 vs. 100 for R3*
	153.0.0.0/8, R2	*chosen due to higher LOCAL_PREF of 100 for R2 vs. 50 for R1, R3 and R4*
	153.10.0.0/16, R1	*chosen as R1 and R2 have the same LOCAL_PREF, but R1 has a lower AS_PATH count of 3 vs. 4*
	153.10.10.0/24,R1	*chosen as R1 is the only one advertising that route*
	153.10.11.0/24, R4	*chosen due to higher LOCAL_PREF of 60 for R4 vs. 50 for R1 and R3*
	72.12.34.0/24, R2	*chosen as all routers have the same LOCAL_PREF value of 50, but R2 has the shortest AS_PATH count of 2*
	72.12.0.0/16, R2	*chosen as R2 is the only one advertising that route*
	53.40.51.0/24, R2	*chosen as R2 is the only one advertising that route*
	14.34.0.0/16, R3	*chosen as R3 has a lower IGP cost (both R3 and R4 have the same LOCAL_PREF and AS_PATH length)*
	14.34.10.0/24, R3	*chosen for the same reasons as 14.34.0.0/16*
R2	*14.34.0.0/16, R4*	*chosen as R4 has the lower IGP cost of 5 (both R3 and R4 have the same LOCAL_PREF and AS_PATH length)*
	14.34.10.0/24, R4	*chosen for the same reasons as 14.34.0.0/16*
R3	*14.34.0.0/16, R3*	*chosen as R3 has the lower IGP cost of 0 (both R3 and R4 have the same LOCAL_PREF and AS_PATH length)*
	14.34.10.0/24, R3	*chosen for the same reasons as 14.34.0.0/16*
R4	*14.34.0.0/16, R3*	*chosen as R4 has the lower IGP cost of 0 (both R3 and R4 have the same LOCAL_PREF and AS_PATH length)*
	14.34.10.0/24, R3	*chosen for the same reasons as 14.34.0.0/16*

b) **[5 points]** For each route known in AS 1234, with the exception of the default route 0.0.0.0/0, identify the AS to which the prefix belongs.

The AS to which the prefix belongs is typically the last one listed in the AS_PATH (for those interested, this is only true when the AS_PATH attribute is of type SEQUENCE and not SET. The latter is used in case of prefix aggregation, but none of the AS_PATH listed in the problem are of type SET). This means that for the nine (9) routes known in AS 1234, the owning ASes are as follows:

- *153.0.0.0/8, 153.10.0.0.16, 153.10.10.0/24, 153.10.11.0/24: AS 6578*
- *72.12.0.0/16, 72.12.34.0/24: AS 1278*
- *53.40.51.0.0/24: AS 5278*
- *14.34.0.0/16, 14.34.10.0/24: AS 784*