CSE 473 – Introduction to Computer Networks *Roch Guérin*

Final Exam Solution

Your Name: *12/17/2014*

PLEASE WRITE LEGIBLY – NO POINTS FOR ILLEGIBLE ANSWERS

1. **[10 points]** Bob has been provided with the following pair of encryption/decryption keys:

(2491,9);(2491,1329).

a) **[3 points]** Justify that Bob was given a valid RSA key pair (**Hint**: 2491 is divisible by 53).

*Knowing that 2491 is divisible by 53, we readily find that the other prime factor in 2491 is 47. Hence, the product (p-1)(q-1)=46*52=2392=2*23*2*2*13, so that the prime factors in (p-1)(q-1) are 2, 13, and 23. The encryption key e=9=3*3, therefore, satisfies the constraint that it be relatively prime with 2392. It remains to verify that 1253 is a valid decryption key, i.e., is of the form 9*1329=(k*2392+1) for some value k. Simple computations show that 9*1329=(5*2392+1)=11961. This establishes that Bob indeed received a valid RSA key.*

b) **[7 points]** Bob receives a message in the form of the number "2". What was the original message/number that was sent to Bob? (Hint 2^{50} mod 2491 = 2319. Taking a divide and conquer approach to perform the required computation will help).

Decrypting the received number 2 to recover the original value requires raising 2 to the power 1329 (the decryption key) and obtaining the remainder after dividing this number by 2491, i.e., performing a modulo 2491 operation.

*From the hint, we know that 2^{50} mod 2491 = 2319. We can use this result to progressively compute 2^{100} mod 2491 = (2319*2319) mod 2491 = 2183, 2^{200} mod 2491 = (2183*2183) mod 2491 = 206, 2^{400} mod 2491 = (206*206) mod 2491 = 89, 2^{800} mod 2491 = (89*89) mod 2491 = 448, 2^{1200} mod 2491 = (448*89) mod 2491 = 16, 2^{1300} mod 2491 = (16*2183) mod 2491 = 54, and finally, 2^{1329} mod 2491 = (54*628) mod 2491 = 1529.*
Hence the original message sent to Bob was 1529.

2. **[15 points]** A corporate network connects to the Internet over a duplex 100 Mbps (10^8 bits/sec) link.

a) **[4 points]** The 100 Mbps connection is shared by four distinct departments, which at peak hour have upload rates of UDP traffic equal to 25 Mbps, 50 Mbps, 75 Mbps, and 100 Mbps, respectively. What fraction of the 100 Mbps Internet link bandwidth would each department approximately get if their traffic shared a single queue?

*Each flow gets a fraction of the link bandwidth proportional to its own rate divided by the sum of the upload rate. Hence, the department with a 25 Mbps upload bandwidth gets 25*100/(25+50+75+100) = 10 Mbps, the department with a 50 Mbps upload bandwidth gets 50*100/(25+50+75+100) = 20 Mbps, the department with a 75 Mbps upload bandwidth gets 75*100/(25+50+75+100) = 30 Mbps, while the department with a 100 Mbps upload bandwidth gets 100*100/(25+50+75+100) = 40Mbps.*

b) **[2 points]** Assume now that each department is assigned its own queue, where each queue is served according to a weighted fair queueing (WFQ) scheduler with equal weights for each queue. Under this assumption, what bandwidth share does each department get?

Under the assumption of equal weights, each department gets up to 25 Mbps.

c) **[4 points]** How should the weights of the scheduler be set to ensure that each flow gets a share of the bandwidth that is proportional to its own rate?

Assigning weights of 0.25, 0.5, 0.75 and 1 to the 25 Mbps, 50 Mbps, 75 Mbps, and 100 Mbps flows, respectively, would achieve the desired outcome. Alternatively, normalized weights of 0.1, 0.2., 0.3 and 0.4 would also realize the same outcome.

d) **[5 points]** Identify a set of scheduler weights that would ensure that the departments with 25 Mbps and 50 Mbps get their upload demand fully satisfied, while the other two departments would share the remaining bandwidth equally.

Assigning weights of 0.25 and 0.5 to the departments with upload bandwidth requirements of 25 Mbps and 50 Mbps, respectively, would ensure their traffic demands are fully met. The other two departments would be assigned weights of 0.125 each.

3. **[15 points]** The network below consists of two routers, R1 and R2, and seven Ethernet switches, A to G. The switches are configured with three VLANs and the labels next to each link show the VLANs active on the link. Some links are in multiple VLANs. VLAN 1 is assigned subnet 1.0.0.0/8, VLAN 2 subnet 2.0.0.0/8 and VLAN 3 subnet 3.0.0.0/8. Note that each VLAN forms a spanning tree to start with. R1 and R2 each belong to two subnets and can send/receive packets on the corresponding VLANs. They are the default gateway for their subnet, *i.e.,* 2.0.0.0/8 for R1 and 3.0.0.0/8 for R2, with router R1 also the default gateway for the subnet they have in common, 1.0.0.0/8. In addition, R1 hosts the DNS server for the network, and offers connectivity to the Internet.

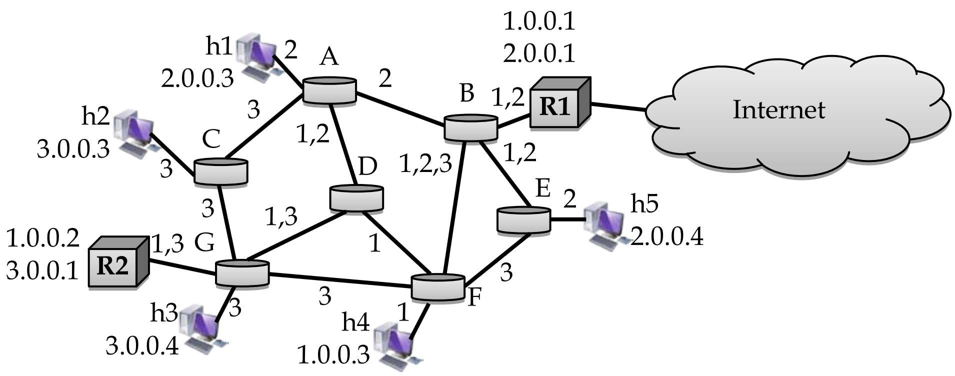

Five hosts, labeled h1 to h5, are shown in the diagram. Hosts belong to the VLAN of their IP subnet.

a) **[5 points]** Host h1 boots up with an empty ARP cache, and needs to communicate with IP addresses: 2.0.0.1, 1.0.0.1, and 5.1.1.1. Identify what MAC addresses are present in which switch after the connections have been established. Your answer should be a list of entries of the form: Switch X: MAC address of IP x, MAC address of IP y, … , etc.

Host h1 and IP address 2.0.0.1 are in the same subnet, so h1 broadcasts an ARP query on VLAN2 and the response comes back directly to it. All switches in VLAN2 create a forwarding entry for the MAC address of 2.0.0.3, and switches A and B create an entry for the MAC address of 2.0.0.1.
IP address 1.0.0.1 is not in the same subnet as host h1. Hence h1 broadcasts an ARP query for its gateway router R1, with the response again coming back directly to it. The switches in VLAN2 already have a forwarding entry for the MAC address of 2.0.0.3, but switches A and B create an entry for the MAC address of 1.0.0.1. Because, IP address 1.0.0.1 is that of R1, no additional ARP query (by R1) is required to establish the connection.
IP address 5.1.1.1 is again not in the subnet of host h1, but it already has the MAC address of R1. Hence, no new ARP query is required, and no additional MAC entries are created in the switches.
The forwarding entries created in the switches are therefore as follows:
Switch A: MAC address of 2.0.0.3, MAC address of 2.0.0.1, MAC address of 1.0.0.1
Switch B: MAC address of 2.0.0.3, MAC address of 2.0.0.1, MAC address of 1.0.0.1
Switch D: MAC address of 2.0.0.3
Switch E: MAC address of 2.0.0.3
Switch F: MAC address of 2.0.0.3.
None of the other switches create any forwarding entry because of the connections.
Note that because IP addresses 2.0.0.1 and 1.0.0.1 are both associated with router R1, they may or may not correspond to different MAC addresses.

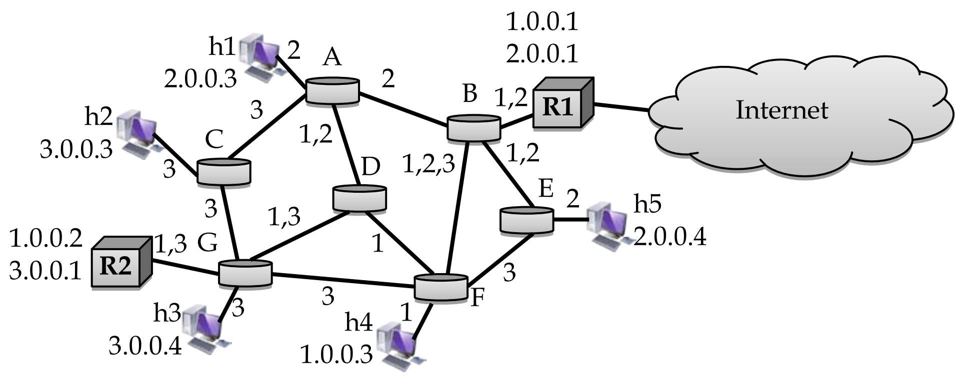

b) **[3 points]** Host h3 has an established communication with host h5. What sequence of switches and routers does it go through? Explain why and list them in order, repeating switches the packet traverses more than once, if any.

Hosts h3 and h5 are in different subnets with R2 serving h3 and R1 serving h5. The path followed by the connection's packets is, therefore: G,R2,G,D,F,B,R1,B,E.

c) **[3 points]** While h3 has a connection to h5, h2 has one to h4, and h1 has a connection to a remote host on the Internet. Assuming that all links are duplex 1 Gbps (10^9 bits/sec), except for link B-R1 that is 10 Gbps and that there is no other meaningful traffic on the network and no other limitation except the Ethernet network itself, what is the maximum bandwidth <u>each connection</u> can get? Justify your answer.

The connection from h2 to h4 follows the path C,G,R2,G,D,F while the connection between h1 and the Internet follows the path A, B,R1. Hence, the links G,R2; R2,G; G,D; and D,F are common to connections h3-h5 and h2-h4. This means that each connection gets approximately 1/2 of the link bandwidth or 500 Mbps. Conversely, link B-R1 is common to connections h3-h5 and h1-Internet, but because its capacity is 10 Gbps it is not a bottleneck link and the connection h1-Internet has a maximum bandwidth of 1 Gbps.

d) **[4 points]** Can you suggest a single change in either VLAN topology or assignment of IP addresses that would allow each connection to reach a maximum bandwidth of 1 Gbps? If yes, what is it? If no, why not?

Consider changing the IP address of h4 to 3.0.0.5. This results in h2 and h4 being on the same subnet, so that the path for connection h2-h4 is now C,G,F. This is disjoint from the other two paths. The path between h3 and h5 and that between h1 and the Internet only share the link between B and R1 that is 10 Gbps. Hence all three paths are now link disjoint in the Ethernet network itself, so that the maximum possible bandwidth of each connection is 1 Gbps.

4. [15 points] Consider the WiFi network shown below with two access points, X and Y, and four active hosts, A, B, C, and D, with A and B associated with X and C and D associated with Y. The two access points are operating on the same channel. Host A can hear host B, but not hosts C and D. Host B can hear hosts A and C, but not host D. Host C can hear hosts B and D, but not host A. Host D can hear host C, but not hosts A and B. All four hosts can hear and be heard by both access points X and Y, which can also hear each other.

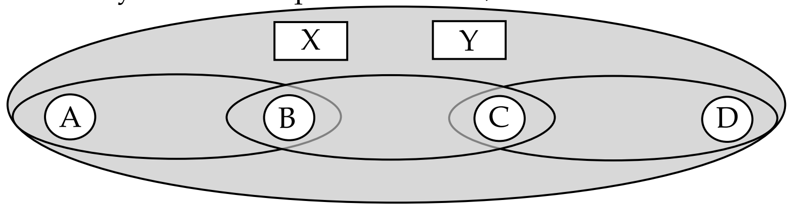

a) **[7 points]** Assume that
At t=0 access point X is transmitting and continues transmitting until t=400μs.
At t=50μs, host A gets a packet to send and initializes its backoff timer to 350μs.
At t=100μs, host B gets a packet to send and initializes its backoff timer to 300μs.
At t=150μs, host C gets a packet to send and initializes its backoff timer to 200μs.
At t= 350μs, host D gets a packet to send and initializes its backoff timer to 50μs.
All packets have a transmission time of 200μs, and RTS/CTS is not used.

Under those assumptions, identify when each host starts its packet transmission, and specify which packets are successfully transmitted in their first attempt. For simplicity, assume that inter-frame spacing times and times to send ACK packets are negligible.

All hosts keep their backoff timer frozen until X stops its transmission at t=400μs. Host D has the smallest backoff timer at, therefore starts transmitting at t=450μs. Host C is the only host that hears the transmission and, therefore, stops its backoff timer with a value of at 150μs. Both hosts A and B continue decrementing theirs. Host D finishes its transmission at t=650μs, at which point the value of host B's backoff timer is 50μs and that of A's backoff timer is 100μs. The backoff timer of host C is still at 150μs, but it now resumes decrementing it. The next timer to expire is that of B that starts transmitting at t=700μs. Both A and C hear B's transmission and, therefore, stop their backoff timers that are at values 50μs and 100μs, respectively. B finishes its transmission at t=900μs, at which point both A and C restart decrementing their backoff timers. A's timer expires first at t=950μs at which point A starts transmitting. Unfortunately, C does not hear A's transmission, so that is keeps decrementing its own backoff timer and eventually starts its own transmission at t=1000μs. This results in a collision with A's transmission so that neither are successful in their first attempt. In summary, packet transmissions and successes/failures are as follows:

Host	*Starts transmission at*	*Successful?*
A	*t=950μs*	*No*
B	*t=700μs*	*Yes*
C	*t=1000μs*	*No*
D	*t=450μs*	*Yes*

b) **[8 points]** Repeat the previous question assuming that RTS/CTS is now enabled for all packet transmissions. In this case, approximately when does each host send its data packet, and which ones are successfully received in their first attempt? <u>For simplicity, assume that the time needed to send RTS, CTS and ACK packets is negligible</u>.

Recall

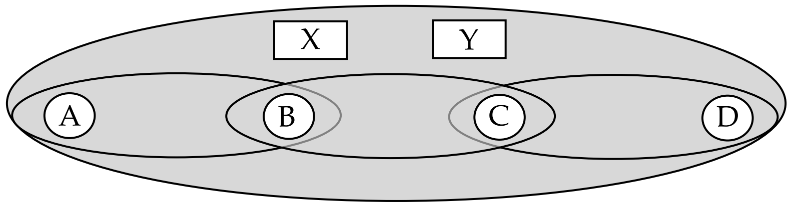

And that

At t=0 access point X is transmitting and continues transmitting until t=400µs.

At t=50µs, host A gets a packet to send and initializes its backoff timer to 350µs.

At t=100µs, host B gets a packet to send and initializes its backoff timer to 300µs.

At t=150µs, host C gets a packet to send and initializes its backoff timer to 200µs.

At t= 350µs, host D gets a packet to send and initializes its backoff timer to 50µs.

Host D sends its RTS at 450µs and gets granted permission to send until 650µs by Y. Host C hears the CTS, and therefore stops decrementing its backoff timer until t=650µs. Because both A and B are associated with X, they ignore the CTS from Y, and therefore continue decrementing their backoff timers, i.e., the use of RTS/CTS does not help in this case. Hence, when D finishes transmitting, the values of their backoff timers are as before equal to 100µs and 50µs, respectively, while C's backoff timer is still at 150µs. The next timer to expire is that of B that sends an RTS at t=700µs and immediately receives a CTS from X that grants it permission to transmit until t=900µs. A hears the CTS and, therefore, stop its backoff timer at 50µs, while C also stops decrementing its backoff timer because it hears the transmission from B. C's backoff timer, therefore, stays at 100µs while B is transmitting. B finishes its transmission at t=900µs, at which point both A and C restart decrementing their backoff timers. A's timer expires first at t=950µs, at which point it sends an RTS and immediately receives a CTS from X that grants it permission to transmit until t=1150µs. Unfortunately, because C is associated with Y, it does not pay attention to the CTS from X and, therefore continues decrementing its backoff timer that eventually expires at t=1000µs. This again results in a collision with the transmission from A, when C sends its RTS to Y (the transmission interferes with the X's reception of A's transmission). Hence, transmission times and transmission successes/failures are identical to those of the previous case.

Host	*Starts transmission at*	*Successful?*
A	*t=950µs*	*No*
B	*t=700µs*	*Yes*
C	*t=1000µs*	*No*
D	*t=450µs*	*Yes*

5. **[15 points]** Consider the network shown below that is configured as a three area OSPF network. As usual, area 0 is the backbone area that provides connectivity between areas 1 and 2. The routers F1 and G1 are two area border routers (ABRs) that are in both areas 0 and 1, and conversely routers F2 and G2 are two area border routers that are in both areas 0 and 2. All links have the same bandwidth and have been assigned an OSPF weight of 1.

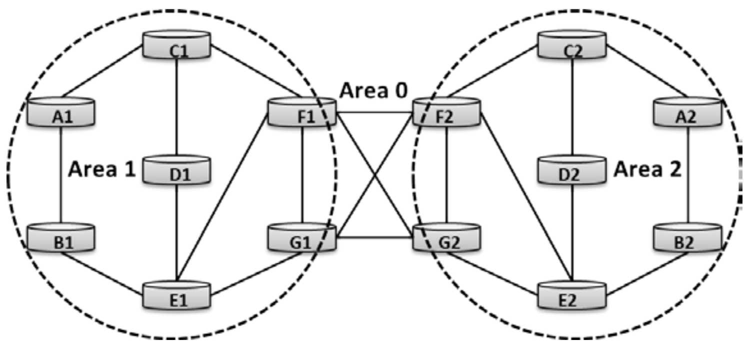

a) **[3 points]** Consider the routerLSA originated by router A2, and for <u>each</u> router in the network, identify the maximum number of copies it can get.

The routerLSA from router A2 is flooded only in area 2, and the maximum number of copies that a router can get is equal to its degree <u>in area 2</u>, i.e., B2 can get two copies, C2 can get three copies, D2 can get two copies, E2 can get four copies, F2 can get three copies, and G2 can get two copies. All other routers get zero copies.

b) **[5 points]** Assume next that router B2 advertises a route to subnet *r* in its routerLSA with a local cost of 1. What is the shortest path distance to *r* computed at router E1 together with its next hop(s)? Explicitly identify <u>**all**</u> steps involved in this computation, *i.e.*, how E1 learns about *r* and how it computes its shortest path and next hop(s) to *r*.

Routers F2 and G2 both compute shortest paths to router B2 using the Dijkstra algorithm and their knowledge of the internal topology of area 2. They will then advertise a T3-summary LSA for r in area 0 with a cost that is that of the shortest path to r. Routers F1 and G1 will receive the T3-summary LSA and compute their shortest path to r by adding the distances to F2 and G2 to the shortest path distance that each advertises for r. They will then flood a T3-summary LSA in area 1 announcing reachability to r with a distance that is that of their shortest path to r. Finally, router E1 will compute its shortest paths to F1 and G1 and determine its own shortest path to r by adding this distance to the shortest path distance that each advertises in their T3-summary LSA.

Specifically, F2 and G2 both have a distance of 2+1=3 to r, which is what they advertise in the T3-summary LSA flooded in area 0. Both F1 and G1 have a distance of 1 to F2 and G2, so that their shortest path to r has a total distance of 4, which is what they advertise in their T3-summary LSA in area 1. Router E1 receives the T3-summary LSAs and has a shortest path distance of 1 to both F1 and G1. Hence, E1 has a shortest path distance of 5 and next hops of F1 and G1.

c) **[3 points]** The link C2-D2 in area 2 fails. Which routers become aware of the fact that a change occurred in the network, and which don't? Justify your answer.

The failure or the link C2-D2 is detected trhough the Hello protocol between C2 and D2 and triggers the flooding in area 2 of updated routerLSAs from routers C2 and D2. The LSAs are received by all routers in area 2, which therefore become aware of the failure (change in topology) All area 2 routers then proceed to recompute shortest paths to destinations in area 2. This includes the two ABRs F2 and G2. However, the failure of the link does not change the minimum distance of any shortest path at either F2 or G2. Hence, they do not flood updated T3-summary LSAs in area 0, and therefore neither F1 nor G1 (nor for that matter any of the routers in area 1) ever become aware of the failure.

d) **[4 points]** Assume now that the following local subnets are included in the routerLSAs of the different routers in area 2, all with a local cost of 1:

- Router A2: 10.1.0.0/24
- Router B2: 10.1.1.0/24
- Router C2: 10.1.2.0/24
- Router D2: 10.1.4.0/24
- Router E2: 10.1.5.0/24
- Router F2: 10.1.3.0/24
- Router G2: 10.1.6.0/23

What is the smallest number of T3-summary LSAs that F2 and G2 need to advertise and what distance(s) do they advertise for each? Justify your answer.

The seven subnets can be aggregated into 10.1.0.0/21, so that F2 and G2 only need to advertise a single T3-summary LSA. The distance advertised in the T3-summary LSA will be the maximum of the distances to the individual subnets being aggregated.
This will, therefore, be a distance of 2+1=3 for F2 and 3+1=4 for G2 (G2 has a distance of 3 to A2).

6. **[10 points]** Consider the network below, where nodes are IP routers and the number next to each link is its OSPF link weight. The network is configured as a single OSPF area.

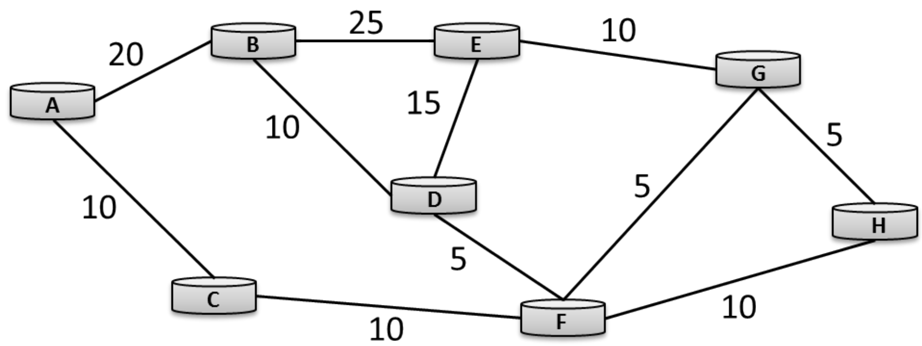

a) **[4 points]** Router A is the PIM DR for some layer 2 subnet, and host *X* in that subnet sends an IGMP report for address 226.1.1.1. Router G is the RP for 226.1.1.1, but no routers are currently participating in 226.1.1.1 (interested in packets for 226.1.1.1), and no host is actively sending packets to 226.1.1.1. What does router A do when receiving the IGMP report from X? Do any routers eventually add forwarding state for 226.1.1.1 as a result of A's action? If yes, which ones? If no, why not?

The receipt of the IGMP report from X triggers A to send a PIM Join packet on the shortest path towards the RP, namely to router C towards router G. Eventually, routers C, F and G would add forwarding state for 226.1.1.1. Note that this is independent of whether or not there are any host actively sending multicast packets to 226.1.1.1. A would also have forwarding state pointing to the local subnet of X.

b) **[3 points]** Assume next that host Y on a layer 2 subnet for which router D is the PIM DR sends an IGMP report for 226.1.1.1. What action would this trigger at router D, and would this result in the creation of any additional forwarding state?

The receipt of the IGMP report from Y triggers D to send a PIM Join packet on the shortest path towards the RP, namely to router F towards router G. This results in the creation of additional forwarding state at F (another branch towards D), but since the multicast tree is already in place at F, the PIM Join message would not be forwarded any further. In addition, as with A, D creates forwarding state towards the local subnet of Y.

c) **[3 points]** A host connected to a layer 2 subnet attached to router C starts sending multicast packets addressed to 226.1.1.1. Identify all the steps this would trigger to ensure the delivery of these packets to subscribers of the multicast group. In particular, identify all routers the packets will traverse and in which form, as well as any additional control messages this may trigger.

Because C is already on the multicast tree for 226.1.1.1, packets can be readily forwarded as native multicast packets. They will be forwarded from C to A and F, and from F to D and G. Once the packets arrive to D and A, the routers will forward them on the local subnets to which hosts X and Y belong. No additional (PIM) control messages are needed. G would simply discard the packets when they arrive, since there are no other tree branches on which to forward them.

7. **[15 points]** Consider the diagram below that displays connectivity between AS11 and several other ASes. AS11 uses OSPF as its internal routing protocol, and BGP to exchange routes with neighboring ASes. The four routers in AS11, A, B, C, and D, are connected by a full iBGP mesh.

AS11 has <u>peering relationships</u> with AS1, AS2 and AS4 and only advertises its own routes to them, while they advertise their own routes and those of their customers to AS11. AS11 is itself a <u>customer</u> of the two ASes labeled ISP1 and ISP2, which offer Internet connectivity to AS11 by advertising all the routes they know on their eBGP connections to routers D and C, respectively. Assume that the sets of routes learned from both ISP1 and ISP2 are identical.

The prefixes shown inside AS1, AS2, AS3, and AS4 identify the internal route they advertise to their eBGP neighbors. In addition, AS2 provides transit service to both AS1 and AS3, *i.e.,* they are customers of AS2. Hence, AS2 also advertises the routes it learns from AS1 and AS3 to AS11, as well as advertises to them the routes it learns from AS11. Finally, like AS11, AS4 is a customer of ISP1 that offers it Internet connectivity by advertising it all its routes.

AS11 assigns a LOCAL_PREF value of 10 to all routes learned by routers C and D over their eBGP connections to ISP2 and ISP1, respectively. It assigns a LOCAL_PREF value of 50 to routes learned by both routers A and D over their eBGP connections with AS1 and AS4, respectively, and a LOCAL_PREF value of 100 to routes learned by router B over its eBGP connection with AS2.

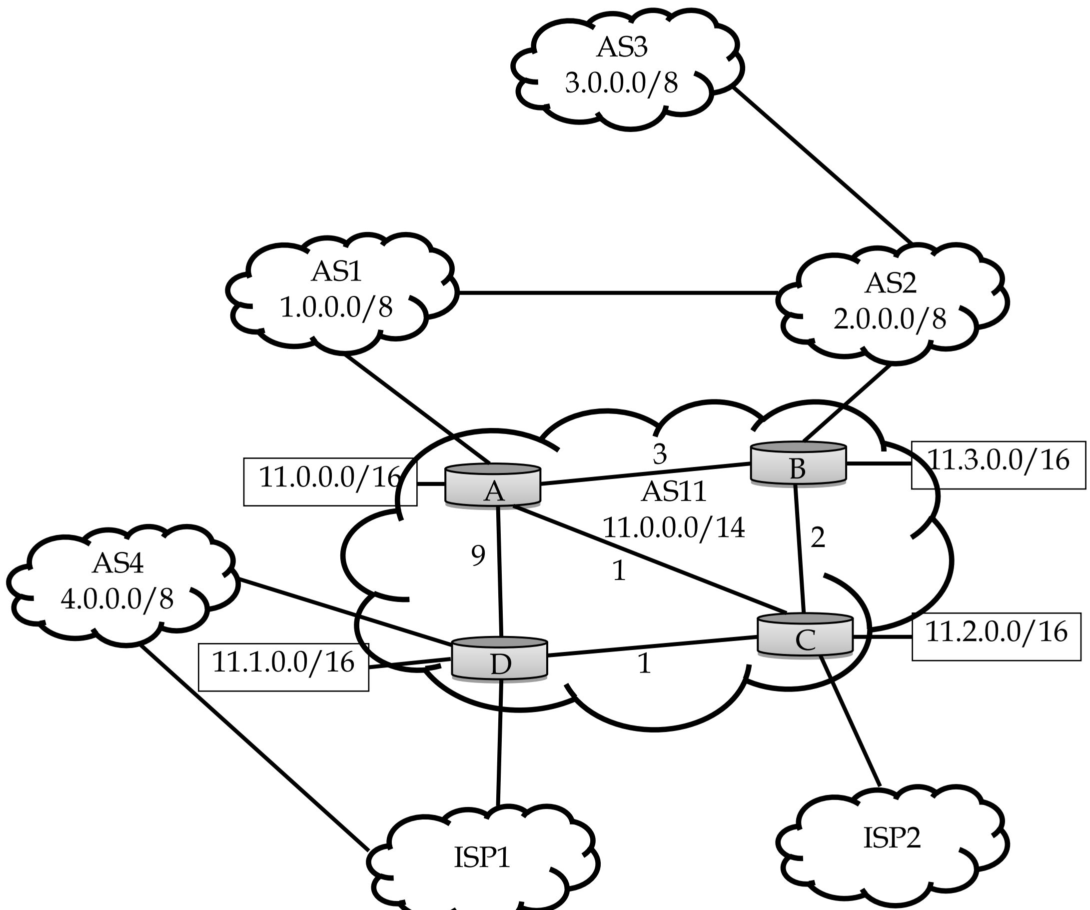

a) **[7 points]** Specify the routing table at **router B**, and for each entry identify the next hop(s) in the form of either a local connection, or connections to internal routers, or a neighboring AS, as applicable. For conciseness, you can represent the set of routes advertised by ISP1 and ISP2 (both advertise the same set) as "ISP". For each entry, explain your choice, *i.e.,* how and why it was selected by router B.

AS2 advertises 1.0.0.0/8, 2.0.0.0/8 and 3.0.0.0/8 to router B. Those routes have AS_PATHs of the form <AS1><AS2>, <AS3><AS2>, and <AS2>, respectively, and are assigned a LOCAL_PREF value of 100.
AS1 advertises 1.0.0.0/8 to router A. The route has an AS_PATH of the form <AS1> and is assigned a LOCAL_PREF value of 50.
AS4 advertises 4.0.0.0/8 to router D. The route has an AS_PATH of the form <AS4> and is assigned a LOCAL_PREF value of 50.
ISP1 advertises "ISP" to router D. AS_PATHs will vary, but LOCAL_PREF values are all 10.
ISP2 advertises "ISP" to router C. AS_PATHs will vary, but LOCAL_PREF values are all 10.

The routing table at B is, therefore, of the form:

Route	***Next hop(s) & Reason***
"ISP"	*C (lowest internal cost of 2 to either C or D that have equal LOCAL_PREF)*
11.0.0.0/16	*A and C (equal lowest internal costs of 3)*
11.1.0.0/16	*C (lowest internal cost of 3)*
11.2.0.0/9	*C (lowest internal cost of 2)*
11.3.0.0/16	*Local*
1.0.0.0/8	*AS2 (highest LOCAL_PREF that trumps the shorter AS_PATH of the route advertised by AS1 to A)*
2.0.0.0/8	*AS2 (highest LOCAL_PREF)*
3.0.0.0/8	*AS2 (highest LOCAL_PREF)*
4.0.0.0/8	*C (lowest internal cost of 3 to exit point, router D)*

b) **[3 points]** How would the routing table at B change, if at all, if the LOCAL_PREF value assigned to routes learned by A was changed from 50 to 100? Justify your answer.

The only change would affect the route 1.0.0.0/8, as the routes learned by routers A and B now have the same LOCAL_PREF value. The shorter AS_PATH of the route learned by A would now ensure that it becomes the preferred exit point to reach 1.0.0.0/8. Hence, the corresponding route entry in B's routing table would now have a next hop of A.

c) **[3 points]** Assume now that AS11 changes its relationship with ASes 1, 2, and 3, and becomes their provider, *i.e.,* offers them Internet connectivity. However, AS11 wants to ensure that all their traffic enters through router B, except for the traffic from AS1 to AS4 that should enter through router A. How would AS11 be able to realize such a goal? Justify your answer.

The easiest option to realize such a goal is to have router A advertise 4.0.0.0/8 to AS1 and have router B only advertise a default route 0.0.0.0/0 to AS2. This will ensure that the traffic from AS1 to AS4 enters through router A (the more specific prefix is preferred) and that all other traffic goes through AS2 and enters through router B. Note that AS1 would not advertise 4.0.0.0/8 to AS2 since it is a route it learned from a provider, and it also has a customer-provider relationship with AS2 (advertising 4.0.0.0/8 to AS2 would allow traffic destined for 4.0.0.0/8 to flow from AS2 through AS1, and consequently make AS1 into AS2's provider, at least for prefix 4.0.0.0/8).

d) **[2 points]** AS11 would like to ensure that traffic destined for the Internet and coming from AS2 (entering through router B) is distributed more or less equally between ISP1 and ISP2. Is this feasible simply by adjusting the OSPF weights of internal links? If yes, explain why. If not, justify why it is not feasible.

Unfortunately, because the BGP protocol only allows the selection of a <u>single</u> path, router B has to pick either ISP1 or ISP2 as its exit point to reach Internet destinations. Hence, no matter how OSPF weights are set, Internet traffic from B will only be forwarded to one exit point.

Note that internal weights can, however, be used to determine which exit point is picked by which internal router, when the internal path cost is the tie-breaker of the BGP decision process. This is not the case here since we have only one internal router to consider, namely, B.