CS 422 (Park) **Final** **Dec. 13 (Mon.), 1999**

PROBLEM 1 (40 pts)

(a) What are the two (complementary) definitions of congestion? How are they related to each other? If a router is fed with variable bit rate (VBR) input and raw throughput—e.g., number of IP packets per second—is measured at an output link, is congestion possible in this set-up? Explain.

(b) The congestion window update in the linear increase phase of TCP's congestion control follows the procedure `CongestionWindow` ← `CongestionWindow` + (1 / `CongestionWindow`). Why is this particular form of update necessitated, and how does it lead to a linear increase? What happens if the increase procedure is replaced by `CongestionWindow` ← `CongestionWindow` + 1?

(c) Why is constant bit rate (CBR) traffic, in general, easier to manage than VBR traffic? Illustrate your answer with an example consisting of two CBR sources with average bit rate 5 Mbps each sharing an output link with bandwidth 10 Mbps, and the same configuration with the CBR sources replaced by VBR sources.

(d) What are the two common modes in which ATM is used in internetworks? Show the structure of the protocol stacks in the two respective cases. What are the pros/cons of the two approaches? What are the roles of the two layers in the ATM protocol stack?

PROBLEM 2 (40 pts)

(a) Shortest path routing is known to suffer from *fluttering*. Define the fluttering problem. Give an example topology with two flows to illustrate its manifestation. What is its root cause? Suggest possible ways to fix it.

(b) What is a leaky bucket? Give a specification of its parameters and how they govern the "traffic shaping" affected. How is leaky bucket used in traffic control? How is it related to characterizing a VBR flow by an average rate and peak rate?

(c) What are the pros/cons of provisioning high-level network services over TCP vs. UDP? In the case of UDP, what are the reasons for refraining from provisioning services directly over IP bypassing UDP altogether? Can you think of instances where provisioning a service directly over the data link layer (i.e., omitting IP) is justified?

(d) Assume you are given a traffic trace measured at an IP router at 10 ms granularity. How would you practically go about checking—using visual inspection—whether the traffic is self-similar or Poisson? Give an intuitive explanation of what "self-similar" means for traffic traces, and how it would show up in the visual inspection.

PROBLEM 3 (20 pts)

(a) A router implementing priority packet scheduling (vis-à-vis FIFO) always services higher priority packets first, i.e., lower priority packets are serviced only when there are no higher priority packets needing attention. What is the drawback of designing routers capable of providing preferential services in this fashion? What is weighted round robin (or weighted fair queueing), and in what sense is it superior to priority scheduling? How would you implement either scheme in the context of IPv4?

(b) How is thermostat control related to congestion control? Why is congestion control a more difficult problem? What can happen if no further amends are made? What fix is required and how is it reflected in a refined control law (i.e., rate update procedure)? What further subtlety exists, and what additional change does it require of the update procedure? How is this related to TCP?

BONUS PROBLEM 4 (20 pts)

The Internet is dominated by cooperative protocols—a primary example being TCP—which exhibit "gentlemanly" backoff behavior when faced with congestion and other adverse network conditions. Describe a noncooperative version of TCP, call it TCP-Greedy, which would exploit the cooperative behavior of other flows to its advantage. Further enhance your protocol so that it does not "hurt" itself (i.e., cause self-congestion) when it is the only flow in the system. That is, there are no other flows to exploit in the system.