CSE 473 – Introduction to Computer Networks *Roch Guérin*

Final Exam

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).

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).

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?

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?

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?

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.

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.

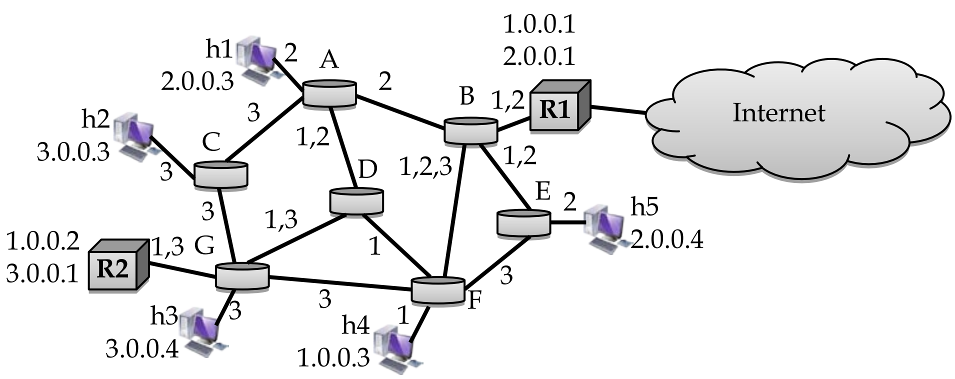

Five hosts, labeled h1 to h5, are shown in the diagram. Hosts belong to the VLAN of their IP subnet.

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 <u>after the connections have been established</u>. Your answer should be a list of entries of the form: Switch X: MAC address of IP x, MAC address of IP y, … , etc.

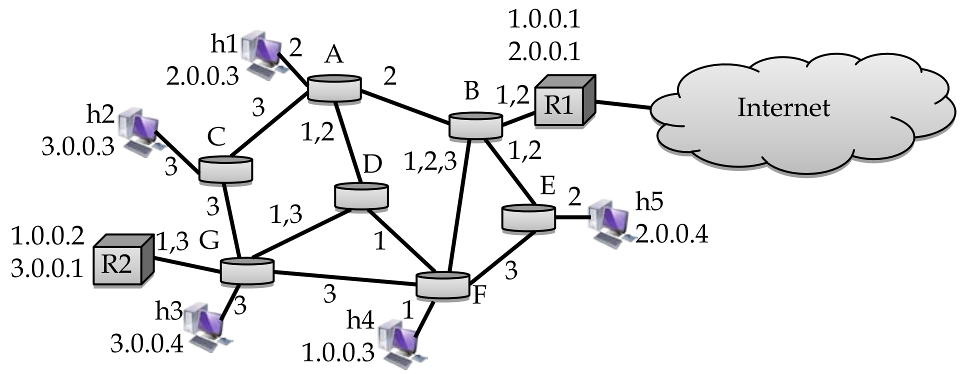

b) **[3 points]** Host h3 has an established communication with host h5. What sequence of switches and routers does it go through? Explain why and list them in order, repeating switches the packet traverses more than once, if any.

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.

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?

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.

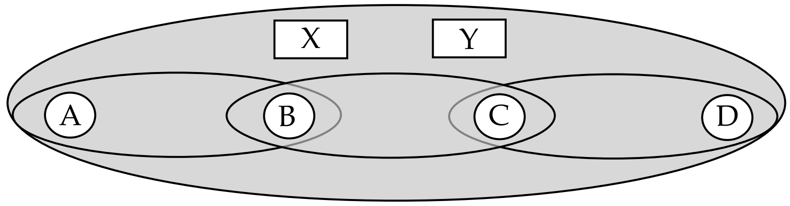

a) **[7 points]** Assume that
At t=0 access point X is transmitting and continues transmitting until t=400μs.
At t=50μs, host A gets a packet to send and initializes its backoff timer to 350μs.
At t=100μs, host B gets a packet to send and initializes its backoff timer to 300μs.
At t=150μs, host C gets 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.

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

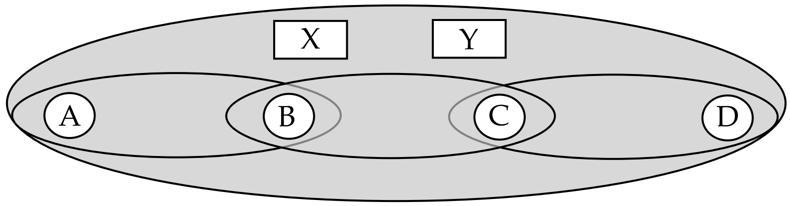

And that

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.

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.

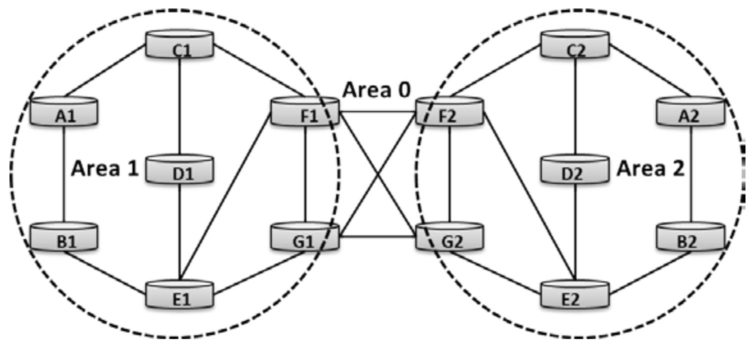

a) **[3 points]** Consider the routerLSA originated by router A2, and for <u>each</u> router in the network, identify the maximum number of copies it can get.

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*.

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.

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.

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.

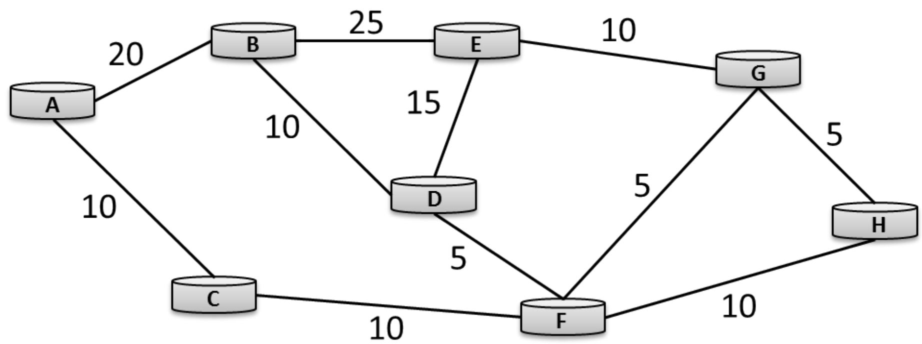

a) **[4 points]** Router A is the PIM DR for some layer 2 subnet, and host *X* in that subnet sends an IGMP report for address 226.1.1.1. Router G is the RP for 226.1.1.1, but no routers are currently participating in 226.1.1.1 (interested in packets for 226.1.1.1), and no host is actively sending packets 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?

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?

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.

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.

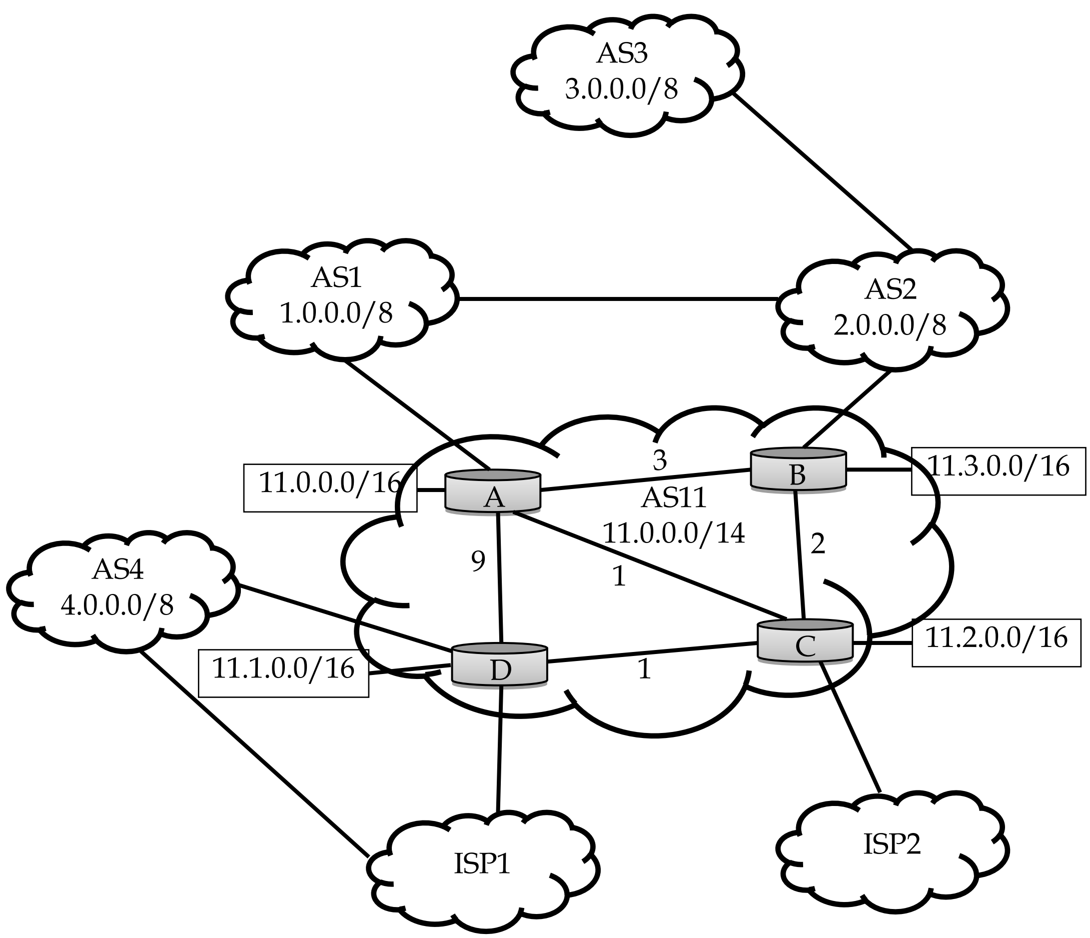

a) **[7 points]** Specify the routing table at **router B**, and for each entry identify the next hop(s) in the form of either a local connection, or connections to internal routers, or a neighboring AS, as applicable. For conciseness, you can represent the set of routes advertised by ISP1 and ISP2 (both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

Route	*Next hop(s) & Reason*

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.

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.

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.