Computer Networks Qualifying Exam
Fall 2011

Department of Computer Science
New Mexico State University

Note: Please attempt all questions. Justify all your answers. This examination is open book and notes. Please write your answers in a separate paper and number them correctly.

[1.] All the sub-questions relate to the application layer of the TCP/IP stack. [**20 pts:** 5 pts + 5 pts + 5 pts + 5 pts]

a. How are the Internet domain name service (DNS) servers arranged? Please draw a diagram to show DNS servers in a three level network, where the lowest level has five nodes, the next level has three nodes, and the top level has one node.

b. If you have access to the canonical name of a website and you want to know it's IP address, and you have access to a linux machine connected to the network having a DNS server, can you identify the IP address? Please explain how.

c. If you know the canonical name of the website you want to connect to and ONLY your local DNS server is down, can you connect to the website through your browser? Please explain your answer.

d. Is there a way to perform a denial of service attack on a DNS server? Please support your answer with explanation.

[2.] All the sub-questions relate to the transport layer of the TCP/IP stack.

a. The amount of unacknowledged data in TCP is limited by three quantities: the congestion window size, the receive window size (flow control), and the RTT-bandwidth product. Is it possible for TCP Reno (the commonly used version) [**10 pts:** 5 pts + 5 pts]

   i. to reach a state with congestion window size much larger than RTT $\times$ bandwidth? Why? How likely is it?

   ii. to reach a state with congestion window size much larger than the receive window? Why? How likely is it?

b. Given that congestion control results in a node's traffic getting throttled, is congestion control a desirable property for an individual node transmitting large amounts of data packets? Please explain your reasoning. [**10 pts**]

c. The UDP protocol does not support congestion control, however, it has other properties that make it a desirable protocol. Assume you have access to all the routers in a network and can change their code such tha they can access the data packet and add/update fields in the packets. Using this additional capability, design a congestion identification protocol with low message and computation complexity, which runs on top of UDP and helps the sender identify congestion. Please explain the protocol clearly and explain why it has low complexity. [**10 pts**]

[3.] All the sub-questions relate to the network layer of the TCP/IP stack.

a. Consider the situation where Host A is on LAN 1 and Host B is on LAN 2, and the Router R directly connects the two LANs. The MTUs of the two LANs are 1500 bytes and 1024 bytes respectively. Suppose an application running on Host A executes 1000 writes to an application running on Host B, each write resulting in a full one-MTU sized IPv4 datagram on LAN 1, with the minimal IP header (no options). How many datagrams from Host A to Host B traverse LAN 2? What are their sizes and offset field values? (Assume that there is no path MTU discovery.) [**10 pts**]

b. Consider routing IPv6 datagrams using an IP tunnel over a path involving IPv4- only routers. Suppose that there are 18 actual hops from the source S to the destination D, but 12 of the hops are through a IPv6-in-IPv4 tunnel. If S sets the TTL of a datagram to 31, what would be the value of TTL when it reaches D? [**5 pts**]

c. Recall that IPv6 datagram can be fragmented only at the time of its formation at the source S, is it possible that the IPv4 datagrams implementing the tunnel get fragmented? Justify your answer briefly with appropriate assumptions. [**5 pts**]

d. Consider the two basic approaches for broadcast, namely unicast emulation and network layer (i.e. router assisted) broadcast; suppose the spanning tree is used to achieve network-layer broadcast. Consider a single sender and 64 receivers. [**15 pts** = 10 pts + 5 pts]

- Suppose the sender is connected to the receivers by a balanced binary tree of routers (with the receivers being the leaves) and assume that if each time a packet (or its copy) is sent over a link, the cost incurred is *one*, then what are the costs incurred in the unicast emulation and network layer techniques.
- If you had a choice of using as many routers as possible, can you find a topology for which the difference between the costs of the unicast emulation and the network-layer broadcast is maximized. If so, name the topology and explain why?

[4.] All the sub-questions relate to the data link layer of the TCP/IP stack.

a. Suppose you have installed an *aware* networking protocol stack as part of a new experimental OS kernel on your Wi-Fi and cellular communication capable PDA. On the PDA, you typically use one primary network application at a time, and traffic due to application mixes is not a very important consideration. You have a choice of three user selectable MAC protocols to be associated with each network application program, namely FDMA, TDMA, CSMA/CA with RTS/CTS disabled, and full CSMA/CA (with RTS/CTS). Assume that the access points support on-the-fly negotiation of MAC choices with each individual mobile station. For each of the following three applications, select the MAC protocol you would like to associate explaining your reasons: [**15 pts**]

i. VoIP phone calls without WiFi access (FDMA. may be even multi-channel)

ii. MPEG mobile TV streaming (TDMA)

iii. Extremely delay sensitive Instant Messenger chat (FDMA)