1 Foster-Lyapunov stability criterion and moment bounds

These are supplementary notes for ECE 567, Fall 2006. -B. Hajek, instructor, 10/18/06

1.1 Introduction

Communication network models can become quite complex, especially when dynamic scheduling, congestion, and physical layer effects such as fading wireless channel models are included. It is thus useful to have methods to give approximations or bounds on key performance parameters. The criteria for stability and related moment bounds discussed in this chapter are useful for providing such bounds.

Aleksandr Mikhailovich Lyapunov (1857-1918) contributed significantly to the theory of stability of dynamical systems. Although a dynamical system may evolve on a complicated, multiple dimensional state space, a recurring theme of dynamical systems theory is that stability questions can often be settled by studying the potential of a system for some nonnegative potential function V. Potential functions used for stability analysis are widely called Lyapunov functions. Similar stability conditions have been developed by many authors for stochastic systems. Below we present the well known criteria due to Foster [3] for recurrence and positive recurrence. In addition we present associated bounds on the moments, which are expectations of some functions on the state space, computed with respect to the equilibrium probability distribution.[1]

Section 1.2 discusses the discrete time tools, and presents examples involving load balancing routing, and input queued crossbar switches. Section 1.3 presents the continuous time tools, and an example. Problems are given in Section 1.4.

1.2 Stability criteria for discrete time processes

Consider an irreducible discrete-time Markov process X on a countable state space $\mathcal{S}$, with one-step transition probability matrix P. If f is a function on $\mathcal{S}$, then Pf represents the function obtained by multiplication of the vector f by the matrix P: $Pf(i) = \sum_{j\in\mathcal{S}} p_{ij} f(j)$. If f is nonnegative, then Pf is well defined, with the understanding that $Pf(i) = +\infty$ is possible for some, or all, values of i. An important property of Pf is that $Pf(i) = E[f(X(t+1)|X(t) = i]$. Let V be a nonnegative function on $\mathcal{S}$, to serve as the Lyapunov function. The drift function of $V(X(t))$ is defined by $d(i) = E[V(X(t+1))|X(t) = i] - V(i)$. That is, $d = PV - V$. Note that $d(i)$ is always well-defined, if the value $+\infty$ is permitted. The drift function is also given by

$$d(i) = \sum_{j:j\neq i} p_{ij}(V(j) - V(i)). \tag{1}$$

[1]The proof of Foster's criteria given here is similar to Foster's original proof, but is geared to establishing the moment bounds, and the continuous time versions, at the same time. The moment bounds and proofs given here are adaptations of those in Meyn and Tweedie [8] to discrete state spaces, and to continuous time. They involve some basic notions from martingale theory, which can be found, for example, in [4]. A special case of the moment bounds was given by Tweedie [12], and a version of the moment bound method was used by Kingman [5] in a queueing context. As noted in [8], the moment bound method is closely related to Dynkin's formula. The works [6, 9–11], and many others, have demonstrated the wide applicability of the stability methods in various queueing network contexts, using quadratic Lyapunov functions.

Proposition 1.1 *(Foster-Lyapunov stability criterion) Suppose $V : \mathcal{S} \to \mathbb{R}_+$ and C is a finite subset of $\mathcal{S}$.*
(a) If $\{i : V(i) \leq K\}$ is finite for all K, and if $PV - V \leq 0$ on $\mathcal{S} - C$, then X is recurrent.
(b) If $\epsilon > 0$ and b is a constant such that $PV - V \leq -\epsilon + bI_C$, then X is positive recurrent.

The proof of the proposition is given after two lemmas.

Lemma 1.2 *Suppose $PV - V \leq -f + g$ on $\mathcal{S}$, where f and g are nonnegative functions. Then for any initial state i_o and any stopping time τ,*

$$E\left[\sum_{k:0\leq k\leq \tau-1} f(X(k))\right] \leq V(i_o) + E\left[\sum_{k:0\leq k\leq \tau-1} g(X(k))\right] \tag{2}$$

Proof. Let $\mathcal{F}_k$ denote the σ-algebra generated by $(X(s) : 0 \leq s \leq k)$. By assumption, $PV + f \leq V + g$ on $\mathcal{S}$. Evaluating each side of this inequality at $X(k)$, and taking the conditional expectation given $\mathcal{F}_k$ yields

$$E[V(x(k+1))|\mathcal{F}_k] + f(X(k)) \leq V(X(k)) + g(X(k)). \tag{3}$$

Let $\tau^n = \min\{\tau, n, \inf\{k \geq 0 : V(X(k)) \geq n\}\}$. The next step is to multiply each side of (3) by $I_{\{\tau^n > k\}}$, take expectations of each side, and use the fact that

$$E\left[E[V(x(k+1))|\mathcal{F}_k]I_{\{\tau^n>k\}}\right] = E\left[V(x(k+1))I_{\{\tau^n>k\}}\right] \geq E\left[V(x(k+1))I_{\{\tau^n>k+1\}}\right].$$

The result is

$$E\left[V(x(k+1))I_{\{\tau^n>k+1\}}\right] + E\left[f(x(k))I_{\{\tau^n>k\}}\right] \leq E\left[V(x(k))I_{\{\tau^n>k\}}\right] + E\left[g(x(k))I_{\{\tau^n>k\}}\right] \tag{4}$$

The definition of τ^n implies that all the terms in (4) are zero for $k \geq n$, and that $E\left[V(x(k))I_{\{\tau^n>k\}}\right] < \infty$ for all k. Thus, it is legitimate to sum each side of (4) and cancel like terms, to yield:

$$E\left[\sum_{k:0\leq k\leq \tau^n-1} f(X(k))\right] \leq V(i_o) + E\left[\sum_{k:0\leq k\leq \tau^n-1} g(X(k))\right]. \tag{5}$$

Letting $n \to \infty$ in (5) and appealing to the monotone convergence theorem yields (2). ■

Lemma 1.3 *Let X be an irreducible, time-homogeneous Markov process. If there is a finite set C, and the mean time to hit C starting from any state in C is finite, then X is positive recurrent.*

Proof. Lemma I.3.9 of Assussen. Variation of Wald's lemma. ■

Proof of Proposition 1.1. (a) The proof of part (a) is based on a martingale convergence theorem. Fix an initial state $i_o \in \mathcal{S} - C$, let $\tau = \min\{t \geq 0 : X(t) \in C\}$, and let $Y(t) = V(X(t \wedge \tau))$ for $t \geq 0$. Note that $Y(t+1) - Y(t) = (V(X(t+1) - V(X(t)))I_{\{\tau>t\}}$. Let $\mathcal{F}_t$ denote the σ-algebra generated by $(X(s) : 0 \leq s \leq t)$, and note that the event $\{\tau > t\}$ is in $\mathcal{F}_t$. Thus,

$$E[Y(t+1)-Y(t)|\mathcal{F}_t] = E[(V(X(t+1))-V(X(t)))I_{\{\tau>t\}}|\mathcal{F}_t] = E[V(X(t+1))-V(X(t))|\mathcal{F}_t]I_{\{\tau>t\}} \geq 0,$$

so that $(Y_t : t \geq 0)$ is a nonnegative supermartingale. Therefore, by a version of the martingale convergence theorem, $\lim_{t\to\infty} Y(t)$ exists and is finite with probability one. By the assumption $\{i : V(i) \leq K\}$ is finite for all K, it follows that $X(t \wedge \tau)$ reaches only finitely many states with probability one, which implies that $\tau < +\infty$ with probability one. Therefore, $P[\tau < \infty | X(0) = i_o] = 1$ for any $i_o \in \mathcal{S} - C$. Therefore, for any initial state in C, the process returns to C infinitely often with probability one. The process watched in C is a finite-state Markov process, and is thus (positive) recurrent. Therefore, all the states of C are recurrent for the original process X. Part (a) is proved.

(b) Let $f \equiv -\epsilon$, $g = bI_C$, and $\tau = \min\{t \geq 1 : X(t) \in C\}$. Then Lemma 1.2 implies that $\epsilon E[\tau] \leq V(i_o) + b$ for any $i_o \in \mathcal{S}$. In particular, the mean time to hit C after time zero, beginning from any state in C, is finite. Therefore X is positive recurrent by Lemma 1.3. ■

Proposition 1.4 *(Moment bound) Suppose V, f, and g are nonnegative functions on $\mathcal{S}$ and suppose*

$$PV(i) - V(i) \leq -f(i) + g(i) \qquad \text{for all } i \in \mathcal{S} \tag{6}$$

In addition, suppose X is positive recurrent, so that the means, $\overline{f} = \pi f$ and $\overline{g} = \pi g$ are well-defined. Then $\overline{f} \leq \overline{g}$. (In particular, if g is bounded, then $\overline{g}$ is finite, and therefore $\overline{f}$ is finite.)

Proof. Fix a state i_o and let T_m be the time of the m^{th} return to state i_o. Then by the equality of time and statistical averages,

$$E\left[\sum_{\{k:0\leq k\leq T_m-1\}} f(X(k))\right] = mE[T_1]\overline{f}$$

$$E\left[\sum_{\{k:0\leq k\leq T_m-1\}} g(X(k))\right] = mE[T_1]\overline{g}$$

Lemma 1.2 applied with stopping time T_m thus yields $mE[T_1]\overline{f} \leq V(i_o) + mE[T_1]\overline{g}$. Dividing through by $mE[T_1]$ and letting $m \to \infty$ yields the desired inequality, $\overline{f} \leq \overline{g}$. ■

Corollary 1.5 *(Combined Foster-Lyapunov stability criterion and moment bound) Suppose V, f, and g are nonnegative functions on $\mathcal{S}$ such that*

$$PV(i) - V(i) \leq -f(i) + g(i) \qquad \text{for all } i \in \mathcal{S} \tag{7}$$

In addition, suppose for some $\epsilon > 0$ that the set C defined by $C = \{i : f(i) < g(i) + \epsilon\}$ is finite. Then X is positive recurrent and $\overline{f} \leq \overline{g}$. (In particular, if g is bounded, then $\overline{g}$ is finite, and therefore $\overline{f}$ is finite.)

Proof. Let $b = \max\{g(i) + \epsilon - f(i) : i \in C\}$. Then V, C, b, and ϵ satisfy the hypotheses of Proposition 1.1(b), so that X is positive recurrent. Therefore the hypotheses of Proposition 1.4 are satisfied, so that $\overline{f} \leq \overline{g}$. ■

The assumptions in Propositions 1.1 and 1.4 and Corollary 1.5 do not imply that $\overline{V}$ is finite. Even so, since V is nonnegative, for a given initial state $X(0)$, the long term average drift of $V(X(t))$

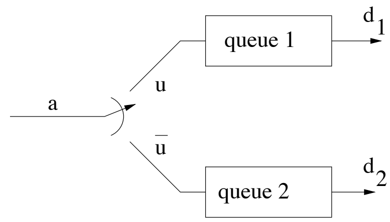

Figure 1: Two queues fed by a single arrival stream.

is nonnegative. This gives an intuitive reason why the mean downward part of the drift, $\overline{f}$, must be less than or equal to the mean upward part of the drift, $\overline{g}$.

Example 1a (Probabilistic routing to two queues) Consider the routing scenario with two queues, queue 1 and queue 2, fed by a single stream of packets, as pictured in Figure 1. Here, $0 \leq a, u, d_1, d_2 \leq 1$, and $\overline{u} = 1 - u$. The state space for the process is $\mathcal{S} = \mathbb{Z}_+^2$, where the state $x = (x_1, x_2)$ denotes x_1 packets in queue 1 and x_2 packets in queue 2. In each time slot, a new arrival is generated with probability a, and then is routed to queue 1 with probability u and to queue 2 with probability $\overline{u}$. Then each queue i, if not empty, has a departure with probability d_i. Note that we allow a packet to arrive and depart in the same slot. Thus, if $X_i(t)$ is the number of packets in queue i at the beginning of slot t, then the system dynamics can be described as follows:

$$X_i(t+1) = X_i(t) + A_i(t) - D_i(t) + L_i(t) \quad \text{for } i \in \{0, 1\} \tag{8}$$

where

- $A(t) = (A_1(t), A_2(t))$ is equal to $(1, 0)$ with probability au, $(0, 1)$ with probability $a\overline{u}$, and $A(t) = (0, 0)$ otherwise.
- $D_i(t) : t \geq 0$, are *Bernoulli*(d_i) random variables, for $i \in \{0, 1\}$
- All the $A(t)$'s, $D_1(t)$'s, and $D_2(t)$'s are mutually independent
- $L_i(t) = (-(X_i(t) + A_i(t) - D_i(t)))_+$ (see explanation next)

If $X_i(t) + A_i(t) = 0$, there can be no actual departure from queue i. However, we still allow $D_i(t)$ to equal one. To keep the queue length process from going negative, we add the random variable $L_i(t)$ in (8). Thus, $D_i(t)$ is the *potential* number of departures from queue i during the slot, and $D_i(t) - L_i(t)$ is the actual number of departures. This completes the specification of the one-step transition probabilities of the Markov process.

A necessary condition for positive recurrence is, for any routing policy, $a < d_1 + d_2$, because the total arrival rate must be less than the total depature rate. We seek to show that this necessary condition is also sufficient, under the random routing policy.

Let us calculate the drift of $V(X(t))$ for the choice $V(x) = (x_1^2 + x_2^2)/2$. Note that $(X_i(t+1))^2 = (X_i(t) + A_i(t) - D_i(t) + L_i(t))^2 \leq (X_i(t) + A_i(t) - D_i(t))^2$, because addition of the variable $L_i(t)$

can only push $X_i(t) + A_i(t) - D_i(t)$ closer to zero. Thus,

$$\begin{aligned}
PV(x) - V(x) &= E[V(X(t+1))|X(t) = x] - V(x) \\
&\leq \frac{1}{2}\sum_{i=1}^{2} E[(x_i + A_i(t) - D_i(t))^2 - x_i^2|X(t) = x] \\
&= \sum_{i=1}^{2} x_i E[A_i(t) - D_i(t)|X(t) = x] + \frac{1}{2}E[(A_i(t) - D_i(t))^2|X(t) = x] \qquad (9)\\
&\leq \left(\sum_{i=1}^{2} x_i E[A_i(t) - D_i(t)|X(t) = x]\right) + 1 \\
&= -\left(x_1(d_1 - au) + x_2(d_2 - a\overline{u})\right) + 1 \qquad (10)
\end{aligned}$$

Under the necessary condition $a < d_1 + d_2$, there are choices of u so that $au < d_1$ and $a\overline{u} < d_2$, and for such u the conditions of Corollary 1.5 are satisfied, with $f(x) = x_1(d_1 - au) + x_2(d_2 - a\overline{u})$, $g(x) = 1$, and any $\epsilon > 0$, implying that the Markov process is positive recurrent. In addition, the first moments under the equlibrium distribution satisfy:

$$(d_1 - au)\overline{X_1} + (d_2 - a\overline{u})\overline{X}_2 \leq 1. \qquad (11)$$

In order to deduce an upper bound on $\overline{X}_1 + \overline{X}_2$, we select u^* to maximize the minimum of the two coefficients in (11). Intuitively, this entails selecting u to minimize the absolute value of the difference between the two coefficients. We find:

$$\begin{aligned}
\epsilon &= \max_{0\leq u \leq 1} \min\{d_1 - au, d_2 - a\overline{u}\} \\
&= \min\{d_1, d_2, \frac{d_1 + d_2 - a}{2}\}
\end{aligned}$$

and the corresponding value u^* of u is given by

$$u^* = \begin{cases} 0 & \text{if } \ d_1 - d_2 \ < -a \\ \frac{1}{2} + \frac{d_1 - d_2}{2a} & \text{if } |d_1 - d_2| \leq a \\ 1 & \text{if } \ \ d_1 - d_2 > a \end{cases}$$

For the system with $u = u^*$, (11) yields

$$\overline{X}_1 + \overline{X}_2 \leq \frac{1}{\epsilon}. \qquad (12)$$

We remark that, in fact,

$$\overline{X}_1 + \overline{X}_2 \leq \frac{2}{d_1 + d_2 - a} \qquad (13)$$

If $|d_1 - d_2| \leq a$ then the bounds (12) and (13) coincide, and otherwise, the bound (13) is strictly tighter. If $d_1 - d_2 < -a$ then $u^* = 0$, so that $\overline{X_1} = 0$, and (11) becomes $(d_2 - a)\overline{X_2} \leq 1$, which implies (13). Similarly, if $d_1 - d_2 > a$, then $u^* = 1$, so that $\overline{X_2} = 0$, and (11) becomes $(d_1 - a)\overline{X_1} \leq 1$, which implies (13). Thus, (13) is proved.

Example 1b (Route-to-shorter policy) Consider a variation of the previous example such that when a packet arrives, it is routed to the shorter queue. To be definite, in case of a tie, the packet is routed to queue 1. Then the evolution equation (8) still holds, but with with the description of the arrival variables changed to the following:

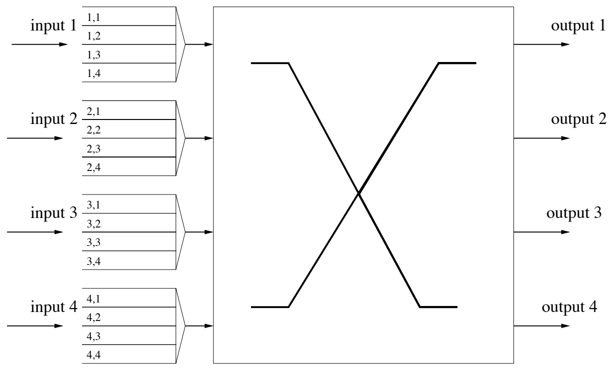

Figure 2: A 4×4 input queued switch

- Given $X(t) = (x_1, x_2)$, $A(t) = (I_{\{x_1 \leq x_2\}}, I_{\{x_1 > x_2\}})$ with probability a, and $A(t) = (0, 0)$ otherwise. Let P^{RS} denote the one-step transition probability matrix when the route-to-shorter policy is used.

Proceeding as in (9) yields:

$$\begin{aligned} P^{RS}V(x) - V(x) &\leq \sum_{i=1}^{2} x_i E[A_i(t) - D_i(t)) | X(t) = x] + 1 \\ &= a \left(x_1 I_{\{x_1 \leq x_2\}} + x_2 I_{\{x_1 > x_2\}} \right) - d_1 x_1 - d_2 x_2 + 1 \end{aligned}$$

Note that $x_1 I_{\{x_1 \leq x_2\}} + x_2 I_{\{x_1 > x_2\}} \leq u x_1 + \overline{u} x_2$ for any $u \in [0, 1]$, with equality for $u = I_{\{x_1 \leq x_2\}}$. Therefore, the drift bound for V under the route-to-shorter policy is less than or equal to the drift bound (10), for V for any choice of probabilistic splitting. In fact, route-to-shorter routing can be viewed as a controlled version of the independent splitting model, for which the control policy is selected to minimize the bound on the drift of V in each state. It follows that the route-to-shorter process is positive recurrent as long as $a < d_1 + d_2$, and (11) holds for any value of u such that $au < d_1$ and $a\overline{u} \leq d_2$. In particular, (12) holds for the route-to-shorter process.

We remark that the stronger bound (13) is not always true for the route-to-shorter policy. The problem is that even if $d_1 - d_2 < -a$, the route-to-shorter policy can still route to queue 1, and so $\overline{X}_1 \neq 0$. In fact, if a and d_2 are fixed with $0 < a < d_2 < 1$, then $\overline{X_1} \to \infty$ as $d_1 \to 0$ for the route-to-shorter policy. Intuitively, that is because occasionally there will be a large number of customers in the system due to statistical fluctuations, and then there will be many customers in queue 1. But if $d_2 << 1$, those customers will remain in queue 2 for a very long time.

Example 2a (An input queued switch with probabilistic switching) [2] Consider a packet switch with N inputs and N outputs, as pictured in Figure 2. Suppose there are N^2 queues – N at each input – with queue i, j containing packets that arrived at input i and are destined for output j, for $i, j \in E$, where $E = \{1, \cdots, N\}$. Suppose the packets are all the same length, and adopt a discrete time model, so that during one time slot, a transfer of packets can occur, such that at

[2] Tassiulas [11] originally developed the results of Examples 2a-b, in the context of wireless networks. The paper [7] presents similiar results in the context of a packet switch.

most one packet can be transferred from each input, and at most one packet can be transferred to each output. A permutation σ of E has the form $\sigma = (\sigma_1, \ldots, \sigma_N)$, where $\sigma_1, \ldots, \sigma_N$ are distinct elements of E. Let Π denote the set of all $N!$ such permutations. Given $\sigma \in \Pi$, let $R(\sigma)$ be the $N \times N$ *switching matrix* defined by $R_{ij} = I_{\{\sigma_i = j\}}$. Thus, $R_{ij}(\sigma) = 1$ means that under permutation σ, input i is connected to output j, or, equivalently, a packet in queue i,j is to depart, if there is any such packet. A state x of the system has the form $x = (x_{ij} : i,j \in E)$, where x_{ij} denotes the number of packets in queue i,j.

The evolution of the system over a time slot $[t, t+1)$ is described as follows:

$$X_{ij}(t+1) = X_{ij}(t) + A_{ij}(t) - R_{ij}(\sigma(t)) + L_{ij}(t)$$

where

- $A_{ij}(t)$ is the number of packets arriving at input i, destined for output j, in the slot. Assume that the variables $(A_{ij}(t) : i,j \in E, t \geq 0)$ are mutually independent, and for each i,j, the random variables $(A_{ij}(t) : t \geq 0)$ are independent, identically distributed, with mean λ_{ij} and $E[A_{ij}^2] \leq K_{ij}$, for some constants λ_{ij} and K_{ij}. Let $\Lambda = (\lambda_{ij} : i,j \in E)$.

- $\sigma(t)$ is the switch state used during the slot

- $L_{ij} = (-(X_{ij}(t) + A_{ij}(t) - R_{ij}(\sigma(t)))_+$, which takes value one if there was an unused potential departure at queue ij during the slot, and is zero otherwise.

The number of packets at input i at the beginning of the slot is given by the row sum $\sum_{j \in E} X_{ij}(t)$, its mean is given by the row sum $\sum_{j \in E} \lambda_{ij}$, and at most one packet at input i can be served in a time slot. Similarly, the set of packets waiting for output j, called the *virtual queue* for output j, has size given by the column sum $\sum_{i \in E} X_{ij}(t)$. The mean number of arrivals to the virtual queue for output j is $\sum_{i \in E} \lambda_{ij}(t)$, and at most one packet in the virtual queue can be served in a time slot. These considerations lead us to impose the following restrictions on Λ:

$$\sum_{j \in E} \lambda_{ij} < 1 \text{ for all } i \quad \text{and} \quad \sum_{i \in E} \lambda_{ij} < 1 \text{ for all } j \tag{14}$$

Except for trivial cases involving deterministic arrival sequences, the conditions (14) are necessary for stable operation, for any choice of the switch schedule $(\sigma(t) : t \geq 0)$.

Let's first explore random, independent and identically distributed (i.i.d.) switching. That is, given a probability distribution u on Π, let $(\sigma(t) : t \geq 0)$ be independent with common probability distribution u. Once the distributions of the A_{ij}'s and u are fixed, we have a discrete-time Markov process model. Given Λ satisfying (14), we wish to determine a choice of u so that the process with i.i.d. switch selection is positive recurrent.

Some standard background from switching theory is given in this paragraph. A *line sum* of a matrix M is either a row sum, $\sum_j M_{ij}$, or a column sum, $\sum_i M_{ij}$. A square matrix M is called *doubly stochastic* if it has nonnegative entries and if all of its line sums are one. Birkhoff's theorem, celebrated in the theory of switching, states that any doubly stochastic matrix M is a convex combination of switching matrices. That is, such an M can be represented as $M = \sum_{\sigma \in \Pi} R(\sigma) u(\sigma)$, where $u = (u(\sigma) : \sigma \in \Pi)$ is a probability distribution on Π. If $\widetilde{M}$ is a nonnegative matrix with all line sums less than or equal to one, then if some of the entries of $\widetilde{M}$ are increased appropriately, a doubly stochastic matrix can be obtained. That is, there exists a doubly stochastic matrix M

so that $\widetilde{M}_{ij} \le M_{ij}$ for all i, j. Applying Birkhoff's theorem to M yields that there is a probability distribution u so that $\widetilde{M}_{ij} \le \sum_{\sigma\in\Pi} R(\sigma)u(\sigma)$ for all i, j.

Suppose Λ satisfies the necessary conditions (14). That is, suppose that all the line sums of Λ are less than one. Then with ϵ defined by

$$\epsilon = \frac{1 - (\text{maximum line sum of } \Lambda)}{N},$$

each line sum of $(\lambda_{ij} + \epsilon : i, j \in E)$ is less than or equal to one. Thus, by the observation at the end of the previous paragraph, there is a probability distribution u^* on Π so that $\lambda_{ij} + \epsilon \le \mu_{ij}(u^*)$, where

$$\mu_{ij}(u) = \sum_{\sigma\in\Pi} R_{ij}(\sigma)u(\sigma).$$

We consider the system using probability distribution u^* for the switch states. That is, let $(\sigma(t) : t \ge 0)$ be independent, each with distribution u^*. Then for each ij, the random variables $R_{ij}(\sigma(t))$ are independent, Bernoulli($\mu_{ij}(u^*)$) random variables.

Consider the quadratic Lyapunov function V given by $V(x) = \frac{1}{2}\sum_{i,j} x_{ij}^2$. As in (9),

$$PV(x) - V(x) \le \sum_{i,j} x_{ij} E[A_{ij}(t) - R_{ij}(\sigma(t))|X_{ij}(t) = x] + \frac{1}{2}\sum_{i,j} E[(A_{ij}(t) - R_{ij}(\sigma(t)))^2|X(t) = x].$$

Now

$$E[A_{ij}(t) - R_{ij}(\sigma(t))|X_{ij}(t) = x] = E[A_{ij}(t) - R_{ij}(\sigma(t))] = \lambda_{ij} - \mu_{ij}(u^*) \le -\epsilon$$

and

$$\frac{1}{2}\sum_{i,j} E[(A_{ij}(t) - R_{ij}(\sigma(t)))^2|X(t) = x] \le \frac{1}{2}\sum_{i,j} E[(A_{ij}(t))^2 + (R_{ij}(\sigma(t)))^2] \le K$$

where $K = \frac{1}{2}(N + \sum_{i,j} K_{ij})$. Thus,

$$PV(x) - V(x) \le -\epsilon \left(\sum_{ij} x_{ij}\right) + K \tag{15}$$

Therefore, by Corollary 1.5, the process is positive recurrent, and

$$\sum_{ij} \overline{X}_{ij} \le \frac{K}{\epsilon} \tag{16}$$

That is, the necessary condition (14) is also sufficient for positive recurrence and finite mean queue length in equilibrium, under i.i.d. random switching, for an appropriate probability distribution u^* on the set of permutations.

Example 2b (An input queued switch with maximum weight switching) The random switching policy used in Example 2a depends on the arrival rate matrix Λ, which may be unknown a priori. Also, the policy allocates potential departures to a given queue ij, whether or not the queue is empty, even if other queues could be served instead. This suggests using a dynamic switching policy, such as the *maximum weight* switching policy, defined by $\sigma(t) = \sigma^{MW}(X(t))$, where for a state x,

$$\sigma^{MW}(x) = \arg\max_{\sigma\in\Pi} \sum_{ij} x_{ij} R_{ij}(\sigma). \tag{17}$$

The use of "arg max" here means that $\sigma^{MW}(x)$ is selected to be a value of σ that maximizes the sum on the right hand side of (17), which is the weight of permutation σ with edge weights x_{ij}. In order to obtain a particular Markov model, we assume that the set of permutations Π is numbered from 1 to $N!$ in some fashion, and in case there is a tie between two or more permutations for having the maximum weight, the lowest numbered permutation is used. Let P^{MW} denote the one-step transition probability matrix when the route-to-shorter policy is used.

Letting V and K be as in Example 2a, we find under the maximum weight policy that

$$P^{MW}V(x) - V(x) \le \sum_{ij} x_{ij}(\lambda_{ij} - R_{ij}(\sigma^{MW}(x))) + K$$

The maximum of a function is greater than or equal to the average of the function, so that for any probability distribution u on Π

$$\begin{aligned} \sum_{ij} x_{ij} R_{ij}(\sigma^{MW}(t)) &\ge \sum_{\sigma} u(\sigma) \sum_{ij} x_{ij} R_{ij}(\sigma) \\ &= \sum_{ij} x_{ij} \mu_{ij}(u) \end{aligned} \tag{18}$$

with equality in (18) if and only if u is concentrated on the set of maximum weight permutations. In particular, the choice $u = u^*$ shows that

$$\sum_{ij} x_{ij} R_{ij}(\sigma^{MW}(t)) \ge \sum_{ij} x_{ij}\mu_{ij}(u*) \ge \sum_{ij} x_{ij}(\lambda_{ij} + \epsilon)$$

Therefore, if P is replaced by P^{MW}, (15) still holds. Therefore, by Corollary 1.5, the process is positive recurrent, and the same moment bound, (16), holds, as for the randomized switching strategy of Example 2a. On one hand, implementing the maximum weight algorithm does not require knowledge of the arrival rates, but on the other hand, it requires that queue length information be shared, and that a maximization problem be solved for each time slot. Much recent work has gone towards reduced complexity dynamic switching algorithms.

1.3 Stability criteria for continuous time processes

Here is a continuous time version of the Foster-Lyapunov stability criteria and the moment bounds. Suppose X is a time-homegeneous, irreducible, continuous-time Markov process with generator matrix Q. The drift vector of $V(X(t))$ is the vector QV. This definition is motivated by the fact that the mean drift of X for an interval of duration h is given by

$$\begin{aligned} d_h(i) &= \frac{E[V(X(t+h))|X(t)=i] - V(i)}{h} \\ &= \sum_{j\in\mathcal{S}} \left(\frac{p_{ij}(h) - \delta_{ij}}{h}\right) V(j) \\ &= \sum_{j\in\mathcal{S}} \left(q_{ij} + \frac{o(h)}{h}\right) V(j), \end{aligned} \tag{19}$$

so that if the limit as $h \to 0$ can be taken inside the summation in (19), then $d_h(i) \to QV(i)$ as $h \to 0$. The following useful expression for QV follows from the fact that the row sums of Q are zero:

$$QV(i) = \sum_{j:j \neq i} q_{ij}(V(j) - V(i)). \tag{20}$$

Formula (20) is quite similar to the formula (1) for the drift vector for a discrete-time process.

Proposition 1.6 *(Foster-Lyapunov stability criterion–continuous time) Suppose $V : \mathcal{S} \to \mathbb{R}_+$ and C is a finite subset of $\mathcal{S}$.*
(a) If $QV \leq 0$ on $\mathcal{S} - C$, and $\{i : V(i) \leq K\}$ is finite for all K then X is recurrent.
(b) Suppose for some $b > 0$ and $\epsilon > 0$ that

$$QV(i) \leq -\epsilon + bI_C(i) \qquad \text{for all } i \in \mathcal{S}. \tag{21}$$

Suppose further that $\{i : V(i) \leq K\}$ is finite for all K, or that X is nonexplosive. Then X is positive recurrent.

The proof of the proposition is similar to that above for discrete time processes. We begin with an analog of Lemma 1.2.

Lemma 1.7 *Suppose $QV \leq -f + g$ on $\mathcal{S}$, where f and g are nonnegative functions. Fix an initial state i_o, let N be a stopping time for the jump process X^J, and let τ^N denote the time of the N^{th} jump of X. Then*

$$E\left[\int_0^{\tau^N} f(X(t))dt\right] \leq V(i_o) + E\left[\int_0^{\tau^N} g(X(t))dt\right] \tag{22}$$

Proof. Let D denote the diagonal matrix with entries $-q_{ii}$. The one-step transition probability matrix of the jump chain is given by $P^J = D^{-1}Q + I$. The condition $QV \leq -f + g$ thus implies $P^J V - V \leq -\widetilde{f} + \widetilde{g}$, were $\widetilde{f} = D^{-1}f$ and $\widetilde{g} = D^{-1}g$. Lemma 1.2 applied to the jump chain thus yields

$$E\left[\sum_{k:0\leq k \leq \tau-1} \widetilde{f}(X^J(k))\right] \leq V(i_o) + E\left[\sum_{k:0\leq k\leq \tau-1} \widetilde{g}(X^J(k))\right] \tag{23}$$

However, by the space-time description of an excursion of X from i_o,

$$E\left[\sum_{k:0\leq k\leq N-1} \widetilde{f}(X^J(k))\right] = E\left[\int_0^{\tau^N} f(X(t))dt\right]$$

and a similar equation holds for $\widetilde{g}$ and g. Substituting these relations into (23) yields (22), as desired. ■

Proof of Proposition 1.6. (a) Let D denote the diagonal matrix with entries $-q_{ii}$. The one-step transition probability matrix of the jump chain is given by $P^J = D^{-1}Q + I$. The condition $QV \leq 0$ thus implies $P^J V - V \leq 0$, so that X^J is recurrent by Proposition 1.1(a). But X^J is recurrent if and only if X is recurrent, so X is also recurrent.

(b) The assumptions imply $QV \leq 0$ on $\mathcal{S} - C$, so if $\{i : V(i) \leq K\}$ is finite for all K then X is recurrent by Proposition 1.6(a), and, in particular, X is nonexplosive. Thus, X is nonexplosive, either by direct assumption or by implication.

Let $f \equiv -\epsilon$, let $g = bI_C$, let $i_o \in C$, let $N = \min\{k \geq 1 : X^J(k) \in C\}$, and let τ_N denote the time of the N^{th} jump. Then Lemma 1.7 implies that $\epsilon E[\tau^N] \leq V(i_o) + b/q_{i_o i_o}$. Since τ_N is finite a.s. and since X is not explosive, it must be that τ_N is the time that X returns to the set C after exiting the initial state i_o. That is, the time to hit C beginning from any state in C is finite. Therefore X is positive recurrent by the continuous time version of Lemma 1.3. ■

Example 3 Suppose X has state space $\mathcal{S} = \mathbb{Z}_+$, with $q_{i0} = \mu$ for all $i \geq 1$, $q_{ii+1} = \lambda_i$ for all $i \geq 0$, and all other off-diagonal entries of the rate matrix Q equal to zero, where $\mu > 0$ and $\lambda_i > 0$ such that $\sum_{i\geq 0} \frac{1}{\lambda_i} < +\infty$. Let $C = \{0\}$, $V(0) = 0$, and $V(i) = 1$ for $i \geq 0$. Then $QV = -\mu + (\lambda_0 + \mu)I_C$, so that (21) is satisfied with $\epsilon = \mu$ and $b = \lambda_0 + \mu$. However, X is not positive recurrent. In fact, X is explosive. To see this, note that $p^J_{ii+1} = \frac{\lambda_i}{\mu+\lambda_i} \geq \exp(-\frac{\mu}{\lambda_i})$. Let δ be the probability that, starting from state 0, the jump process does not return to zero. Then $\delta = \prod_{i=0}^{\infty} p^J_{ii+1} \geq \exp(-\mu \sum_{i=0}^{\infty} \frac{1}{\lambda_i}) > 0$. Thus, X^J is transient. After the last visit to state zero, all the jumps of X^J are up one. The corresponding mean holding times of X are $\frac{1}{\lambda_i+\mu}$ which have a finite sum, so that the process X is explosive. This example illustrates the need for the assumption just after (21) in Proposition 1.6.

As for the case of discrete time, the drift conditions imply moment bounds. The proofs of the following two propositions are minor variations of the ones used for discrete time, with Lemma 1.2 used in place of Lemma 1.7, and are omitted.

Proposition 1.8 *(Moment bound–continuous time) Suppose V, f, and g are nonnegative functions on $\mathcal{S}$, and suppose $QV(i) \leq -f(i) + g(i)$ for all $i \in \mathcal{S}$. In addition, suppose X is positive recurrent, so that the means, $\overline{f} = \pi f$ and $\overline{g} = \pi g$ are well-defined. Then $\overline{f} \leq \overline{g}$.*

Corollary 1.9 *(Combined Foster-Lyapunov stability criterion and moment bound–continuous time) Suppose V, f, and g are nonnegative functions on $\mathcal{S}$ such that $QV(i) \leq -f(i) + g(i)$ for all $i \in \mathcal{S}$, and, for some $\epsilon > 0$, the set C defined by $C = \{i : f(i) < g(i) + \epsilon\}$ is finite. Suppose also that $\{i : V(i) \leq K\}$ is finite for all K. Then X is positive recurrent and $\overline{f} \leq \overline{g}$.*

Example 4.a (Random server allocation with two servers) Consider the system shown in Figure 3. Suppose that each queue i is fed by a Poisson arrival process with rate λ_i, and suppose there are two potential departure processes, D_1 and D_2, which are Poisson processes with rates m_1 and m_2, respectively. The five Poisson processes are assumed to be independent. No matter how the potential departures are allocated to the permitted queues, the following conditions are necessary for stability:

$$\lambda_1 < m_1, \qquad \lambda_3 < m_2, \qquad \text{and} \quad \lambda_1 + \lambda_2 + \lambda_3 < m_1 + m_2 \tag{24}$$

That is because server 1 is the only one that can serve queue 1, server 2 is the only one that can serve queue 3, and the sum of the potential service rates must exceed the sum of the potential arrival rates for stability. A vector $x = (x_1, x_2, x_2) \in \mathbb{Z}_+^3$ corresponds to x_i packets in queue i for each i. Let us consider random selection, so that when D_i has a jump, the queue served is chosen at random, with the probabilities determined by $u = (u_1, u_2)$. As indicated in Figure 3, a potential service by server 1 is given to queue 1 with probability u_1, and to queue 2 with probability $\overline{u_1}$.

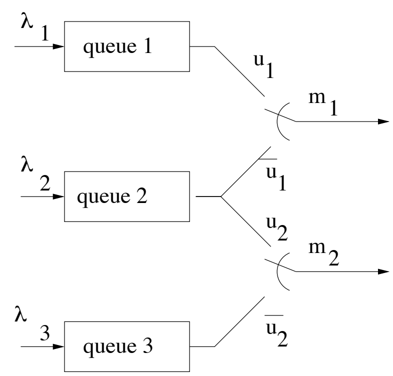

Figure 3: A system of three queues with two servers

Similarly, a potential service by server 2 is given to queue 2 with probability u_2, and to queue 3 with probability $\overline{u_2}$. The rates of potential service at the three stations are given by

$$\begin{aligned}\mu_1(u) &= u_1 m_1\\ \mu_2(u) &= \overline{u_1} m_1 + u_2 m_2\\ \mu_3(u) &= \overline{u_2} m_2.\end{aligned}$$

Let $V(x) = \frac{1}{2}(x_1^2 + x_2^2 + x_3^2)$. Using (20), we find that the drift function QV is given by

$$QV(x) = \frac{1}{2}\left(\sum_{i=1}^{3}((x_i+1)^2 - x_i^2)\lambda_i\right) + \frac{1}{2}\left(\sum_{i=1}^{3}((x_i-1)_+^2 - x_i^2)\mu_i(u)\right)$$

Now $(x_i - 1)_+^2 \leq (x_i - 1)^2$, so that

$$QV(x) \leq \left(\sum_{i=1}^{3} x_i(\lambda_i - \mu_i(u))\right) + \frac{\gamma}{2} \tag{25}$$

where γ is the total rate of events, given by $\gamma = \lambda_1+\lambda_2+\lambda_3+\mu_1(u)+\mu_2(u)+\mu_3(u)$, or equivalently, $\gamma = \lambda_1 + \lambda_2 + \lambda_3 + m_1 + m_2$. Suppose that the necessary condition (24) holds. Then there exists some $\epsilon > 0$ and choice of u so that

$$\lambda_i + \epsilon \leq \mu_i(u) \quad \text{for} \quad 1 \leq i \leq 3$$

and the largest such choice of ϵ is $\epsilon = \min\{m_1 - \lambda_1, m_2 - \lambda_3, \frac{m_1+m_2-\lambda_1-\lambda_2-\lambda_3}{3}\}$. (See excercise.) So $QV(x) \leq -\epsilon(x_1 + x_2 + x_3) + \gamma$ for all x, so Corollary 1.9 implies that X is positive recurrent and $\overline{X}_1 + \overline{X}_2 + \overline{X}_3 \leq \frac{\gamma}{2\epsilon}$.

Example 4.b (Longer first server allocation with two servers) This is a continuation of Example 4.a, concerned with the system shown in Figure 3. Examine the right hand side of (25). Rather than taking a fixed value of u, suppose that the choice of u could be specified as a function of the state x. The maximum of a function is greater than or equal to the average of the function, so that for any probability distribution u,

$$\begin{aligned}\sum_{i=1}^{3} x_i \mu_i(u) &\leq \max_{u'} \sum_i x_i \mu_i(u') \qquad (26)\\ &= \max_{u'} m_1(x_1 u_1' + x_2 \overline{u_1'}) + m_2(x_2 u_2' + x_3 \overline{u_2'})\\ &= m_1(x_1 \vee x_2) + m_2(x_2 \vee x_3)\end{aligned}$$

with equality in (26) for a given state x if and only if a longer first policy is used: each service opportunity is allocated to the longer queue connected to the server. Let Q^{LF} denote the one-step transition probability matrix when the longest first policy is used. Then (25) continues to hold for any fixed u, when Q is replaced by Q^{LF}. Therefore if the necessary condition (24) holds, ϵ can be taken as in Example 4a, and $Q^{LF}V(x) \leq -\epsilon(x_1 + x_2 + x_3) + \gamma$ for all x. So Corollary 1.9 implies that X is positive recurrent under the longer first policy, and $\overline{X}_1 + \overline{X}_2 + \overline{X}_3 \leq \frac{\gamma}{2\epsilon}$. (Note: We see that

$$Q^{LF}V(x) \leq \left(\sum_{i=1}^{3} x_i \lambda_i\right) - m_1(x_1 \vee x_2) - m_2(x_2 \vee x_3) + \frac{\gamma}{2},$$

but for obtaining a bound on $\overline{X}_1 + \overline{X}_2 + \overline{X}_3$ it was simpler to compare to the case of random service allocation.)

References

[1] S. Asmussen, *Applied Probability and Queues*, Wiley, New York, 1987.

[2] J. Edmonds, “Submodular functions, matroids and certain polyhedra, *Proc. Calgary Int. Conf. Combinatorial Structures and Applications*, Calgary, Alta, June 1969, pp. 6987.

[3] F. G. Foster, “On the stochastic matrices associated with certain queueing processes,” *Ann. Math Statist.*, Vol. 24, pp. 355-360, 1953.

[4] G.R. Grimmett and D.R. Stirzaker, *Probability and Random Processes*, 3rd Ed., Oxford University Press, Oxford and New York, 2001.

[5] J.F.C. Kingman, “Some inequalitiies for the queue GI/G/1,” *Biometrika*, vol. 49, no. 3 and 4, p. 315-324, 1962.

[6] P.R. Kumar and S.P. Meyn, “Stability of queueing networks and scheduling policies,” *IEEE Trans. Automatic Control*, vol. 40, 1995, pp. 251-260.

[7] N. McKeown, A. Mekkittikul, V. Ananthraram, and J. Walrand, “Achieving 100% throuhput in an input-queued switch,” *IEEE Transactions on Communications,* vol. 47, no. 8, 1999.

[8] S.P. Meyn and R.L. Tweedie, *Markov chains and stochastic stability,* Springer-Verlag, London, 1993.

[9] L. Tassiulas and A. Ephremides, “Stability properties of constrained queueing systems and scheduling for maximum throughput in multihop radio networks,” *IEEE Trans. Automatic Control*, vol. 37, no. 12, December 1992.

[10] L. Tassiulas and A. Ephremides, “Dynamic server allocation to parallel queues with randomly varying connectivity,” *IEEE Trans. Information Theory*, vol. 39, March 1993.

[11] L. Tassiulas, “Scheduling and performance limits of networks with constantly changing topology,” *IEEE Trans. on Info. Theory,* Vol. 43, No. 3, pp. 1067-1073, May 1997.

[12] R. L. Tweedie, The existence of moments of stationary Markov chains, *J. Appl. Prob.*, vol. 20, pp. 191 - 196, 1983.

1.4 Problems

1. Recurrence of mean zero random walks
(a) Suppose $B_1, B_2, \ldots$ is a sequence of independent, mean zero, integer valued random variables, which are bounded, i.e. $P[|B_i| \leq M] = 1$ for some M.
(a) Let $X_0 = 0$ and $X_n = B_1 + \cdots + B_n$ for $n \geq 0$. Show that X is recurrent.
(b) Suppose $Y_0 = 0$ and $Y_{n+1} = Y_n + B_n + L_n$, where $L_n = (-(Y_n + B_n))_+$. The process Y is a reflected version of X. Show that Y is recurrent.

2. Positive recurrence of reflected random walk with negative drift
Suppose $B_1, B_2, \ldots$ is a sequence of independent, integer valued random variables, each with mean $\overline{B} < 0$ and second moment $\overline{B^2} < +\infty$. Suppose $X_0 = 0$ and $X_{n+1} = X_n + B_n + L_n$, where $L_n = (-(X_n + B_n))_+$. Show that X is positive recurrent, and give an upper bound on the mean under the equilibrium distribution, $\overline{X}$. (Note, it is not assumed that the B's are bounded.)

3. Routing with two arrival streams
(a) Generalize Example 1.a to the scenario shown.

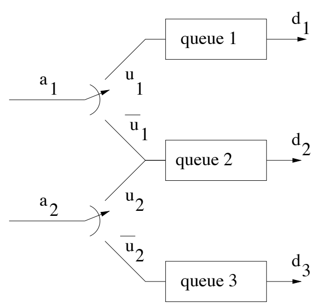

where $a_i, d_j \in (0, 1)$ for $1 \leq i \leq 2$ and $1 \leq j \leq 3$. In particular, determine conditions on a_1 and a_2 that insure there is a choice of $u = (u_1, u_2)$ which makes the system positive recurrent. Under those conditions, find an upper bound on $\overline{X_1} + \overline{X}_2 + \overline{X_3}$, and select u to mnimize the bound.
(b) Generalize Example 1.b to the scenario shown. In particular, can you find a version of route-to-shorter routing so that the bound found in part (a) still holds?

4. Allocation of service
Prove the claim in Example 4a about the largest value of ϵ.

5. Opportunistic scheduling (Tassiulas and Ephremides [10])
Suppose N queues are in parallel, and suppose the arrivals to a queue i form an independent, identically distributed sequence, with the number of arrivals in a given slot having mean $a_i > 0$ and finite second moment K_i. Let $S(t)$ for each t be a subset of $E = \{1, \ldots, N\}$ and $t \geq 0$. The random sets $S(t) : t \geq 0$ are assumed to be independent with common distribution w. The interpretation is that there is a single server, and in slot i, it can serve one packet from one of the queues in $S(t)$. For example, the queues might be in the base station of a wireless network with packets queued for N mobile users, and $S(t)$ denotes the set of mobile users that have working channels for time slot $[t, t+1)$. See the illustration:

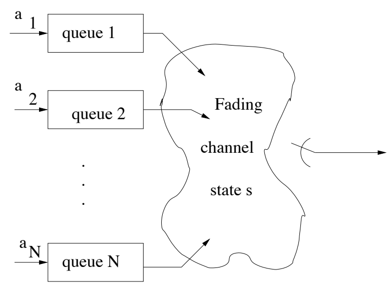

(a) Explain why the following condition is necessary for stability: For all $s \subset E$ with $s \neq \emptyset$,

$$\sum_{i \in s} a_i < \sum_{B : B \cap s \neq \emptyset} w(B) \tag{27}$$

(b) Consider u of the form $u = (u(i,s) : i \in E, s \subset E)$, with $u(i,s) \geq 0$, $u(i,s) = 0$ if $i \notin s$, and $\sum_{i \in E} u(i,s) = I_{\{s \neq \emptyset\}}$. Suppose that given $S(t) = s$, the queue that is given a potential service opportunity has probability distribution $(u(i,s) : i \in E)$. Then the probability of a potential service at queue i is given by $\mu_i(u) = \sum_s u(i,s)w(s)$ for $i \in E$. Show that under the condition (27), for some $\epsilon > 0$, u can be selected to that $a_i + \epsilon \leq \mu_i(u)$ for $i \in E$. (Hint: Apply the min-cut, max-flow theorem, given in Chapter 6 of the notes, to an appropriate graph.)
(c) Show that using the u found in part (b) that the process is positive recurrent.
(d) Suggest a dynamic scheduling method which does not require knowledge of the arrival rates or the distribution w, which yields the same bound on the mean sum of queue lengths found in part (b).

6. Routing to two queues – continuous time model
Give a continuous time analog of Examples 1.a and 1.b. In particular, suppose that the arrival process is Poisson with rate λ and the potential departure processes are Poisson with rates μ_1 and μ_2.

1.5 Solutions

1. Recurrence of mean zero random walks
(Note: The boundedness condition is imposed to make the problem easier – it is not required for the results.) (a) Let $V(x) = |x|$, the absolute value of x. Then for $x \geq M$ and any $k \geq 0$, $|x + B_k| = x + B_k$ with probability one, by the boundedness of B_k. Thus, if $x \geq M$, $PV(x) - V(x) = E[|x + B_k|] - x = E[x + B_k] - x = 0$. Similarly, if $x \leq -M$, $PV(x) - V(x) = E[|x + B_k|] - |x| = E[|x| - B_k] - |x| = -E[B_k] = 0$. Therefore, the Foster-Lyapunov criteria for recurrence are satisfied by P, V, and the set $C = \{x : |x| < M\}$, so that X is recurrent.
(b) Let $V(y) = y$. As in part (a), we see that $PV(y) - V(y) = 0$ off the finite set $C = \{y : 0 \leq y < M\}$, so that Y is recurrent.

2. Positive recurrence of reflected random walk with negative drift
Let $V(x) = \frac{1}{2}x^2$. Then

$$\begin{aligned}
PV(x) - V(x) &= E[\frac{(x + B_n + L_n)^2}{2}] - \frac{x^2}{2} \\
&\leq E[\frac{(x + B_n)^2}{2}] - \frac{x^2}{2} \\
&= x\overline{B} + \frac{\overline{B^2}}{2}
\end{aligned}$$

Therefore, the conditions of the combined Foster stability criteria and moment bound corollary apply, yielding that X is positive recurrent, and $\overline{X} \leq \frac{\overline{B^2}}{-2\overline{B}}$. (This bound is somewhat weaker than Kingman's moment bound, disussed later in the notes: $\overline{X} \leq \frac{\mathrm{Var}(B)}{-2\overline{B}}$.)

3. Routing with two arrival streams
We begin by describing the system by stochastic evolution equations. For $i \in \{0, 1\}$, let $B_i(t)$ be a Bernoulli(a_i) random variable, and let $U_i(t)$ be a Bernoulli(u_i) random variable, and for $j \in \{1, 2, 3\}$ let $D_j(t)$ be a Bernoulli(d_j) random variable. Suppose all these Bernoulli random variables are independent. Then the state process can be described by

$$X_i(t + 1) = X_i(t) + A_i(t) - D_i(t) + L_i(t) \quad \text{for } i \in \{1, 2, 3\} \tag{28}$$

where

$$\begin{aligned}
A_1(t) &= U_1(t)B_1(t) \\
A_2(t) &= (1 - U_1(t))B_1(t) + U_2(t)B_2(t) \\
A_3(t) &= (1 - U_2(t))B_2(t)
\end{aligned}$$

and $L_i(t) = (-(X_i(t) + A_i(t) - D_i(t)))_+$.

Assume that the parameters satisfy the following three conditions:

$$a_1 + a_2 < d_1 + d_2 + d_3 \tag{29}$$
$$a_1 < d_1 + d_2 \tag{30}$$
$$a_2 < d_2 + d_3 \tag{31}$$

These conditions are necessary for any routing strategy to yield positive recurrence. Condition (29) states that the total arrival rate to the network is less than the sum of service rates in the network. Condition (30) states that the total arrival rate to servers 1 and 2 is less than the sum of the service rates of those servers. Condition (31) has a similar interpretation. We will show that conditions (29)-(31) are also sufficient for positive recurrence for random routing for a particular choice of u. To that end we consider the Lyapunov-Foster criteria.

Letting $V = (x_1^2+x_2^2+x_3^2)$ and arguing as in Example 1a, and using the fact $(A_i(t)-D_i(t))^2 \leq 1$ for $i = 1, 3$ and $(A_2(t) - D_2(t))^2 \leq 4$ yields (writing $\overline{u_i}$ for $1 - u_i$):

$$\begin{aligned} PV(x) - V(x) &\leq \left(\sum_{i=1}^{3} x_i E[A_i(t) - D_i(t)|X(t) = x] \right) + 3 \\ &= x_1(a_1u_1 - d_1) + x_2(a_1\overline{u_1} + a_2u_2 - d_2) + x_3(a_2\overline{u_2} - d_3) + 3 \quad (32) \\ &\leq -(x_1 + x_2 + x_3) \min\{d_1 - a_1u_1, d_2 - a_1\overline{u_1} - a_2u_2, d_3 - a_2\overline{u_2}\} + 3 \quad (33) \end{aligned}$$

To obtain the best bound possible, we select $u = (u_1, u_2)$ to maximize the min term in (33). For any u, $\min\{d_1 - a_1u_1, d_2 - a_1\overline{u_1} - a_2u_2, d_3 - a_2\overline{u_2}\}$, is the minimum of three numbers, which is less than or equal to the average of the three numbers, is less than or equal to the average of any two of the three numbers, and is less than or equal to any one of the three numbers. Therefore, for any choice of $u \in [0, 1]^2$,

$$\min\{d_1 - a_1u_1, d_2 - a_1\overline{u_1} - a_2u_2, d_3 - a_2\overline{u_2}\} \leq \epsilon \quad (34)$$

where

$$\epsilon = \min\{\frac{d_1 + d_2 + d_3 - a_1 - a_2}{3}, \frac{d_1 + d_2 - a_1}{2}, \frac{d_2 + d_3 - a_2}{2}, d_1, d_2, d_3\}$$

Note that $\epsilon > 0$ under the necessary conditions (29)-(31). Taking $u_1^* \in [0, 1]$ as large as possible subject to $d_1 - a_1u_1^* \geq \epsilon$, and taking $\overline{u_2^*} \in [0, 1]$ as large as possible subject to $d_3 - a_2\overline{u_2^*} \geq \epsilon$, yields the choice

$$u_1^* = \frac{d_1 - \epsilon}{a_1} \wedge 1 \qquad \overline{u_2^*} = \frac{d_3 - \epsilon}{a_2} \wedge 1$$

It is not hard to check that $d_2 - a_1\overline{u_1^*} - a_2u_2^* \geq \epsilon$, so that equality holds in (34) for $u = u^*$.

This paragraph shows that the above choice of ϵ and u^* can be better understood by applying the max-flow min-cut theorem to the flow graph shown.

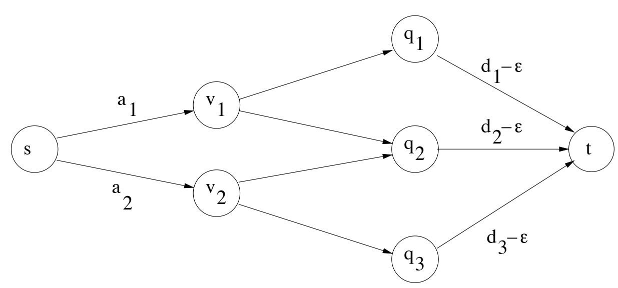

In addition to the source node s and sink node t, there are two columns of nodes in the graph. Nodes v_1 and v_2 correspond to the two arrival streams, and nodes q_1, q_2 and q_3 correspond to the three queues. There are three stages of links in the graph. The capacity of a link (s,v_i) in the first stage is a_i, the capacities of the links in the second graph are very large, and the capacity of a link (q_j, t) in the third stage is $d_j - \epsilon$. The choice of ϵ above is the largest possible so that (1) all links have nonnegative capacity, and (2) the capacity of any $s-t$ cut is greater than or equal to a_1+a_2. Thus, the min flow max cut theorem insures that there exists an $s-t$ flow with value a_1+a_2. Then, u_i^* can be taken to be the fraction of the flow into v_i that continues to q_i, for $i=1$ and $i=2$.

Under (29)-(31) and using u^*, we have $PV(x)-V(x) \leq -\epsilon(x_1+x_2+x_3)+3$. Thus, the system is positive recurrent, and $\overline{X}_1+\overline{X}_2+\overline{X}_3 \leq \frac{3}{\epsilon}$.[3]

(b) Suppose each arrival is routed to the shorter of the two queues it can go to. For simplicity, let such decisions be based on the queue lengths at the beginning of the slot. As far as simple bounds are concerned, there is no need, for example, to take into account the arrival in the top branch when making the decision on the bottom branch, although presumably this could improve performance. Also, it might make sense to break ties in favor of queues 1 and 3 over queue 2, but that is not very important either. We can denote this route-to-shorter policy by thinking of u in part (a) as a function of the state, $u=\phi^{RS}(x)$. Note that for given x, this choice of u minimizes the right hand side of (32). In particular, under the route-to-shorter policy, the right hand side of (32) is at least as small as its value for the fixed control u^*. Therefore, writing P^{RS} for the one-step transition matrix for the route to shorther policy, $P^{RS}V(x)-V(x) \leq -\epsilon(x_1+x_2+x+3)+3$, and under the necessary condtions (29)-(31), the system is positive recurrent, and $\overline{X}_1+\overline{X}_2+\overline{X}_3 \leq \frac{3}{\epsilon}$.

4. Allocation of service

Let $\epsilon>0$. The largest we can make $\mu_1(u)$ is m_1 and the largest we can make $\mu_3(u)$ is m_2. Thus, we can select u so that

$$\lambda_1+\epsilon \leq \mu_1(u) \quad \text{and} \quad \lambda_3+\epsilon \leq \mu_3(u) \tag{35}$$

if and only if $m_1 \geq \lambda_1+\epsilon$ and $m_2 \geq \lambda_3+\epsilon$. If we select u so that equality holds in (35) then the value of $\mu_2(u)$ is maximized subject to (35) , and it is given by

$$\mu_2(u) = m_1+m_2-\mu_1(u)-\mu_3(u) = m_1+m_2-\lambda_1+\lambda_3-2\epsilon$$

Therefore, in order that $\lambda_2+\epsilon \leq \mu_2(u)$ hold in addition to (35), it is necessary and sufficient that $\epsilon \leq \frac{\lambda_1+\lambda_2+\lambda_3-\mu_1-\mu_2}{3}$ and $m_1 \geq \lambda_1+\epsilon$ and $m_2 \geq \lambda_2+\epsilon$. Thus, the value of ϵ given in the example is indeed the largest possible.

5. Opportunistic scheduling (Tassiulas and Ephremides [10])

(a) The left hand side of (27) is the arrival rate to the set of queues in s, and the righthand side is the probability that some queue in s is eligible for service in a given time slot. The condition is necessary for the stability of the set of queues in s.

(b) Fix $\epsilon>0$ so that for all $s \in E$ with $s \neq \emptyset$,

$$\sum_{i\in s}(a_i+\epsilon) \leq \sum_{B:B\cap s\neq\emptyset} w(B)$$

Consider the flow graph shown.

[3] By using the fact that $\overline{X_i}=0$ if the arrival rate to queue i is zero, it can be shown that for $u=u^*$, $\overline{X}_1+\overline{X}_2+\overline{X}_3 \leq \frac{3}{\epsilon'}$, where $\epsilon' = \min\{\frac{d_1+d_2+d_3-a_1-a_2}{3}, \frac{d_1+d_2-a_1}{2}, \frac{d_2+d_3-a_2}{2}\}$.

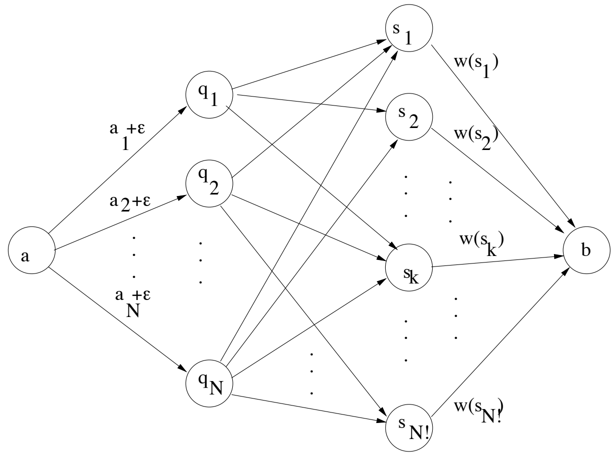

In addition to the source node a and sink node b, there are two columns of nodes in the graph. The first column of nodes corresponds to the N queues, and the second column of nodes corresponds to the 2^N subsets of E. There are three stages of links in the graph. The capacity of a link (a,q_i) in the first stage is $a_i + \epsilon$, there is a link (q_i, s_j) in the second stage if and only if $q_i \in s_j$, and each such link has capacity greater than the sum of the capacities of all the links in the first stage, and the weight of a link (s_k, t) in the third stage is $w(s_k)$.

We claim that the minimum of the capacities of all $a-b$ cuts is $v* = \sum_{i=1}^N (a_i + \epsilon)$. Here is a proof of the claim. The $a-b$ cut $(\{a\} : V - \{a\})$ (here V is the set of nodes in the flow network) has capacity v^*, so to prove the claim, it suffices to show that any other $a-b$ cut has capacity greater than or equal to v^*. Fix any $a-b$ cut $(A : B)$. Let $\widetilde{A} = A \cap \{q_1, \ldots, q_N\}$, or in words, $\widetilde{A}$ is the set of nodes in the first column of the graph (i.e. set of queues) that are in A. If $q_i \in \widetilde{A}$ and $s_j \in B$ such that (q_i, s_j) is a link in the flow graph, then the capacity of $(A : B)$ is greater than or equal to the capacity of link (q_i, s_j), which is greater than v^*, so the required inequality is proved in that case. Thus, we can suppose that A contains all the nodes s_j in the second column such that $s_j \cap \widetilde{A} \neq \emptyset$. Therefore,

$$
\begin{aligned}
C(A:B) &\geq \sum_{i \in \{q_1,\ldots,q_N\} - \widetilde{A}} (a_i + \epsilon) \;+ \sum_{s \subset E : s \cap \widetilde{A} \neq \emptyset} w(s) \\
&\geq \sum_{i \in \{q_1,\ldots,q_N\} - \widetilde{A}} (a_i + \epsilon) \;+ \sum_{i \in \widetilde{A}} (a_i + \epsilon) = v^*, \qquad (36)
\end{aligned}
$$

where the inequality in (36) follows from the choice of ϵ. The claim is proved.

Therefore there is an $a-b$ flow f which saturates all the links of the first stage of the flow graph. Let $u(i,s) = f(q_i, s)/f(s,b)$ for all i, s such that $f(s,b) > 0$. That is, $u(i,s)$ is the fraction of flow on link (s,b) which comes from link (q_i, s). For those s such that $f(s,b) = 0$, define $u(i,s)$ in some arbitrary way, respecting the requirements $u(i,s) \geq 0$, $u(i,s) = 0$ if $i \notin s$, and $\sum_{i \in E} u(i,s) = I_{\{s \neq \emptyset\}}$. Then $a_i + \epsilon = f(a, q_i) = \sum_s f(q_i, s) = \sum_s f(s,b) u(i,s) \leq \sum_s w(s) u(i,s) = \mu_i(u)$, as required.
(c) Let $V(x) = \frac{1}{2} \sum_{i \in E} x_i^2$. Let $\delta(t)$ denote the identity of the queue given a potential service at time t, with $\delta(t) = 0$ if no queue is given potential service. Then $P[\delta(t) = i | S(t) = s] = u(i,s)$. The dynamics of queue i are given by $X_i(t+1) = X_i(t) + A_i(t) - R_i(\delta(t)) + L_i(t)$, where $R_i(\delta) = I_{\{\delta = i\}}$.

Since $\sum_{i\in E}(A_i(t)-R_i(\delta_i(t)))^2 \le \sum_{i\in E}(A_i(t))^2+(R_i(\delta_i(t)))^2 \le N+\sum_{i\in E}A_i(t)^2$ we have

$$PV(x)-V(x) \le \left(\sum_{i\in E} x_i(a_i-\mu_i(u))\right)+K \tag{37}$$

$$\le -\epsilon\left(\sum_{i\in E} x_i\right)+K \tag{38}$$

where $K=\frac{N}{2}+\sum_{i=1}^N K_i$. Thus, under the necessary stability conditions we have that under the vector of scheduling probabilities u, the system is positive recurrent, and

$$\sum_{i\in E}\overline{X}_i \le \frac{K}{\epsilon} \tag{39}$$

(d) If u could be selected as a function of the state, x, then the right hand side of (37) would be minimized by taking $u(i,s)=1$ if i is the smallest index in s such that $x_i=\max_{j\in s}x_j$. This suggests using the *longest connected first* (LCF) policy, in which the longest connected queue is served in each time slot. If P^{LCF} denotes the one-step transition probability matrix for the LCF policy, then (37) holds for any u, if P is replaced by P^{LCF}. Therefore, under the necessary condition and ϵ as in part (b), (38) also holds with P replaced by P^{LCF}, and (39) holds for the LCF policy.

6. Routing to two queues – continuous time model

Suppose $\lambda<\mu_1+\mu_2$, which is a necessary condition for positive recurrence.

(a) Under this condition, we can find u so that $\mu_1>\lambda u$ and $\mu_2>\lambda\overline{u}$. If each customer is independently routed to queue 1 with probability u, and if $V(x)=(x_1^2+x_2^2)/2$, (20) becomes

$$2QV(x)=\left((x_1+1)^2-x_1^2\right)\lambda u+\left((x_2+1)^2-x_2^2\right)\lambda\overline{u}+\left((x_1-1)_+^2-x_1^2\right)\mu_1+\left((x_2-1)_+^2-x_2^2\right)\mu_2.$$

Since $(x_1-1)_+^2\le(x_i-1)^2$, it follows, with $\gamma=\lambda+\mu_1+\mu_2$, that

$$QV(x)\le -\left(x_1(\mu_1-u\lambda)+x_2(\mu_2-\overline{u}\lambda)\right)+\frac{\gamma}{2}. \tag{40}$$

Thus, the combined stability criteria and moment bound applies, yielding that the process is positive recurrent, and $\overline{X_1}(\mu_1-\lambda)+\overline{X_2}(\mu_2-\overline{u}\lambda)\le\frac{\gamma}{2}$.

In analogy to Example 1a, we let

$$\begin{aligned}\epsilon &= \max_{0\le u\le 1}\min\{\mu_1-\lambda u,\mu_2-\lambda\overline{u}\}\\ &= \min\{\mu_1,\mu_2,\frac{\mu_1+\mu_2-\lambda}{2}\}\end{aligned}$$

and the corresponding value u^* of u is given by

$$u^*=\begin{cases}0 & \text{if } \mu_1-\mu_2<-\lambda\\ \frac{1}{2}+\frac{\mu_1-\mu_2}{2\lambda} & \text{if } |\mu_1-\mu_2|\le\lambda\\ 1 & \text{if } \mu_1-\mu_2>\lambda\end{cases}$$

For the system with $u=u^*$, $\overline{X}_1+\overline{X}_2\le\frac{\gamma}{2\epsilon}$. The remark at the end of Example 1a also carries over, yielding that for splitting probability $u=u^*$, $\overline{X}_1+\overline{X}_2\le\frac{\gamma}{\mu_1+\mu_2-\lambda}$

(b) Consider now the case that when a packet arrives, it is routed to the shorter queue. To be definite, in case of a tie, the packet is routed to queue 1. Let Q^{RS} denote the transition rate matrix in case the route to short queue policy is used. Note that for any u, $\lambda(x_1 u + x_2 \overline{u}) \geq \lambda(x_1 \wedge x_2)$. Thus, (40) continues to hold if Q is replaced by Q^{RS}. In particular, if $\lambda < \mu_1 + \mu_2$, then the process X under the route-to-shorter routing is positive recurrent, and $\overline{X}_1 + \overline{X}_2 \leq \frac{\gamma}{2\epsilon}$.