PS #1 – EE 599: Stochastic Network Optimization Due: Monday, Sept. 15, 2008

Reading:

- Sections 3.1 and 4.4 of text.
- "Section III: Convexity" of the "Limsups, Sets, and Convexity" handout.

I. Sample Paths and T-slot drift

Consider a discrete time queue that serves fixed length packets according to the queueing law:

$$U(t+1) = \max[U(t) - \mu(t), 0] + A(t) \quad \text{for } t \in \{0, 1, 2, \ldots\} \tag{1}$$

where $A(t)$ and $U(t)$ take integer units of packets, and $\mu(t)$ takes integer units of packets/slot. Assume $U(0) = 0$.

a) The $A(t)$ and $\mu(t)$ values for $t \in \{0, 1, \ldots, 10\}$ are shown. Complete the table by filling in the appropriate $U(t)$ and $\tilde{\mu}(t)$ values.[1]

	t	0	1	2	3	4	5	6	7	8	9	10
Arrivals	$A(t)$	3	3	0	2	1	0	0	2	0	0	0
Current Rate	$\mu(t)$	4	2	1	3	3	2	2	4	0	2	1
Backlog	$U(t)$	0										
Transmitted	$\tilde{\mu}(t)$											

Verify that $U(6) - U(0) = \sum_{\tau=0}^{5} A(\tau) - \sum_{\tau=0}^{5} \tilde{\mu}(\tau)$. Verify that $U(8) - U(2) = \sum_{\tau=2}^{7} A(\tau) - \sum_{\tau=2}^{7} \tilde{\mu}(\tau)$. Verify that $U(9) - U(0) = \sum_{\tau=0}^{8} A(\tau) - \sum_{\tau=0}^{8} \tilde{\mu}(\tau)$. Verify that $\tilde{\mu}(t) \leq \mu(t)$ for all t.

b) Suppose that the same arrival sequence is repeated periodically every 11 slots (so that $A(11) = 3, A(12) = 2, A(13) = 0$, etc.). Show that $A(t)$ is *admissible*, and compute its rate (note: because $A(t)$ in this case is deterministic, all expectations are equal to the actual values with probability 1, so that $\mathbb{E}\{A(11)\} = 3$).

c) Argue that for any discrete time queue that satisfies (1), we have for any time t and any positive integer T:

$$U(t+T) \leq \max\left[U(t) - \sum_{k=0}^{T-1} \mu(t+k), 0\right] + \sum_{k=0}^{T-1} A(t+k)$$

d) Using $T = 11$, compute $\left(\sum_{k=0}^{T-1} A(t+k)\right)^2$ (for any time t).

e) Assume $\mu(t)$ is i.i.d. over timeslots, with $Pr[\mu(t) = 0] = 1/11$, $Pr[\mu(t) = 1] = 2/11$, $Pr[\mu(t) = 2] = 4/11$, $Pr[\mu(t) = 3] = 2/11$, $Pr[\mu(t) = 4] = 2/11$. Compute $\overline{\mu}$. Compute $\mathbb{E}\left\{\left(\sum_{k=0}^{T-1} \mu(t+k)\right)^2\right\}$ for general T, t.

f) Define the *T-slot conditional Lyapunov drift* $\Delta_T(U(t))$ as follows:

$$\Delta_T(U(t)) \triangleq \mathbb{E}\left\{U^2(t+T) - U^2(t) \mid U(t)\right\}$$

Compute an upper bound on $\Delta_T(U(t))$ for $T = 11$ and assuming the $A(t)$ and $\mu(t)$ processes as defined in parts (b) and (e). The bound should have the form:

$$\Delta_T(U(t)) \leq B - \epsilon U(t) \tag{2}$$

for all timeslots t and for some fixed values $B > 0$ and $\epsilon > 0$. What are the values for B and ϵ?

g) Show that if (2) holds for all t, then for any integer $M > 0$ and any time $t_0 \in \{0, 1, \ldots, T-1\}$, we have:

$$\frac{\mathbb{E}\left\{U^2(MT)\right\} - \mathbb{E}\left\{U^2(t_0)\right\}}{M} \leq B - \epsilon \frac{1}{M} \sum_{m=0}^{M-1} \mathbb{E}\left\{U(t_0 + mT)\right\}$$

[1]In all parts of this document where you are asked to fill in a chart (this problem and in problem V), you can make your own (similar) chart to hand in for your solutions (you do not need to hand in this exact page).

h) Conclude that: If (2) holds for all t, then for any integer $M > 0$:

$$\frac{1}{MT}\sum_{\tau=0}^{MT-1}\mathbb{E}\left\{U(\tau)\right\} \leq \frac{B}{\epsilon} + \frac{1}{M\epsilon}\frac{1}{T}\sum_{t_0=0}^{T-1}\mathbb{E}\left\{U^2(t_0)\right\}$$

and hence (taking the limit as $M \to \infty$):[2]

$$\limsup_{t\to\infty}\frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\left\{U(\tau)\right\} \leq B/\epsilon$$

II. Convexity

a) Let A be a $M \times N$ matrix and let $\boldsymbol{b}$ be a M dimensional column vector. Consider the N-dimensional set $\mathcal{X}$ defined as follows:

$$\mathcal{X} \triangleq \{\boldsymbol{x} \in \mathbb{R}^N \mid A\boldsymbol{x} \leq \boldsymbol{b}\}$$

Show that $\mathcal{X}$ is a convex set.

b) Let $\mathcal{X}$ be a convex subset of $\mathbb{R}^N$ and let $f(\boldsymbol{x})$ be a convex function over $\mathcal{X}$. Define the $N+1$ dimensional set $\mathcal{Z}$ as follows:

$$\mathcal{Z} \triangleq \{(\boldsymbol{x}, y) \mid \boldsymbol{x} \in \mathcal{X}, y \geq f(\boldsymbol{x})\}$$

This set $\mathcal{Z}$ is called the *epigraph* of the function $f(\boldsymbol{x})$. Prove that $\mathcal{Z}$ is a convex set.

c) Plot the epigraph of the function $f(x) = x^2$ over $0 \leq x \leq 1$. (Note that the epigraph will be a set in $\mathbb{R}^2$).

III. Queueing with a Concave Rate-Power Curve

Consider a discrete time queue where unfinished work $U(t)$ has dynamics for slots $t \in \{0, 1, 2, \ldots\}$:

$$U(t+1) = \max[U(t) - \mu(t), 0] + A(t)$$

Assume $A(t)$ is i.i.d. over slots with mean $\mathbb{E}\left\{A(t)\right\} = \lambda$ and second moment $\mathbb{E}\left\{A(t)^2\right\} < \infty$. Suppose that the transmission rate $\mu(t)$ is affected by power allocation decisions every slot t, so that:

$$\mu(t) = C(P(t))$$

where $C(p)$ is a continuous, concave, and strictly increasing function of power (with units of bits/slot), with $C(0) = 0$. Power is constrained every slot t by a maximum power P_{max}, so that we have the following control decision constraint:

$$0 \leq P(t) \leq P_{max} \text{ for all } t$$

a) Suppose we have some power allocation algorithm that chooses $P(t)$ over time and stabilizes the system. Suppose the algorithm has a well defined time average expected power $\overline{P}$:

$$\overline{P} \triangleq \lim_{t\to\infty}\frac{1}{t}\sum_{\tau=0}^{t-1}\mathbb{E}\left\{P(\tau)\right\}$$

Prove that $\overline{P} \geq C^{-1}(\lambda)$. This shows that the minimum possible time-average power required for stability is $C^{-1}(\lambda)$. (Hint: Use the fact that stability implies $\mathbb{E}\left\{U(t)/t\right\} \to 0$, together with Jensen's inequality for concave functions).

b) Suppose we have a *linear* rate-power curve $C(P) = \alpha P$ for $0 \leq P \leq P_{max}$. Thus, $\mu_{max} = \alpha P_{max}$. Design a simple power allocation algorithm that stabilizes the system whenever $\lambda < \mu_{max}$, and always yields the *exact* minimum average power expenditure $\overline{P} = \lambda/\alpha$. Compute a bound on the average congestion $\overline{U}$ and average bit delay $\overline{W} = \overline{U}/\lambda$ for your algorithm. (Hint: Your power allocation algorithm can be so simple that the queue dynamics are the same as a discrete GI/D/1 queue.) (Note: This special case of a purely linear rate-power curve is the only case when we can achieve *exact* minimum average power, with no delay tradeoff).

[2] In this case, it can be shown that "renewals" occur when the system empties on timeslots that are multiples of 11, and so renewal theory can be used to show the regular limit exists and hence is equal to the $\limsup$.

c) Suppose we have $C(P) = \log(1+P)$ for $0 \leq P \leq P_{max}$. Define $\mu_{max} = \log(1+P_{max})$. Suppose $\lambda < \mu_{max}$, and define P_{av}^* as the optimal average power expenditure (so that $P_{av}^* = e^{\lambda} - 1$ by part (a)). In this case, it is impossible to achieve the *exact* value of P_{av}^* for random traffic without incurring infinite average delay. Design a simple power allocation algorithm that, for any given $\epsilon > 0$, stabilizes the system (with finite $\overline{U}$) with time average power $\overline{P} \leq P_{av}^* + \epsilon$. Compute a bound on the average queue backlog $\overline{U}$, and show that it goes to infinity as $\epsilon \to 0$. (Hint: Use a *very simple* power allocation algorithm. Your answer should have $\overline{U} \leq O(1/\epsilon)$). (Note: For random traffic, the *optimal* congestion tradeoff for strictly concave rate-power curves is $O(1/\sqrt{\epsilon})$ [1], and the optimal congestion tradeoff for piecewise linear rate-power curves (with at least two pieces) is $O(\log(1/\epsilon))$ [2]).

IV. Routing to Parallel Queues

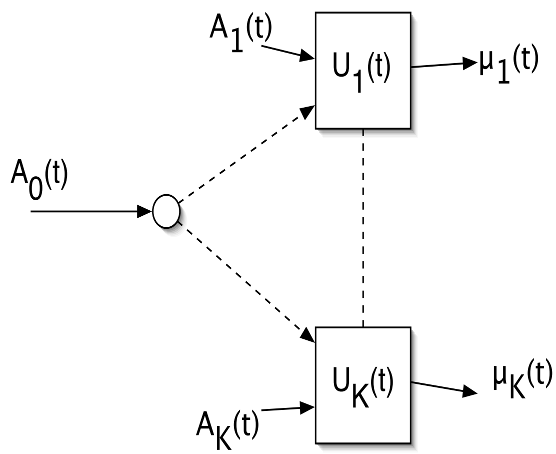

Fig. 1. A problem of routing to K queues.

Consider a system of K parallel queues with individual server rates $\mu_k(t)$ for $k \in \{1, \ldots, K\}$. Let $\boldsymbol{\mu}(t) = (\mu_1(t), \ldots, \mu_K(t))$ be the service rate vector, and suppose that these vectors are i.i.d. over timeslots with $\mathbb{E}\{\mu_k(t)\} = \overline{\mu}_k$ and $\mathbb{E}\{\mu_k(t)^2\} = \mathbb{E}\{\mu_k^2\}$ for all t (note that the individual server rates $\mu_k(t)$ may be *correllated* over different queues k on the same timeslot t). Consider the vector arrival process $\boldsymbol{A}(t) = (A_0(t), A_1(t), \ldots, A_K(t))$ that is i.i.d. over timeslots, such that $\mathbb{E}\{A_i(t)\} = \lambda_i$, $\mathbb{E}\{A_i^2(t)\} = \mathbb{E}\{A_i^2\}$ for all t. The process $A_k(t)$ goes directly into queue k (for $k \in \{1, \ldots, K\}$), and the arrivals $A_0(t)$ must be *routed* to one of the K queues immediately every timeslot. Define the set Λ as the set of all $(\lambda_0, \lambda_1, \ldots, \lambda_K)$ vectors that satisfy:

$$\lambda_k \leq \overline{\mu}_k \text{ for all } k \in \{1, \ldots, K\}$$
$$\lambda_0 + \sum_{k=1}^{K} \lambda_k \leq \sum_{k=1}^{K} \overline{\mu}_k$$

We want to show that Λ is the closure of all rate vectors for which there exists a stabilizing routing policy, and hence Λ is the *capacity region* for this network.

a) Show that the above constraints are *necessary* for strong stability. Start by assuming there exists a routing policy that strongly stabilizes the system, and then prove the rates must satisfy all of the above constraints.

b) Suppose now we have a rate vector $(\lambda_0, \ldots, \lambda_K)$ that is strictly interior to the above region Λ, in the sense that there is some positive value $\epsilon > 0$ such that $(\lambda_0 + \epsilon, \lambda_1 + \epsilon, \ldots, \lambda_K + \epsilon) \in \Lambda$ (i.e., adding ϵ to all entries of the input rate vector yields another vector that is still inside the capacity region). Prove that there exists a stabilizing policy that randomly (and independently every slot t) routes all $A_0(t)$ traffic to queue k with probability p_{0k}, for some probabilities that satisfy $\sum_{k=1}^{K} p_{0k} = 1$. You must explicitly compute the p_{0k} probabilities.

c) Let $\boldsymbol{A}_0(t) = (A_{01}(t), A_{02}(t), \ldots, A_{0K}(t))$ be the routing control vector, representing the amount of bits that the controller decides to route to queue k during slot t (for all $k \in \{1, \ldots, K\}$). Every slot the routing vector $\boldsymbol{A}_0(t)$

must satisfy:

$$A_{0k}(t) \geq 0 \text{ for all } k \in \{1, \ldots, K\} \quad (3)$$

$$\sum_{k=1}^{K} A_{0k}(t) = A_0(t) \quad (4)$$

The queueing dynamics for each queue $k \in \{1, \ldots, K\}$ are thus given by:

$$U_k(t+1) = \max[U_k(t) - \mu_k(t), 0] + A_k(t) + A_{0k}(t)$$

Let $\boldsymbol{U}(t)$ represent the queue backlog vector, and define the quadratic Lyapunov function $L(\boldsymbol{U}) = \sum_{k=1}^{K} U_k^2$. Show that the one-step conditional Lyapunov drift $\Delta(\boldsymbol{U}(t)) \triangleq \mathbb{E}\left\{L(\boldsymbol{U}(t+1)) - L(\boldsymbol{U}(t)) \mid \boldsymbol{U}(t)\right\}$ satisfies:

$$\Delta(\boldsymbol{U}(t)) \leq B - 2\mathbb{E}\left\{\sum_{k=1}^{K} U_k(t)[\overline{\mu}_k - \lambda_k - A_{0k}(t)] \mid \boldsymbol{U}(t)\right\} \quad (5)$$

where $B < \infty$. Compute the value of B. Thus, the right hand side depends on the *routing policy* for choosing $\boldsymbol{A}_0(t) = (A_{01}(t), \ldots, A_{0K}(t))$. (Note that expectations can often be "pushed" or "pulled" as convenient, so that if X and Y are two random variables with $\mathbb{E}\{X\} = \overline{x}$, then $\mathbb{E}\{2(X + Xf(Y))\} = 2\mathbb{E}\{X\} + 2\mathbb{E}\{Xf(Y)\} = 2\overline{x} + 2\mathbb{E}\{Xf(Y)\} = \mathbb{E}\{2(\overline{x} + Xf(Y))\}$.)

d) Define $\Phi(\boldsymbol{A}_0(t), \boldsymbol{U}(t)) \triangleq \sum_{k=1}^{K} U_k(t)[\overline{\mu}_k - \lambda_k - A_{0k}(t)]$. Derive a dynamic routing policy that observes $\boldsymbol{U}(t)$ every slot, and then chooses a routing decision vector $\boldsymbol{A}_0(t)$ to maximize $\Phi(\boldsymbol{A}_0(t), \boldsymbol{U}(t))$ over all possible $\boldsymbol{A}_0(t)$ vectors (given the observed value of $\boldsymbol{U}(t)$). That is, the policy yields decision vectors $\boldsymbol{A}_0^{dyn}(t)$ that satisfy (3)-(4) every slot t, and such that:

$$\Phi(\boldsymbol{A}_0^{dyn}(t), \boldsymbol{U}(t)) \geq \Phi(\boldsymbol{A}_0^{*}(t), \boldsymbol{U}(t))$$

for all $\boldsymbol{U}(t)$ and for any other alternative routing policy that yields a decision vector $\boldsymbol{A}_0^{*}(t)$ that satisfies (3)-(4) every slot. By taking conditional expectations of the above inequality with respect to the random $A_0(t)$ value and the (potentially random) routing decisions, conclude that:

$$\mathbb{E}\left\{\Phi(\boldsymbol{A}_0^{dyn}(t), \boldsymbol{U}(t)) \mid \boldsymbol{U}(t)\right\} \geq \mathbb{E}\left\{\Phi(\boldsymbol{A}_0^{*}(t), \boldsymbol{U}(t)) \mid \boldsymbol{U}(t)\right\}$$

e) Use your stationary randomized routing policy of part (b) to show that there exists a routing policy that chooses decisions $\boldsymbol{A}_0^{*}(t)$ independent of the queue backlog, and that satisfies (3)-(4) every slot, and that $\mathbb{E}\left\{A_{0k}^{*}(t) \mid \boldsymbol{U}(t)\right\} = \lambda_0 p_{0k} \leq \overline{\mu}_k - \lambda_k - \delta$ (for some positive constant $\delta > 0$).

f) Conclude that for every slot t and for all $\boldsymbol{U}(t)$, we have:

$$\Delta^{dyn}(\boldsymbol{U}(t)) \leq B - 2\delta \sum_{k=1}^{K-1} U_k(t)$$

where $\Delta^{dyn}(\boldsymbol{U}(t))$ is the conditional Lyapunov drift from the policy of part (d) that chooses decision vectors $\boldsymbol{A}_0^{dyn}(t)$. Conclude that the policy of part (d) yields a strongly stable system with $\overline{U}_{tot} \leq B/(2\delta)$ whenever the input rates $(\lambda_1, \ldots, \lambda_K)$ are strictly interior to the capacity region Λ. (It turns out that, experimentally, this policy also yields *delay* that is much better than the stationary randomized policy of part (b)).

g) List one or two other advantages/disadvantages of the dynamic policy of part (d) in comparison with the stationary randomized policy of part (b).

V. Simulation of Bernoulli Queues

This is a simulation problem. You can write your own simulation, or you can use the C program on the course webpage.

Consider a packet based B/B/1 queue, where $A(t)$ is i.i.d. over slots with $Pr[A(t) = 1] = \lambda$, $Pr[A(t) = 0] = 1 - \lambda$, and $\mu(t)$ is i.i.d. over slots with $Pr[\mu(t) = 1] = \overline{\mu}$, $Pr[\mu(t) = 0] = 1 - \overline{\mu}$. The queueing equation is $U(t+1) = \max[U(t) - \mu(t), 0] + A(t)$. Assume $U(0) = 0$.

a) Let $\lambda = 0.2, \overline{\mu} = 0.9$. Simulate the system over T timeslots, for $T = 10^6$, to obtain the empirical time average backlog $\overline{U}$. Compare with the exact steady state answer $\overline{U} = \frac{\lambda(1-\lambda)}{\overline{\mu}-\lambda}$. Note: You may want to use double-precision real numbers when computing the time average.

b) For $\mu = 0.9$ and $T = 10^6$, fill in the following chart:

c) For $\mu = 0.9$, fill in the following chart:

λ	Simulation	Exact Steady State Value
0.2		0.228571
0.3		
0.4		
0.5		
0.6		
0.7		
0.8		
0.85		
0.89		

λ	$T = 10^5$	$T = 10^6$	$T = 10^7$	Exact Steady State Value
0.5				
0.9				
0.91				

VI. Practice Problem (ungraded — do not turn in): ON/OFF Sources

Consider a discrete time Markov chain $M(t) \in \{ON, OFF\}$. The transition probabilities are given by ϵ and δ, as shown in the figure.

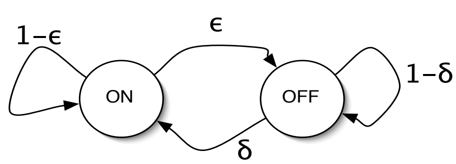

Fig. 2. A 2-state Markov chain in discrete time.

Let $p_{ON}(t)$ represent the probability that $M(t) = ON$ during slot t, and let $p_{OFF}(t) = 1 - p_{ON}(t)$ represent the probability that $M(t) = OFF$.

a) Compute $p_{ON}(t+1)$ in terms of $p_{ON}(t)$.

b) Show that $p_{ON}(t) = A + B\gamma^t$ solves your equation of (a), for some constants A, B, γ. Suppose that $p_{ON}(0) = 1$ (so the chain is initially ON). Compute $p_{ON}(t)$ exactly as a function of t for $t \in \{0, 1, 2, \ldots\}$.

c) Suppose we have an arrival process $A(t)$ such that $A(t) = 1$ if $M(t) = ON$, and $A(t) = 0$ if $M(t) = OFF$. Show that $A(t)$ is admissible with rate λ, and compute the value λ. For a fixed value $\sigma > 0$, compute an integer T_σ that ensures for any time t_0:

$$\frac{1}{T_\sigma} \sum_{k=0}^{T_\sigma - 1} \mathbb{E}\left\{A(t_0 + k) \mid H(t_0)\right\} \leq \lambda + \sigma$$

Hint: Note that $\mathbb{E}\left\{A(t_0 + k) \mid H(t_0)\right\} = \mathbb{E}\left\{\mathbb{E}\left\{A(t_0 + k) \mid H(t_0), M(t_0)\right\}\right\}$.

d) Using the expression from chapter 3 in the text (for T-slot Lyapunov analysis), state an upper bound on the time average queue backlog for a queue with this type of input and with a Bernoulli server with rate $\overline{\mu} > \lambda$.

e) How well does this bound compare with simulation?

References

[1] R. Berry and R. Gallager. Communication over fading channels with delay constraints. *IEEE Transactions on Information Theory*, vol. 48, no. 5, pp. 1135-1149, May 2002.

[2] M. J. Neely. Optimal energy and delay tradeoffs for multi-user wireless downlinks. *IEEE Transactions on Information Theory*, vol. 53, no. 9, pp. 3095-3113, Sept. 2007.