CS 690N (Park) **Midterm** **March 28 (Tue.), 2000**

Remarks: Keep the answers compact, yet precise and to-the-point. Long-winded answers that do not address the key points are of limited utility. Time is not meant to be overly plentiful. Make sure not to get bogged down on a single problem.

PROBLEM 1 (25 pts = 10 + 15)

(a) Consider a traffic source $X(t)$. $X(t)$ may be dependent in time, or, as a special case, independent. What are the implications of independent traffic for traffic control? Does it make traffic control "easy," "difficult," or some combination thereof?

Consider the case when $X(t)$ is dependent. As $X(t)$ becomes more strongly dependent or correlated in time, what impact does it have on traffic control? Is there more than one side to the story? Is there a trade-off?

(b) It is often said that self-similar burstiness renders buffering as a resource provisioning strategy at routers ineffective. Give a explanation of why this is so using the set-up where self-similar traffic $X(t)$—in the form of on/off traffic with heavy-tailed on-periods (but exponential off-periods)—enters an output-buffered router with buffer capacity B (bits) and bandwidth μ (bps) on a designated output port.

Conversely, if input traffic is bursty but not long-range dependent (e.g., exponential on-periods), explain why buffering is effective.

PROBLEM 2 (45 pts = 15 + 25 + 5)

(a) Why is feedback congestion control, in general, a difficult control problem? Why is it, for example, more difficult than controlling room temperature using a thermostat? What can go wrong if you apply the thermostat control to congestion control? When can congestion control degenerate to the more simple thermostat control (i.e., the same control law can be effectively used)? Enumerate the principal reasons why instability can potentially arise. What are the countermeasures that can be employed to contain the corresponding instabilities?

(b) Consider the congestion control law discussed in class in its asymmetric form where control law is given by

$$\frac{d\lambda(t)}{dt} = \begin{cases} \epsilon(Q^* - Q(t)), & \text{if } Q(t) \le Q^*, \\ -a\lambda(t), & \text{otherwise.} \end{cases}$$

As usual, the change in network load $Q(t)$ is determined by the update equation $dQ(t)/dt = \lambda(t) - \mu(Q)$. Using the tools developed in class, give a complete analysis of the stability/instability of the optimal operating point (Q^*, λ^*) where $\lambda^* = \mu(Q^*)$. What is the condition for stability? Is it asymptotically stable?

(c) Assume that the asymmetric congestion control of part (b) is used to control the client-side buffer of pseudo-real-time MPEG video/audio players (e.g., *RealPlayer*, *Windows Media Player*) such that a target buffer content (containing MPEG video/audio frames) is maintained at the client for uninterrupted playback. Is the control law implementable, as is (except for discretization), without incurring unreasonable assumptions? What if the control is used to transport bulk data end-to-end on the IP Internet where throughput maximization is the objective?

PROBLEM 3 (45 pts = 10 + 15 + 20) [part (b) is a bonus question]

(a) Assume n traffic flows $X_i(t)$, $i = 1, 2, \ldots, n$, with peak rate λ_i^* and average rate λ_i are admitted to a network system using open-loop control. Assume the connections share a bottleneck router with bandwidth μ (the flows are routed to the same output port) where the router implements GPS such that a flow can be reserved its individual minimal guaranteed bandwidth μ_i with $\sum_{i=1}^{n} \mu_i \le \mu$. Assuming a bufferless system, if peak rate allocation is used resulting in zero packet loss, what is the admission control test that a subset $S \subseteq \{1, 2, \ldots, n\}$ of flows must pass to have their QoS requirements satisfied with probability 1? As a service provider, if network utilization were your objective, how would you choose S such that your objective is optimized? Give the formula for utilization when S is the set of flows admitted. Given the n flows, how would you choose $S \subseteq \{1, 2, \ldots, n\}$ such that utilization is maximized?

Assume you are a service provider whose main objective is profit. Assume that you charge a price p for a unit flow (1 bps) admitted such that revenue is given by $\sum_{i \in S} p\lambda_i$. If this is your revised criterion, does your optimal admitted set S change from that of utilization maximization?

(b) Assume you are offering a more generalized service where each user (or flow) i is allowed to specify a target packet loss rate c_i^* $(0 \leq c_i^* \leq 1)$, and, if admitted, the actual packet loss rate c_i that you deliver to user i is guaranteed to be $c_i \leq c_i^*$. Thus zero packet loss rate is a special case of this more general—but still guaranteed service semantics—set-up. Give the revised admission control test for user set $S \subseteq \{1, \ldots, n\}$. What is the corresponding utilization when S has passed admission? What is the gain in utilization vis-à-vis the special case of zero packet loss in (b) by allowing the more generalized QoS service? How would you change your pricing scheme in the revised set-up, and why?

(c) Assume n flows are mutually independent. What will the distribution of the aggregate flow $X(t) = \sum_{i=1}^{n} X_i(t)$ resemble when n is "large" (note that finite peak rate implies finite variance, call it σ_i^2)? Give the expression for the distribution that $X(t)$ converges to as $n \to \infty$. Compared to part (a), by allowing an overall packet loss rate ε, what is the revised admission control test that can be administered to decide whether a flow set $S \subseteq \{1, \ldots, n\}$ can be admitted? Why is this statistical guarantee, strictly speaking, not a guaranteed in the sense of (a) and (b)?

Now, to the most important question of Problem 3: As far as packet loss rate is concerned, is admission control affected by whether the flows are correlated, in particular, self-similar? Here we assume that, whether correlated or not, the peak rate, average rate, and variance remain the same, and the system is bufferless. Carefully explain your reasoning.