CS 422 (Park) **Final** **Dec. 13 (Thu.), 2007**

Remarks: Keep the answers compact, yet precise and to-the-point. Long-winded answers that do not address the key points are of limited value. Binary answers that give little indication of understanding are no good either. Time is not meant to be plentiful. Make sure not to get bogged down on a single problem.

PROBLEM 1 (40 pts)

(a) What is the fundamental difference between inter-domain and intra-domain routing? In the latter, one popular method implemented by several protocols is finding a shortest-path. Why are there different protocols—all used in real-world networks—for implementing shortest-path? Which protocol is preferred by large ISPs and why?

(b) Why is flat (i.e., CBR) traffic preferred over fluctuating (i.e., VBR) traffic by traffic engineers? What is a key reason that causes Internet traffic to be VBR? How do data and telephone traffic differ from each other with respect to their time dependent characteristics? What impact does this have on traffic engineering?

(c) TCP's main congestion control component is congestion avoidance which implements linear increase/exponential decrease. Why does increasing the congestion window by 1 when an ACK packet arrives not lead to linear increase? How should the congestion window be updated to effect linear increase? When TCP is ported to router operating systems, what additional implementation issue should one be concerned with?

(d) What is the role of IP-over-IP tunneling in Mobile IP? How does it work? Why is there a need for a reverse tunnel from a mobile to its home agent?

PROBLEM 2 (40 pts)

(a) What are the key mechanisms that have prevented IPv4 address depletion over the years? Provide a ranking of the mechanisms with respect to their impact or importance. Explain your ranking. What are some undesirable by-products of the mechanisms that residential users using cable/DSL access lines have been saddled with?

(b) A DoS attack may be directed at a server by sending a rapid succession of bogus packets, TCP or UDP, that eat up bandwidth and consume the server's CPU cycles. For a web server listening on port 80, which form of attack is more damaging and why? If an attacker were sending UDP packets targeted at port 80, what defense might work at the gateway router of the organization wherein the server is located so that it is shielded from outside attacks? Would the solution extend to the case when an attacker sends TCP packets directed at port 80?

(c) What is the 2-person consensus problem? Why does the problem not have a solution? How does the problem manifest itself during TCP connection set-up and tear-down? How is the problem dealt with in practice?

(d) What is the hidden station problem in wireless LANs? How is the hidden station problem related to the starvation problem? How is RTS/CTS intended to solve the hidden station problem? Is it an effective solution? Explain your reasoning. Can RTS/CTS help address the starvation problem? What feature of IEEE 802.11 WLANs can definitely help alleviate the starvation problem?

PROBLEM 3 (20 pts)

What is the technical meaning of congestion? Give an illustration in the specific context of WLANs. How is congestion control effected in IEEE 802.11 WLAN? How is it effected in Ethernet? How is it effected in TCP? What is the goal of congestion control in peudo real-time multimedia streaming? Is the goal the same as in WLAN, Ethernet, or TCP? Explain your reasoning. HTTP does not have its own congestion control mechanism, instead relying on TCP to pace its traffic. However, HTTP does employ caching to reduce overall traffic on the Internet as well as improve users' web response time. What are the main approaches to achieving HTTP or web caching? Which one might be best? What role can DNS play to reduce congestion and improve user experience?

BONUS PROBLEM (10 pts)

On the Internet, it is not sufficient to have a sound technical solution to achieve actual deployment and use of services enabled by the solution. In many (perhaps most) instances backward compatibility is a necessary requirement. Solutions that are not backward compatible with legacy systems tend to fall by the wayside. Describe examples of network technologies and solutions that have succeeded by following the legacy compatibility constraint "to the extreme." What are examples that have not succeeded?