CS 690N (Park) **Final** **May 2 (Tue.), 2000**

PROBLEM 1

(a) What are the key differences between packet-level FEC and bit-level FEC? Recall that the latter is commonly used in CD players (i.e., storage), video transmission by deep space probes journeying to the outer edges of our solar system, and so forth. What makes packet-level FEC a more difficult problem—from a traffic control perspective—than bit-level FEC?

(b) Assume a leaky bucket with parameters (σ, ρ) is used to shape an incoming traffic $X(t)$ with average rate λ, thus yielding a departure process whose *cumulative* process $A(t)$ satisfies $A(t) \leq \rho t + \sigma$. Let $Y(t)$ denote the cumulative process corresponding to $X(t)$, i.e.,

$$Y(t) = \int_0^t X(s)ds.$$

Assuming zero buffering (a packet is discarded by the leaky bucket regulator if no tokens are present upon arrival), what is the expression for "packet loss rate" (or discard rate) as affected by the leaky bucket? Assuming $X(t)$ is short-range dependent (e.g., Poisson), how would you go about setting (σ, ρ) such that the discard rate is bounded above by ε? Your procedure should be efficient in the resource allocation sense that ρ and σ are made minimal. Sketch, in some detail, the computations needed when $X(t)$ is an i.i.d. process and minimizing ρ is the primary objective.

(c) As a continuation of part (b), how is your answer to the leaky bucket parameter assignment problem affected if $X(t)$ is self-similar with the same average rate λ? For the case where $X(t)$, in addition to being self-similar, is Gaussian (i.e., for each t, $X(t)$ has a normal distribution with mean λ and variance ν^2), show how you would go about computing an efficient assignment for (σ, ρ) that satisfies the discard rate ε.

PROBLEM 2

(a) Consider periodic hard real-time applications such as real-time MPEG video with frame rate f (frames per second), e.g., 25 fps, where successive video frames must arrive (and be decoded) by intervals of period $T = 1/f$ (e.g., 40 ms) to be of positive utility to the receiver. Let γ $(0 \leq \gamma \leq 1)$ denote the hit rate, i.e., fraction of deadlines over a stretch of time that have been met. Alternatively, γ may be viewed as the probability of a frame meeting its deadline. Assume the sender transmits the video stream over a FIFO router configured with bandwidth μ and buffer capacity B. For the present exercise, ignore the link latencies and decoding times (very important). Let the delay be given as a function $d(\lambda, \mu, B)$ where d is increasing in λ (average bit rate, an increasing function of f) and B, but decreasing in μ. Let $c(\lambda, \mu, B)$ denote the packet loss rate where c is increasing in λ but decreasing in B and μ. Thus the hit rate γ is a function of c and d, hence a function of λ, B, and μ. What is the qualitative dependence of γ on λ, B, and μ? Be as concrete as you can be when uncovering the underlying relationships.

(b) As a continuation of part (a), assume the MPEG video sender engages in FEC such that, on average, h additional bits (packetized) are transmitted per frame. Given frame rate f, what is the new total bit rate λ' entering the network? What is the average total bit rate (including redundant bits) corresponding to a single frame? Identify the conditions under which FEC (i.e., $h > 0$) can help increase the hit rate γ. Can increasing h "too much" hurt γ?

(c) A router deemed capable of providing "differentiated services" per-hop scheduling support has been associated with the three properties denoted (A1), (A2), and (B). What are the roles of (A1), (A2), and (B)? Is (A2) a "desirable" property? In what sense is it different from (A1) and (B) from a design perspective? Does a priority queue satisfy the above properties? Give a sketch of the proofs. What about FIFO? Does satisfying the above properties suffice to make a "good" differentiated services router? Explain your reasoning.

PROBLEM 3

(a) Assume a router satisfies properties (A1), (A2), and (B). Assume you are given the task of writing an end-to-end protocol—adaptive label control—that dynamically adjusts the label value such that a target throughput λ^* (as experienced by the receiver) is achieved. The receiver sends back (i.e., feedback) to the sender its measured throughput over a small time window. Write down the control law for adjusting the label values (e.g., enscribed in the TOS field of IP) to achieve λ^*. What are the principal factors that determine whether adaptive label control will attain λ^* or not? When compared to feedback congestion control, is feedback (or adaptive) label control a simpler or more difficult control problem? For example, is there a need for exponential back-off? Explain your reasoning. How does adaptive redundancy control (i.e., the adaptive varying of h in Problem 2 (b)) compare in difficulty with the above controls?

(b) Assume you have been hired as Chief Technology Officer of a major ISP and you are faced with the daunting task of providing quality of service to a diverse customer base including those performing e-commerce transactions, IP telephony, on-line video rental, and so forth, circa 2003. Upon performing measurement observation, you realize that the Internet has become like the "Wild Wild West" (WWW) where end users engage in greedy-TCP (doesn't back off expecting to soak up bandwidth relinquished by other cooperative TCP flows), use packet-level FEC to amplify QoS, among other end-to-end traffic control techniques picked up, in part, at Purdue. Due to the abuse of these techniques by certain elements, however, you find that your network experiences concentrated periods of congestion and your high profit customer base (i.e., business users) is opting to lease expensive private lines instead of using your network. As a first step to imposing some order onto the system, you decide to implement *usage pricing* geared toward maintaining the overall traffic load $Q(t)$—which, you, as the ISP, can reasonably monitor—near some target traffic load Q^* at which services are reasonably fast, even for the typical business user. You install traffic meters at every access point (i.e., customer premise) and charge a cost of $p\lambda_i$ to user i for having sent traffic λ_i. Since users value "money," you observe that λ_i is a function of price, i.e., $\lambda_i = \lambda_i(p)$, where $\lambda_i(p)$ is a decreasing function of p. Assuming your sole objective is to promote an overall traffic load of Q^*, write down the control law for updating p (i.e., expressed as $dp/dt = F(\cdot)$) which achieves your stated goal. Compared to part (a), what kind of a control problem is adaptive usage pricing as formulated above with respect to its difficulty? If there are two types of users—business and casual—with demand functions $\lambda_B(p)$ and $\lambda_C(p)$, respectively, and your objective is revenue maximization independent of Q^*, how would you go about determining the optimal price p?