Exam - Midterm

Due: 11:50am, 21 Oct 2014

Closed Book. Maximum points: 100

NAME:

1. Transmission on a link [20 pts] Did you write your name above? Please do! Now we can start: Your company has offices in Lisbon and Barcelona. The two cities are 1000 km (10^6m) apart. You currently have a link with bandwidth of 10Mbps (10^7bits/s). You currently have a straight-line dedicated fiber link between two computers, one at each office.

a. **What is the total time to send a 1250-byte packet and receive an acknowledgment?** Assume the speed of light in the fiber is 2×10^8m/s, and that the acknowledgment only incurs propagation delay (as if it were 0-bytes long). Draw a time diagram to help (two vertical lines with for each end, with time increasing from top to bottom). [5 pts]

b. Assume that the transmission protocol used offers reliability by retransmitting if an acknowledgment is not received. **Draw in two time diagrams two distinct situations that would result in the receiver getting two copies of the same packet**.[5 pts]

c. **What is the minimum size in bytes of a sliding window that fully utilizes this link?** [4 pts]

d. You have some budget to upgrade the link, and there are two options. One is to go with the same provider and upgrade the endpoints to be able to transmit at 100Mpbs, i.e., with 10x the bandwidth. The other is to change providers to one that uses microwave technology. While the latter offers the same bandwidth as the old link, the speed of microwaves over the air is 1.5× higher, 3×10^8m/s.

In order to help the company decide, you are tasked with a report of the impact of the link upgrade on two use cases: (i) the time to send one 1250-byte packet and wait for the ack – as in (a) above, and (ii) the time to send a 1.25MB (1,250,000B) file. For the file, assume a full pipeline using a sliding window and no losses. **What would be the total times for cases (i) and (ii), for the two links above (100Mpbs on fiber and 10Mbps over the air)?** (Write the 4 numbers for the answer in the table below) [6 pts]

	Fiber	Microwave
1250 B packet		
1.25MB file		

2. IP Addresses and Anycast [20 pts]

a. **Why does longest prefix matching work on IP addresses and not on Ethernet addresses?** [5 pts]

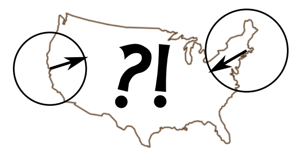

Google uses 8.8.8.8 as one of its public DNS servers, set up as an anycast address.

Before leaving for her internship in San Francisco, a Brown student measures the round trip time from her dorm to 8.8.8.8 to be 17ms, using `ping`. When she gets to her internship in Mountain view, she measures the RTT to 8.8.8.8 from there, and obtains 12ms. Puzzled, she measures the RTT between Mountain View and Brown, and finds it to be 113ms. She knows you are taking 168, so she asks you how this is possible. Are we violating some law of physics?

b. **What is your explanation to her of how this is possible?** [5 pts]

c. **How can Google implement this using BGP?** [5 pts]

d. **Are there any potential problems with using anycast?** [5 pts]

3. IP over Ethernet [20 pts] You want to set up your network as follows.

You want the IP addresses listed in the table. You also want {A, B, X} and {C,Y} to be two broadcast domains. S1 and S2 are learning switches. Assume that all tables are empty at first, except for the routing tables. These are set up correctly. Each of the nodes A, B, and C have only one entry in their routing tables.

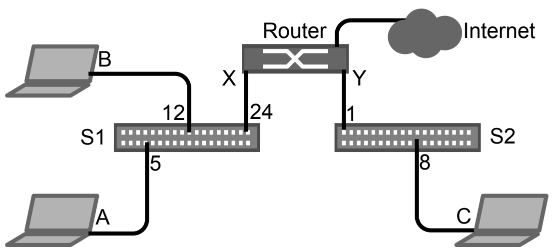

Interface	IP/mask	MAC
A	192.168.1.2/__	3e:eb:12:00:13:12
B	192.168.0.2/__	3e:eb:12:00:13:02
X	192.168.0.1/__	23:12:ef:12:01:02
C	192.168.2.2/__	3e:eb:12:00:14:1f
Y	192.168.2.1/__	23:12:ef:12:01:06

a. **Complete the network masks in the table using the CIDR prefix notation** (e.g., /24) so that the broadcast domains are respected and IP routing works. [4 pts]

b. Assume that the nodes A, B, C are now correctly configured with their network masks and gateways. **What is the only entry in A's routing table** (prefix and next hop)? [4 pts]

c. What entry(ies) does node A add to its **ARP cache** after it sends a UDP packet to B's IP address? Note that each entry has to have two components, as it is a mapping. [4 pts]

d. What entry(ies) does S1 have in its **learning switch table** after A's UDP packet to B above is sent? (The numbers close to switch 1 – 5, 12, 24 – are switch port numbers.) [4 pts]

e. What entry(ies) does node A add to its **ARP cache** after it sends a UDP packet to C's IP address? [4 pts]

4. Network Address Translation [15 + 8(bonus) pts]

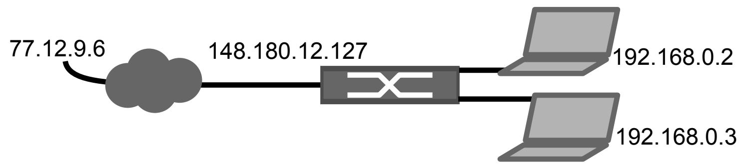

a. **Why does sending IP packets directly from 192.168.0.2 to 77.12.9.6 not work in the figure?** [5 pts]

b. **How does NAT solve the problem?** In this specific example, **list the address changes that NAT causes to an outgoing IP packet** from 192.168.0.2 to 77.12.9.6, **and to the corresponding response packet**. [5 pts]

c. FTP, the File Transfer Protocol, in its original method of working, broke with the presence of NATs. FTP uses two connections, one for control, and another for data. This used to be a popular pattern with old protocols. Without NAT, originally the client opened a control TCP connection to the external server on TCP port 21, with local port N. The client then sent, over this connection, its IP address and port N+1, as arguments to the PORT command. This would cause the server to initiate a data connection to the IP and port indicated. Data can flow in both directions in the data and control connections.

i. Why does this scheme not work with NAT? [5 pts]

(continued)

ii. [BONUS] How would you make FTP work with NAT? More points if your solution doesn't require the NAT to understand the FTP protocol! [8 pts]

5. BGP Routing [25 pts] The Gao-Rexford policies are a simplification of the BGP route propagation among ASes:

1. Routes learned from customers are exported to all neighbors
2. Routes learned from peers or providers are only exported to customers
3. Of the routes you know, prefer routes through customers

a. **What is the reasoning for rule #1?** [5 pts]

b. **What can happen if a multihomed customer violates rule #2?** [5 pts]

(continued)

The graph below represents the connections between ASes A through H. The single arrows link customer to providers (e.g., F is a customer of C) and the double arrows represent peering links (e.g., D is a peer of E).

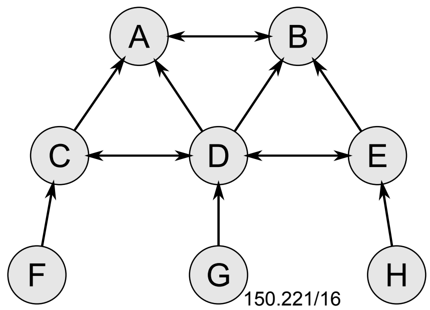

Assume that the ASes in the figure strictly follow the Gao-Rexford rules above, and keep them in mind to answer the following questions:

c. If G owns prefix 150.221/16, **what are all of the BGP announcements that A receives for this prefix?** [3 pts]

d. **What route (sequence of ASes) does a packet take from F to H? Why?** [3 pts]

e. **What route does a packet take from F to G? Why?** [3 pts]

f. If ASes prefer routes through customers, the route between A and H is A,D,E,H. **True or false? Why?** [3 pts]

g. If A and B got into a dispute and decided to not peer anymore, **which set of ASes would not be able to reach which other set of ASes? What one link (and type) would you add to fix this** (other than adding A-B to the graph)? [3 pts]