Carnegie Mellon
Computer Science Department.
15-744 Spring 2015 Final Exam

Name: ______________________________

Andrew ID: ______________________________

INSTRUCTIONS:

- There are 9 pages (numbered at the bottom). Make sure you have all of them.
- Please write your name on this cover and put your initials at the top of each page in this booklet.
- If you find a question ambiguous, be sure to write down any assumptions you make.
- It is better to partially answer a question than to not attempt it at all.
- Be clear and concise. Limit your answers to the space provided.

Section	**Earned**	**Possible**
A		30
B		12
C		12
D		15
E		12
F		18
Total		99 + 1 (free bonus)

A Short Answers

1. Which of the following is true about the Content-Centric Networking (CCN) architecture?

   T F It is based on the observation that users are more interested in retrieving content than contacting hosts.

   T F Unlike bittorrent, CCN does not allow any host or router to provide the response to a content request

   T F CCN uses hashes to convert content names to fixed length identifiers to use in the network.

   Solution: T,F,F

   +1 for each correct answer

2. What are the four widely used metrics used to quantify the quality of Internet video? (Naming three of them will suffice.)

   Solution: Average bitrate, join time, buffering ratio, rate of buffering.

   +1 for each correct item. max of 3.

3. Give one reason that makes reflector-based denial-of-service attacks very powerful?

   Solution: The traffic reflected to the victim is significantly larger in bandwidth than the traffic an attacker has to send to the amplifiers. (2pts)

4. Give one reason that makes TCP perform poorly on wireless networks?

   Solution: packet losses are not congestion related (2pts)

5. Which of the following are true about the wireless systems we covered in class:
   - T F Since multi-path reflections may reduce signal strength, MIMO-based radios perform best in environments with no reflections.
   - T F Roofnet found that most links on their testbed surprisingly fell in one of two categories: 1) most packet received or 2) most packets lost.
   - T F WhiteSpace networks must avoid active transmission by TV stations and wireless microphones.
   - T F With appropriate topology and workload assumptions, the coding + MAC gains from a system like COPE can be unbounded.

Solution:

F, F, T, T

+1 for each correct answer

6. Describe one difference between the abstract model of the Internet described "Internet Inter-Domain Traffic" paper and the traditional model of the Internet?

Solution: A flatter and much more densely interconnected AS-level topology is emerging.

content providers are playing a more significant role in the network (2pts)

7. What are the three metrics that are considered in the FORTE design for request routing and traffic engineering to data centers?

Solution: Carbon emissions, average distance (or request latency), and electricity costs.

+1 for each correct item. max of 3.

8. Based on the key findings of the paper on IXPs, circle True/False:
   - T F Tier-1 ISPs do not have public peering at IXPs.
   - T F Many types of businesses (e.g., small ISPs, enterprise networks, CDNs, Tier-1 ISPs) peer at IXPs.
   - T F IXP peering links are primarily for backup and observe little traffic.

Solution: False, True, False

+1 for each correct answer

9. A naive implementation of IP traceback would require a packet to carry the ID of each on-path router in the packet. Describe one technique that they use to reduce the number of bits needed in each packet.

Solution:

1- edge sampling, 2- encoding partial information on IDs in each packet and collect many sampled packets to recover router IDs at the destination.

+2 for each correct answer, +2 for the description of each correct answer. max of .4

10. Consider the following AS topology:

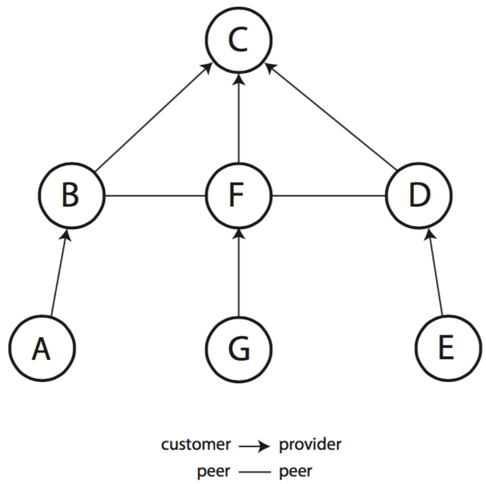

Which of the following BGP routes will AS F advertise to AS B?

(a) F

(b) G — F

(c) E — D — F

(d) E — D — C — F

> **Solution:** 1 and 2 (Peers do not advertise routes from other peers or providers; they only advertise routes to their own prefixes and to their customers.)
>
> +2 for each correct circle. -2 for each correct circle. min of 0.

B Flat Address Internet

11. Brian is thinking about a future Internet architecture. He has an idea to remove the IP layer and run TCP/UDP on top of Ethernet directly. Therefore, packets are addressed and forwarded by their MAC address. Routers use routing tables of MAC addresses to decide which packets should be forwarded to which interface.

    (a) What is the impact of using a flat address on router design?

    > **Solution:** Generally routers need to store much larger routing tables as the MAC addresses are less likely to be aggregated by their prefixes as IP addresses. (6pts)

    (b) With this flat address, for a malicious host is it easier or harder to spoof an other host in other subnets? Please briefly explain your answer.

    > **Solution:** It becomes easier. In an IP network, a malicious host claiming an IP address outside its subnet will receive no packet. (6pts)

C Hashing and Caching

12. Bovik is trying to figure out a scheme his clients should use, so that given a URL, they can find the appropriate CDN node to fetch the content from.

    Bovik has come up with a hash function h that takes a string and maps it to a real number in the range $[0, 1)$. Assume there are 3 CDN nodes with names such that $h(node_1) = 0.1, h(node_2) = 0.85, h(node_3) = 0.5$. When a client needs to fetch a URL and has to decide which replica to query, it picks $node_i$, such that the absolute value of the difference between $h(node_i)$ and $h(URL)$ is minimum. This scheme does *not* use circular mapping - it's just numeric closeness. This technique is "scheme 1"

    (a) Assuming all URLs are equi-popular, which node is likely to see the highest load?

    > **Solution:** 0-.3 go to node1, .3-.675 go to node3 and .675 - 1 go to node2. This makes $node_3$ the most loaded.
    >
    > If you mapped it in a circular fashion, $node_1$ gets .975 - .3, $node_3$ still gets .3-.675 and $node_2$ gets .675 to .975. $node_3$ is still the most loaded. (6pts)

    (b) Since load can be unevenly distributed in the above scheme, Bovik is not satisfied with the scheme. Instead, he thinks of a new arrangement. Let there be m CDN nodes in all; sort them using the $h(node_i)$ values. If the rank of a node is r, $(0 \leq r \leq m - 1)$, it is responsible for storing all URLs that map to the interval $[r/m, (r + 1)/m)$. This new scheme is called "scheme 2". Why might a CDN with a large number of nodes (that occasionally crash and are later repaired) choose scheme 1 over scheme 2?

    > **Solution:** Scheme 2 forces you to frequently move content between the nodes. Scheme 1 only forces the two neighboring ranges to adjust. This is one of the important properties of consistent hashes. (An argument based on node failures is acceptable too.) (6pts)

D Distributing a File to Internet

13. Dan wants to share a file of size S MB to 100 friends. So in this problem, he is looking for different methods to send the file across Internet. For simplicity, You can assume the downloading capacity is never the bottleneck.

   (a) If Dan sets up a FTP server at his home and its max upload rate is 10MB/s. The server splits its bandwidth across all concurrent users evenly. How long does it take to finish downloading for all his friends when $S = 1$ and $S = 500$?

   Solution:

   $$1MB/(10MB/s/100) = 10s$$
   $$500MB/(10MB/s/100) = 5000s$$

   (5pts)

   (b) Dan thinks a CDN might help, so he decides to try Akamai for his file transfers. Akamai distributes Dan's file on 5 different Akamai servers so that each server happens to serve 20 friends. Assume the RTT to the Akamai servers is 100ms and using Akamai involves 3 RTTs to resolve the closest Akamai server IP. Assume that each Akamai server can upload at 20 MB/s and that there are no other users except Dan's friends. Again, how long does it take to finish for $S = 1$ and $S = 500$?

   Solution:

   $$20MB/s/20 = 1MB/s$$
   $$1MB/1MB/s + 3RTT = 1.3s$$
   $$500MB/1MB/s + 3RTT = 500.3s$$

   (5pts)

   (c) After receiving the bill from Akamai, Dan decides to switch to BitTorrent. Basically Dan and his friends are joining the same torrent where Dan is acting as a seed and the others are peers downloading from the seed. All BitTorrent clients (including the seed) set their max uploading rate to 2MB/s. Assume the chunk size is 256KB and it takes about 10 seconds for each peer to connect to tracker, download the peer list and connect to all peers. To bootstrap, the seed sends an initial chunk to each client in parallel. After receiving the first chunk, all peers can fully utilize and evenly share their uploading capacity. How long does it take to finish for $S = 1$ and $S = 500$?

   Solution:

   $$0.25MB/(2MB/s/100) + 0.75MB/(2MB/s * 101/100) + 10 = 22.9s$$
   $$0.25MB/(2MB/s/100) + 499.75MB/(2MB/s * 101/100) + 10 = 269.9s$$

   (5pts)

E Wireless Networks

14. Wireless radios that transmit with higher power will have a larger range. Consider the wireless topology below:

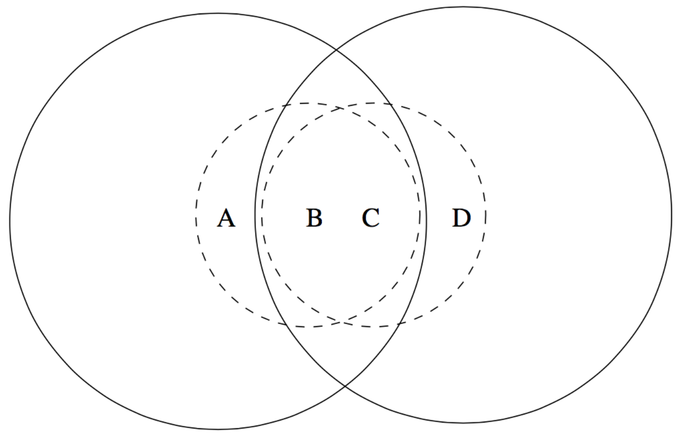

The solid circles represent the transmission radius of nodes A and D, respectively, and the dashed circles represent the transmission range of B and C, respectively. Assume that if two nodes' transmissions will interfere at a location if and only if they transmit at the same time and their transmission areas overlap. In these problems, assume that losses only occur due to collisions.

(a) When node A transmits to node B, list the potential hidden terminals (in either direction - those who might clobber A's transmission or those who A's transmission might clobber) and exposed terminals.

> **Solution:** C and D are both hidden terminals. If they were transmitting, A would not see them and would clobber reception by either B or C. Similarly, if A was transmitting, D would not see the transmission, and might clobber reception at B.
>
> C is an exposed terminal. When A is transmitting to B, C could (in the absence of ACKs) transmit to D. (6pts)

(b) What about when node B transmits to node C?

> **Solution:** D is a hidden terminal. It might broadcast, not hearing B's transmission, and clobber reception at C. There are no exposed terminals. (6pts)

F Carrier Sensing

The figure below shows the receiver preference region of CSMA. In the figure, S is the sender and I is the interferer, assume the clients are within the circle around S (but S does not know the exact locations of the clients). In the dark shaded areas, clients prefers concurrency, and in the light shaded areas, the clients prefers multiplexing. In case 1 and case 2, both S and I could hear each other, but in case 3, neither S or I could hear each other.

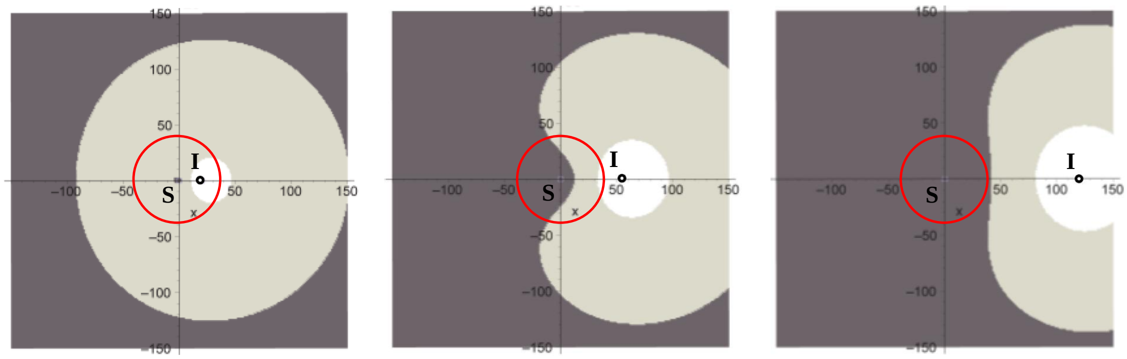

Case 1 Case 2 Case 3

For this problem, assume frames will be received in the dark shaded areas and will be lost in the light shaded areas. Assume both S and I are persistently backlogged (always have packets to send). Suppose the physical data rate on the is R.

15. What is the expected receive rate for S's clients if both S and I defer to each others transmissions
    (a) Case 1:
    (b) Case 2:
    (c) Case 3:

> **Solution:** (a) R/2, (b) R/2, (c) R/2 (or R)
> +2 for each correct item.

16. What is the expected receive rate for S's clients if both S and I transmit concurrently
    (a) Case 1:
    (b) Case 2:
    (c) Case 3:

> **Solution:** (a) 0, (b) R/2, (c) R
> +2 for each correct item.

17. In terms of both fairness and performance, should S and I differ to each other or transmit concurrently in each of the below cases:
    (a) Case 1:
    (b) Case 2:
    (c) Case 3:

Solution: (a) defer, (b) defer, (c) concurrent (or does not matter)
+2 for each correct item.