CS144 Practice Problems For Final
Fall 2010

Note: These problems cover a **subset** of the material that we expect you to be familiar with for the final. In particular, these problems primarily cover material taught in the second half of the course. Although the final will be weighted more heavily towards no material, old material will also be well-represented on the final.

Question 1

In the first lecture on security, it was said that you cannot trust an HTTP URL when on an open wireless network. List one man-in-the-middle attack that an adversary could launch against each of DHCP, DNS, and HTTP (a total of 3 attacks).

Question 2

Eric is designing a new file transfer protocol on top of UDP. Eric's initial design is simple: immediately send all blocks of data to the remote side, and wait for the remote side to send a file checksum. If the checksum is incorrect, the sender retransmits the whole file.

Part (a)

Suppose the network guarantees that no packets are lost, and the protocol handles packet reordering. Remember that UDP has a checksum. With these guarantees, is the checksum in Eric's protocol necessary? Why or why not?

Part (b)

Suppose Eric wants to his protocol to environments where receiver-to-sender communication is expensive or impossible (such as IP Multicast or wide-area wireless). How can he adapt his protocol so that receivers get the complete file efficiently and with high probability?

Part (c)

While receiving a large file over the Internet using his protocol, Eric notices that his existing HTTP download of the latest Debian Linux ISO from the same source slows down considerably. Eric expected for both file transfers to use about the same amount of bandwidth, but he finds that the transfer using his protocol uses more bandwidth than the HTTP download. What might be happening? How could Eric change his protocol to act more like what he expected?

Question 3

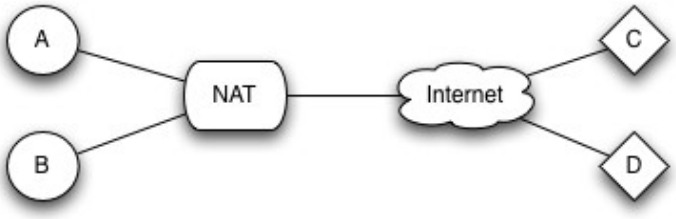

You are tasked with analyzing the NAT device in the network shown above. Hosts A and B are internal to the NAT, and hosts C and D are connected directly to the Internet and not behind NATs. For this problem, host and port pairs are represented as host:port, e.g., C:80 might be a webserver running on C.

You observe the following (sequential) UDP packets:

1. A:45920 sends a packet to C:5060. Allowed. The packet appears (to C) to be coming from NAT:25223.
2. C:5060 sends a packet to NAT:25223. Allowed. NAT forwards packet to A:45920.
3. A:45920 sends a packet to D:5061. Allowed. The connection appears (to D) to be coming from NAT:25223.
4. B:54429 sends a packet to C:5060. Allowed. The packet appears (to C) to be coming from NAT:27558.
5. D:5061 sends a packet to NAT:27558. The packet is denied at the NAT and is not forwarded.
6. B:54429 sends a packet to D:5061. Allowed. The packet appears (to D) to be coming from NAT:27558.
7. D:7058 sends a packet to NAT:27558. The packet is denied at the NAT and is not forwarded.
8. D:5061 sends a packet to NAT:27558. Allowed. NAT forwards packet to B:54429.

Part (a)

Which of the following policies does this NAT use?

a) Full Cone
b) Restricted Cone
c) Port Restricted Cone
d) Symmetric

Justify your answer.

Part (b)

Suppose packet 7 above was allowed at the NAT and was forwarded to B:54429. Would this change your answer to part (a)? Explain why or why not.

Question 4

You and your friend are on opposite sides of a large lecture hall (200' apart) and are both using the wireless network. The access point is between the two of you, but very close to your friend (20' away) and far from you (180'). Your network is very slow, giving you 50kbps, while your friend is able to download at >1Mbps. Assume there are other people actively using the wireless network in the lecture hall. Why might this be the case?

Question 5

Suppose you are running a webserver and you get a hit from an IP address that has a PTR record of "client-42.ext.google.com". The user agent string contains "Chrome", so you can assume it's not a Google web crawling bot. Was the hit from a Google-owned computer? What else might you do to prove it, one way or another?

Question 6

Consider the following message authentication coding algorithm, MOD_MAC.

MOD_MAC takes in a message, M, and a key, K. K is 80 bits and known by both sender and receiver beforehand.

MOD_MAC divides a message into 80 bit chunks (padding 0s onto the end of the message if the message length is not evenly divisible).

The MAC is not a prespecified length. Instead, each bit of the MAC MOD_MAC produces maps to an 80-bit chunk of the message (so, a 8000-bit long message has a 100-bit mac). More specifically, each bit of the resulting mac is given by:

MAC[bit i] = (k xor message_chunk_i) mod 2;

Does this MAC provide reasonable integrity? Why or why not?

Question 7

Assume you have the following topology:

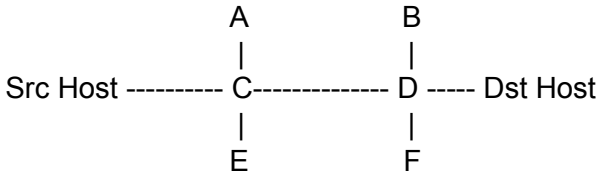

For parts a-c, assume no fragmentation, and that only data packets are dropped (ie acknowledgment packets are never dropped)

Part (a)

Assume that each link does not drop packets. What is the ETX for the path from src to dst?

Part (b)

Assume that each link drops a packet with a probability of 50%. What is the ETX for the path from src to dst? (Note: tricky question!!!)

Part (c)

Assume that each link drops a packet independently with the following probabilities: src-to-c drops with probability of 1/3; c-to-d drops with probability 1/2; and d-dst drops with probability 1/5. What is the ETX for the path from src to dst?

Question 8

(partially inspired by http://www.tomkleinpeter.com/2008/03/17/programmers-toolbox-part-3-consistent-hashing/):

Your friend, Alex, is starting a company. Alex's company has 50 machines that serve customers whose ip addresses uniformly range from 18.0.0.0 to 24.255.255.255. Each machine contains customer records, which are frequently backed up to external storage. However, none of Alex's customer data is replicated between his/her 50 machines.

If Alex's company does well, he/she knows that the number of requests to his/her servers will go up, and he'll/she'll need to add additional machines one-by-one. Similarly, Alex knows that each of his/her machines may fail (infrequently). Transferring customer records between machines or from external storage to a machine is expensive.

Alex recommends the following algorithm for mapping customers to internal machines:
1) Assign each machine a unique number from 0 to 49.
2) Assign each ip address in the given range a unique number from 0 to (7*(2^24) -1)
3) To decide how to route each customer's traffic (and where to store his/her record), take the number associated with the customer's ip address and modding it by the number of available machines (initially 50).

When a machine fails or is added, re-run the algorithm with the new available number of machines, moving records between machines/external storage appropriately.

Is this algorithm a good idea? What could Alex do instead?

Question 9

Why is variable bit rate support important in wireless? What tradeoff is being made by changing the bit rate?

Question 10

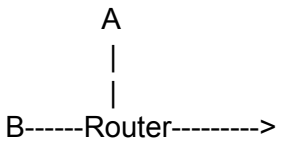

At the end of every second, the router outputs a single packet from the front of an internal queue that is 100 packets long. If there is no packet in the queue, the router sends no packet. This internal queue uses a drop tail policy.

At the end of every half second, A sends a single packet to the router. If there is room in the router's internal queue, the packet is appended to the queue. If there is not, the packet is dropped.

At the end of every minute, B sends an instantaneous 120 packet burst to the router. If there is room in the router's internal queue, the packet is appended to the queue. If there is not, the packet is dropped.

Assume that this system has reached equilibrium (ie the Router and A and B have been running for a while).

Part (a)

Will packets ever be dropped? If so, about how many packets will be dropped per minute? If not what is the output rate of the router?

Part (b)

For every 60 packets that the router sends out, how many will be from A (+/- 2 packets) and how many from B (+/- 2 packets)? (The +/- 2 packets is included to simplify any issues from packets from A and B arriving simultaneously etc.)