Exam - Midterm

Due: 7:00pm, 24 Oct 2016

Closed Book. Maximum points: 100

NAME:

1. Flying high [25 pts]

Did you write your name above? Please do! Now we can start:

You are designing the network for a company that wants to use satellites to provide Internet access.

These satellites will be initially at an altitude of 750km. Packets in your protocol are 1250 bytes in size, and the bandwidth you can get out of your equipment 100Mpbs, or 10^8 bits/s. (Make sure you don't mix bits with bytes).

a. **What is the two-way propagation delay in this link**, considering the speed of light to be 3×10^8 m/s? [4 pts]

$$\frac{2 \times 7.5 \times 10^5 m \cdot s}{3 \times 10^8 m} = 5 \times 10^{-3} s \text{ or } 5\ ms$$

b. If you use a stop-and-wait protocol, **what would be the average throughput you get**? (You can express it as a fraction of the link bandwidth). [5 pts]

Bytes/s: $\frac{1250B}{5 \times 10^{-3} s} = 2.5 \times 10^5 B/s$
bits/s: $10^4 b/5 \times 10^{-3} s = 2 \times 10^6 bps$
As a fraction: Maximum: $10^8 bps \times 10^-4 pkts/bit = 10^4 packets/second$. Actual: $1pkt/5 \times 10^{-3} s = 200 packets/s$.
$2 \times 10^2/10^4 = 2 \times 10^{-2} = 2\%$ of the link bandwidth.

c. You quickly realize that this is terrible, and decide to use a sliding window. **What is the best size of the sending window, *in packets*, to maximally utilize the link**?[5 pts]

Bandwidth × Delay = $10^8 bps \times 5 \times 10^{-3} s \times packets/10^4 bits$ = 50 packets.
This is how many packets fit in one round trip time.

d. When you are ready to do a git commit of your calculations, Sam and Xueyang from the orbital calculation department come to you and say that you will only have authorization to have the satellites orbiting at 2× the original distance. You immediately realize that your window calculations will have to change.

You were very proud of your modulation. You had used Shannon/Hartley's law to compute how many levels of modulation you could use ($M \leq \sqrt{1 + S/N}$), and sized your antenna to get $S/N = 1024$. It is too bad that the signal strength will reduce proportionally to the square of the distance. Given that you will likely have to adjust your modulation (and thus the bandwidth), and that the latency is also changing, **what is the new size of the window to fully utilize this new link**? [6 pts]

Call M, S, B, and D the original M, Signal, Bandwidth, and latency, and M', S', B', D' their respective values in the new setting.
$M \leq \sqrt{1 + 1024} \approx 32 \rightarrow 5$ levels.
$M' \leq \sqrt{1 + 1024/4} = \sqrt{257} \approx 16 \rightarrow 4$ levels.
So, $B' = \frac{4}{5}B$.
Now, D' = 2D.
So, the new window $W' = B' \times D' = \frac{4}{5}B \times 2D = \frac{8}{5}B \times D = \frac{8}{5}W$ = 80 packets.

2. MAC Unlearning [25 pts]

Consider the following network.

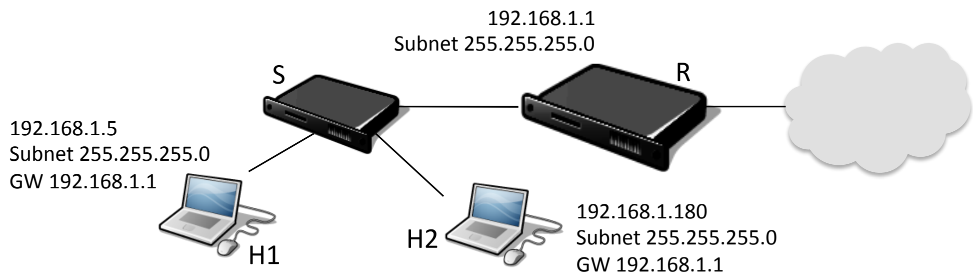

a. Suppose H2 wants to snoop on all communications between H1 and some other host S, without disrupting the communication. H2 sets its interface to promiscuous mode, but, alas, gets at most one packet. **Why**? [6 pts]

Because the switch is doing MAC learning: after the first frame it receives from H1 it stops broadcasting frames destined to H1, only sending them on the right port.

b. **Describe what H2 can do** (an attack) to force its way into getting most packets from H1 to S, again without preventing S from getting its packets. [6 pts]

H2 can flood the switch with frames with forged source MAC addresses, forcing S to evict H1's entry from its cache. This way, the next frames to H1 will be broadcast, and H2 will get them.

c. **Cite two differences** between this switched network and an old-fashioned Ethernet in which the switch S is replaced by a wire. [6 pts]

Original Ethernet has collisions, switched Ethernet does not; Switched Ethernet has MAC learning, so unicast frames are not broadcast.

d. If the network mask of H1 is accidentally changed to 255.255.255.128, **does the way in which it sends packets to H2's IP address change? If yes, how does it change? If not, why not?** [7 pts]

Yes. From H1's perspective, H2's address will no longer be in the local subnet. Thus, packets will have to go to the router instead, which will then forward them to H2.

3. Intra-domain Routing [25 pts]

a. In a network with Distance Vector routing, if the diameter (the length of the longest shortest path among any pair of nodes) of the network increases to large values, you have to adjust the protocol so that it keeps working. **What is this adjustment, and what is a bad side effect of this adjustment**? [5 pts]

Change the value of 'infinity'. The effect is that it will now take longer for the protocol to count to infinity, converging when a link disappears.

b. In DV routing, **can count-to-infinity happen when a link is added to the network? Why or why not?** (Assume nothing else changes in the network during this time.)[5 pts]

No, because when a link is added the length of the shortest path can only decrease, never increase.

c. **Mention one advantage** of Link State routing over Distance Vector. [5 pts]

No loops; no count-to-infinity; faster convergence; more robust to faulty or malicious nodes.

d. **Mention one advantage** of Distance Vector routing over Link State. [5 pts]

Smaller update traffic; simpler to implement.

e. One way to scale Link State is to use a hierarchical set of zones. If this solves scaling to large networks, **why is this not a good solution for the global Internet?** [5 pts]

On the Internet, shortest paths are not the main criterion, policy is important. You do not want to force all domains to implement the same routing algorithm.

4. IP Routing Tables [25 pts]

Having gotten tired of satellite links from question 1, you are now working with a Tier-1 ISP, and are constantly monitoring the size of the core BGP tables.

a. As a Tier-1 ISP, **why can't you have a default route in your routing tables**? [5 pts]

Because you are at the top of the hierarchy, there is no neighbor that knows more than you. You have to have the complete table, and know how to reach all prefixes.

b. For each of the items below, **say if it <u>increases</u>, <u>decreases</u>, or <u>doesn't affect</u> the number of entries in your routing tables, and briefly explain why:** (Assume that any announcements generated will eventually be reflected in your routing table.) [5 pts each]

(a) An organization with a spare /8 address allocation decides to sell blocks of this space to anyone that wants IP addresses.

Increases, because these blocks were previously aggregated in the same organization announcement, and now have to be specifically announced by different ASes.

(b) An organization that was a leaf in the BGP graph becomes multi-homed.

Increases, because the prefix of the organization now has to be specifically advertised in the second AS to which it connected. It was likely aggregated in the original upstream provider's announcements.

(c) An organization defending against a BGP prefix hijacking announces the same set of addresses as it had before, but using more specific prefixes.

Increases, because the previously aggregated announcement gets split into multiple announcements.

(d) An organization which had already had IP prefix 15.81.128/17 manages to buy 15.81.0/17.

Decreases, because it can now aggregate these to prefixes into one – 15.81/16.