

Bounding Stationary Expectations of Markov Processes

Peter W. Glynn¹ and Assaf Zeevi^{2,*}

Stanford University and Columbia University

Abstract: This paper develops a simple and systematic approach for obtaining bounds on stationary expectations of Markov processes. Given a function f which one is interested in evaluating, the main idea is to find a function g that satisfies a certain “mean drift” inequality with respect to f , which in turn leads to bounds on the stationary expectation of the latter. The approach developed in the paper is broadly applicable and can be used to bound steady-state expectations in general state space Markov chains, continuous time chains, and diffusion processes (with, or without, reflecting boundaries).

1. Introduction

Consider an irreducible non-explosive Markov jump process $X = (X(t) : t \geq 0)$ on a discrete state space S (otherwise known as a continuous-time Markov chain on S). Let $f : S \rightarrow \mathbb{R}_+$ be a cost function on S , in which $f(x)$ represents the instantaneous rate at which cost accrues when X is in state $x \in S$. (Here and in what follows, all functions are assumed to be finite valued.) Then,

$$C(t) = \int_0^t f(X(s)) ds$$

is the total cost of running X over the time horizon $[0, t]$. Computing the exact distribution of $C(t)$ (or even its expectation) is difficult. However, when X is positive recurrent, it is well known that there exists a distribution $\pi = (\pi(x) : x \in S)$ for which

$$\frac{1}{t} C(t) \rightarrow \sum_{x \in S} \pi(x) f(x) \quad \text{a.s.}$$

as $t \rightarrow \infty$; see, for example, Asmussen (2003). This justifies the approximation

$$C(t) \approx t\alpha := t \cdot \sum_{x \in S} \pi(x) f(x)$$

for t large. Of course, for this approximation to be practically useful, we need to be able to compute α or (at least) bound it.

The distribution π is the unique stationary distribution of X , so that π satisfies

$$(1) \quad \begin{aligned} & \sum_{x \in S} \pi(x) Q(x, y) = 0, \quad y \in S \\ \text{s.t. } & \sum_{x \in S} \pi(x) = 1; \quad \pi(y) \geq 0, \quad y \in S, \end{aligned}$$

*Research partially supported by NSF grant DMI-0447562

¹Management Science and Engineering, Stanford University, e-mail: glynn@stanford.edu

²Graduate School of Business, Columbia University, e-mail: assaf@gsb.columbia.edu

AMS 2000 subject classifications: Primary 68M20, 60J10, 60J15

Keywords and phrases: Markov processes, diffusions, stationary, bounds, Poisson’s equation, queueing

where $Q = (Q(x, y) : x, y \in S)$ is the rate matrix of X . If $|S|$ is finite and small, π can be computed numerically. If S is large, π can typically not be computed numerically, and in this setting one may need to be satisfied with computing bounds on α .

Assuming that π is encoded as a row vector, the linear system (1) can be re-written in matrix/vector notation as

$$(2) \quad \pi Q = 0$$

subject to π being a probability distribution on S . To obtain a bound on α , note that when $|S| < \infty$, it follows from (2) that

$$(3) \quad \pi Q g = 0$$

for any column vector g . Hence, if we can find a vector g and a constant c for which

$$(4) \quad Qg \leq -f + ce$$

(where the function $f = (f(x) : x \in S)$ is now encoded as a column vector and $e = (1, \dots, 1)^\top$ is the column vector in which all the entries are 1s), it is evident that we arrive at the upper bound:

$$(5) \quad \pi f \leq c.$$

Similarly, if we can find a \tilde{g} and \tilde{c} for which

$$(6) \quad Q\tilde{g} \geq -f + \tilde{c}e,$$

we arrive at the lower bound:

$$(7) \quad \pi f \geq \tilde{c}.$$

While the bounds (5) and (7) are trivial to derive, we are unaware of any specific literature that presents these bounds (although it seems like such bounds have appeared previously); see further comments in the literature review at the end of this section.

Our objective, in this paper, is to extend the above bounds to infinite state spaces, as well as in the direction of more general Markov processes. To offer a hint of the difficulties that can arise, suppose that X is an irreducible non-explosive birth-death process on $\mathbb{Z}_+ = \{0, 1, \dots\}$. For this class of jump processes, the so-called Poisson's equation

$$(8) \quad Qg = -k$$

has a solution for all right-hand sides k . (This solution can be computed by setting $g(0) = 0$, and then using the tri-diagonal structure of Q to recursively solve for the $g(k)$ s.). Since πk is typically non-zero, it is evident that (3) fails badly for arbitrary g , even in the setting of simple birth-death processes. (For some discussion of sufficient conditions on the functions g and k that ensure that (3) holds see, e.g., Kumar and Meyn (1996).) Note that when $|S| < \infty$, the Poisson's equation (8) is solvable for g only when $\pi k = 0$ (as can be seen by pre-multiplying both sides of (8) by the row vector π), so that the above difficulty disappears. Thus, to some degree, the complications associated with the validity of the bounds (5) and (7) have to do with issues of non-uniqueness and solvability of Poisson's equation when S is infinite, and with related potential-theoretic issues.

Since our interest is in obtaining computable bounds for $\alpha = \pi f$, our focus here is on deriving sufficient conditions under which the bounds (5) and (7) are valid. In Section 2, we explore such conditions that support the upper bound (5), and in Section 3 we develop conditions that support the derivation of the lower bound (7). Section 4 deals with linear programming formulations and its connections with the basic inequalities described above. Section 5 provides several applications of the method to queueing-related processes.

Related literature: The references mentioned below are not meant to be an exhaustive survey, but rather touch upon strands of work that are connected directly with the main theme of our paper; for further reading and connections to bounds related to the ones mentioned above, the reader is referred to Meyn and Tweedie (1993) and Borovkov (2000). In the former, the function g appearing above is referred to as a *Lyapunov function*, and similar inequalities to (4), otherwise known as drift conditions, are used as sufficient conditions to establish f -regularity of the chain; see also the “comparison theorem” in Meyn and Tweedie (Theorem 14.2.2, 1993). The primary thrust of Meyn and Tweedie (1993) is the use of such drift conditions for purposes of establishing stochastic stability and recurrence properties of Markov chains [a similar treatment can be found in Borovkov (2000), and for stochastic differential equations in Hazminskii (1979)]. We note that passing from inequality (4) to (5) bears some similarities to the analysis of Poisson’s equation in Glynn and Meyn (1996), and some connections will be made to in this paper as well.

By contrast to much of the above work, this paper is concerned with the use of the aforementioned drift conditions to develop computable bounds on various expectations of Markov processes. A particular focus is on clarifying the role that non-negativity plays in the application of such bounds. We further provide easily applied concrete hypotheses under which our bounds apply to discrete time Markov chains, Markov jump processes, and diffusion processes. In addition, we also illustrate these ideas on some queueing-related examples, and indicate how one may tighten such bounds via linear programming formulations.

A related analysis that focuses on bounding the tails of the stationary distribution can be found in Hajek (1982), Lasserre (2002), Bertsimas, Gamarnik and Tsitsikilis (2001), and Gamarnik and Zeevi (2006); see also further references in the latter two papers. One important application area that has historically driven the need for such bounds is the queueing network context. A significant number of papers have focused on deriving performance bounds for such networks. In that setting, the goal is typically to bound the steady-state queue lengths or workload [see, e.g., Bertsimas, Paschalidis and Tsitsikilis (1994), Kumar and Kumar (1994), Sigman and Yao (1997), Bertsimas, Gamarnik and Tsitsikilis (2001) and Gamarnik and Zeevi (2006), as well as references therein]. We provide some examples in this paper that are related to the derivation of such bounds.

2. The Upper Bound

Our goal is to exploit the inequality (4) so as to arrive at the bound (5). An equivalent perspective is to seek conditions on g under which

$$\pi Qg \geq 0.$$

In this case, if f is a function for which there exists c such that

$$Qg \leq -f + ce,$$

then we arrive at the bound $\pi f \leq c$. Here is our main result for Markov jump processes.

Proposition 1. *Let $X = (X(t) : t \geq 0)$ be a non-explosive Markov jump process with rate matrix Q . If $g : S \rightarrow \mathbb{R}$ is non-negative and*

$$(9) \quad \sup_{x \in S} (Qg)(x) < \infty,$$

then the inequality

$$(10) \quad \pi Qg \geq 0$$

holds for any stationary distribution π of X .

Remark 1. Note that in the presence of (9), we can write

$$Qg = -h + ce,$$

where $c = \sup\{(Qg)(x) : x \in S\}$ and h is a non-negative function. The inequality (10) then asserts that $\pi h \leq c$.

Proposition 1 is actually a special case of a result that holds for much more general Markov processes. To state this result, assume that $X = (X(t) : t \geq 0)$ is a strong Markov process taking values in a Polish space S and having càdlàg (i.e., right-continuous with left-limits) paths. We say that g belongs to the domain of the (extended) generator A of the process X and write $g \in D(A)$ if there exists a function k for which the process

$$(11) \quad M(t) = g(X(t)) + \int_0^t k(X(s)) ds$$

is a local martingale (adapted to the filtration of X) with respect to $\mathbb{P}_x(\cdot) := \mathbb{P}(\cdot | X(0) = x)$ for each $x \in S$. Furthermore, we then write $Ag = -k$, where k is any given member selected from the class of functions satisfying (11).

For the Markov processes that arise in typical applications, it is straightforward to offer conditions guaranteeing that $g \in D(A)$ and to compute explicitly Ag .

Markov jump processes: Suppose that $X = (X(t) : t \geq 0)$ is a non-explosive Markov jump process living on a discrete state S , with associated rate matrix $Q = (Q(x, y) : x, y \in S)$. Then, any function $g : S \rightarrow \mathbb{R}$ for which $\sum_{y \in S} |Q(x, y)g(y)| < \infty$ for each $x \in S$ lies in $D(A)$. Furthermore, for such a function g , $(Ag)(x) = (Qg)(x)$ for $x \in S$. To see this, let $\{K_n : n \geq 1\}$ be a sequence of subsets of S with $K_n \nearrow S$ as $n \rightarrow \infty$ and $|K_n| < \infty$ for all $n \geq 1$. We then define $T_n = \inf\{t \geq 0 : X(t) \in K_n^c\}$. Using the Kolmogorov forward equations, it follows that $M_n(t) = g(X(\min\{t, T_n\})) - \int_0^{\min\{t, T_n\}} (Qg)(X(s)) ds$ is a \mathbb{P}_x -martingale for all $x \in S$; see also further discussion in Karlin and Taylor (1981).

Stochastic differential equations (SDEs): Let $B = (B(t) : t \geq 0)$ denote standard Brownian motion in \mathbb{R}^d . Let $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times r}$ be functions that are assumed to satisfy the “usual” Lipschitz and linear growth conditions. In particular, we require the existence of constants c_1 and c_2 such that

$$(12) \quad \begin{aligned} \|\mu(x) - \mu(y)\| + \|\sigma(x) - \sigma(y)\| &\leq c_1 \|x - y\|, \\ \|\mu(x)\|^2 + \|\sigma(x)\|^2 &\leq c_2(1 + \|x\|^2) \end{aligned}$$

for all $x, y \in \mathbb{R}^d$; the (vector) coefficient function μ constitutes the *drift* of the process, and $\sigma(x)$ is known as the volatility matrix. Let $X = (X(t) : t \geq 0)$ denote

the unique \mathbb{R}^d -valued strong solution of the following stochastic differential equation (SDE)

$$(13) \quad dX(t) = \mu(X(t))dt + \sigma(X(t))dB(t) \quad ,$$

where $X(0) = x \in \mathbb{R}^d$. If g is a twice continuously differentiable, then $g \in D(A)$ and $Ag = Lg$, where L is a the second order differential operator

$$(14) \quad L := \sum_{i=1}^d \mu_i(x) \frac{\partial}{\partial x_i} + \frac{1}{2} \sum_{i,j=1}^d b_{ij}(x) \frac{\partial^2}{\partial x_i \partial x_j} \quad ,$$

where $b(x) = \sigma(x)\sigma(x)^\top$ is the *diffusion matrix*. The localizing sequence of stopping times $\{T_n : n \geq 1\}$ can be taken to be

$$T_n = \inf \left\{ t \geq 0 : \|X(t)\| \geq n, \text{ or } \int_0^t \sigma_{i,j}^2(X(s))ds \geq n \text{ for some } i = 1, \dots, d, j = 1, \dots, r \right\}.$$

see pp. 312–313 of Karatzas and Shreve (1991).

Jump diffusion processes: For simplicity, consider the one-dimensional case, i.e., $S = \mathbb{R}$. Let μ and σ be such that they satisfy (12), and consider the following process:

$$(15) \quad X(t) = \int_0^t \mu(X(s))ds + \int_0^t \sigma(X(s))dB(s) + \left(\sum_{i=1}^{N(t)} Y_i \right)$$

where $B = (B(t) : t \geq 0)$ is a standard Brownian motion, $N = (N(t) : t \geq 0)$ is a Poisson process with constant rate $\lambda > 0$, and $\{Y_i\}$ are iid random variables with common distribution function F and finite second moment. The above represents one of the more standard formulations of a jump-diffusion process, where the jump component is given by a compound Poisson process. It is assumed that N and the sequence $\{Y_i\}$ are mutually independent, as well as independent of $X(0)$ and the Brownian motion B . A sufficient condition that $g \in D(A)$ is that it be twice continuously differentiable and that $\int |g(x+y)|dF(y)$ be bounded on compact sets. For such functions g

$$(Ag)(x) := (Lg)(x) + \lambda \left(\int_{\mathbb{R}} g(x+y)dF(y) - g(x) \right) \quad ,$$

where L is given in (14). In this jump diffusion setting, the localizing sequence of stopping times $\{T_n : n \geq 1\}$ can be taken to be

$$T_n = \inf \left\{ t \geq 0 : \int_0^t (g'(X(s)))^2 \sigma^2(X(s))ds \geq n, \text{ or } \int_{\mathbb{R}} |g(X(t-) + y)|F(dy) \geq n \right\}.$$

Discrete-time Markov chains (DTMCs): Suppose that $X = (X(t) : t \geq 0)$ is an S -valued Markov chain with one-step transition kernel $P = (P(x,y) : x,y \in S)$, so that $P(x,dy) = \mathbb{P}(X_1 \in dy \mid X_0 = x)$. If $g : S \rightarrow \mathbb{R}$ is such that $\int_S P(x,dy)|g(y)| < \infty$ for each $x \in S$, then

$$M_n = g(X_n) + \sum_{j=0}^{n-1} k(X_j)$$

is a \mathbb{P}_x -local martingale for each $x \in S$, where $k(x) = g(x) - \int_S P(x, dy)g(y)$ for $x \in S$, thus, if we set $A = P - I$, then $k = -Ag$; for a localizing sequence, set $T_n = \inf\{m \geq 1 : (P|g|)(X_m) > n, \text{ or } k(X_m) > n\}$.

Recall that π is a stationary distribution for X if, given that $X(0)$ has distribution π , then $X(t)$ has distribution π for each $t \geq 0$. Let $\mathbb{E}_x[\cdot] := \mathbb{E}[\cdot \mid X(0) = x]$. Here is our main upper bound result.

Theorem 1. *Suppose that $g \in D(A)$ is a non-negative function for which*

$$(16) \quad \sup_{x \in S} (Ag)(x) < \infty.$$

Then:

i.) For each $x \in S$ and $t \geq 0$,

$$-\mathbb{E}_x \int_0^t (Ag)(X(s)) ds \leq g(x).$$

ii.) For each stationary distribution π of X ,

$$\int_S \pi(dx) (Ag)(x) \geq 0.$$

Proof. First note that by definition of g ,

$$g(X(t)) - \int_0^t (Ag)(X(s)) ds$$

is a \mathbb{P}_x -local martingale for each $x \in S$. Let $\{T_n : n = 1, 2, \dots\}$ be the localizing sequence of stopping times under \mathbb{P}_x . Then,

$$-\mathbb{E}_x \int_0^{\min\{t, T_n\}} (Ag)(X(s)) ds = g(x) - \mathbb{E}_x g(X(\min\{t, T_n\})).$$

Since g is by assumption non-negative, we obtain the inequality:

$$-\mathbb{E}_x \int_0^{\min\{t, T_n\}} (Ag)(X(s)) ds \leq g(x),$$

which holds for each $x \in S$. Put $C := \sup_{x \in S} (Ag)(x) < \infty$, then $C - Ag(x) \geq 0$. Rewriting the above inequality we have

$$\mathbb{E}_x \int_0^{\min\{t, T_n\}} (C - (Ag)(X(s))) ds \leq g(x) + C \mathbb{E}_x \min\{t, T_n\}.$$

Since, by definition, $T_n \uparrow \infty$, letting $n \rightarrow \infty$ and applying monotone convergence to each side of the inequality above we get that

$$-\mathbb{E}_x \int_0^t (Ag)(X(s)) ds \leq g(x)$$

which proves i.).

To prove ii.) we proceed as follows. Let π be any stationary distribution of X . Adding and subtracting Ct to each side of the inequality in i.) we have

$$\mathbb{E}_x \int_0^t (C - (Ag)(X(s))) ds \leq g(x) + Ct.$$

Now, dividing both sides of the inequality by t , sending $t \rightarrow \infty$ and using Fatou's lemma we have that

$$\mathbb{E}_x \liminf_{t \rightarrow \infty} \frac{1}{t} \int_0^t (C - Ag)(X(s)) ds \leq C.$$

Because $C - (Ag)(x)$ is non-negative for all $x \in S$, we may integrate the left-hand side against π . It now follows from Birkhoff's ergodic theorem that the left-hand side above is given by $C - \mathbb{E}_\pi \mathbb{E}[(Ag)(X(0)) \mid \mathcal{I}]$, where \mathcal{I} denotes the invariant sigma-field of X . Hence, we have that $\pi Ag \geq 0$. This concludes the proof. \square

Remark 2. The finite-time bound i.) is well known, see for example Meyn and Tweedie (p. 337,1993) for the discrete time result, and a version of the continuous time bound is implicit in Meyn and Tweedie (1993a). Note that Proposition 1 is a direct consequence of the above theorem.

Remark 3. A simpler way to obtain ii.) from i.) would be to require that g be π -integrable and to directly integrate both sides of the finite time bound i.) against π . Dividing by t and sending $t \rightarrow \infty$ then yields the inequality ii.). This approach has two disadvantages. Firstly, it requires an additional step from an applied standpoint, as one must now check π -integrability of g . Secondly, such a hypothesis would weaken the result, as Example 1 below shows that the functions g satisfying the hypothesis of Theorem 1 need not be π -integrable.

Example 1. Let X be the number-in-system process corresponding to the $M/M/1$ queue, so that X is a birth-death process in \mathbb{Z}_+ with birth rates $\lambda(x) = \lambda$ for $x \geq 0$ and death rates $\mu(x) = \mu$ for $x \geq 1$. If $\lambda < \mu$ then X has a unique stationary distribution $\pi(x) = (1 - \lambda/\mu)(\lambda/\mu)^x$ for $x \geq 0$ with $\rho := \lambda/\mu$. Given $\theta > 0$, the function

$$(17) \quad g(x) = \frac{\theta(\mu/\lambda)^{x+1}}{\lambda(\mu/\lambda - 1)^2} - \frac{\theta x}{\lambda(\mu/\lambda - 1)}$$

satisfies $Qg = \theta e$, and is non-negative. The function g therefore satisfies the hypothesis of Theorem 1. On the other hand, g is not π -integrable.

Remark 4. Note that Example 1 implies that the conclusions of Theorem 1 cannot be strengthened to

$$\int_S \pi(dx)(Ag)(x) = 0$$

under the hypothesis stated in Theorem 1. In other words, the inequality statement in ii.) is the best possible under the assumptions of the theorem.

Theorem 1 leads immediately to the following corollaries.

Corollary 1. *Let $X = (X(t) : t \geq 0)$ be a non-explosive Markov jump process and suppose that $f : S \rightarrow \mathbb{R}$ is non-negative. If there exists a non-negative function g and a constant c for which*

$$Qg \leq -f + ce,$$

then $\pi f \leq c$ for any stationary distribution π of X .

Corollary 2. *Let $X = (X(t) : t \geq 0)$ be a solution of the SDE (12), and suppose that $f : S \rightarrow \mathbb{R}$ is non-negative. If there exists a non-negative twice continuously differentiable function g and a constant c for which*

$$(Lg)(x) \leq -f(x) + c,$$

for $x \in S$, then

$$\int_S \pi(dx) f(x) \leq c$$

for any stationary distribution π of X .

Corollary 3. Let $X = (X(t) : t \geq 0)$ be a jump-diffusion process as in (15), and suppose that $f : S \rightarrow \mathbb{R}$ is non-negative. If there exists a non-negative twice continuously differentiable function g with $\int |g(x+y)|dF(y) < \infty$ for all $x \in S$, and a constant c for which

$$(Ag)(x) \leq -f(x) + c,$$

for $x \in S$, then

$$\int_S \pi(dx) f(x) \leq c$$

for any stationary distribution π of X .

Corollary 4. Let $X = (X_n : n \geq 0)$ be a discrete-time S -valued Markov chain with transition kernel P , and suppose $f : S \rightarrow \mathbb{R}$ is non-negative. If there exists a non-negative function $g : S \rightarrow \mathbb{R}$ and a constant c for which

$$\int_S P(x, dy) g(y) \leq g(x) - f(x) + c,$$

for $x \in S$, then

$$\int_S \pi(dx) f(x) \leq c$$

for any stationary distribution π of X .

Another important applications domain is that of diffusions with boundaries. Very similar results to Theorem 1 hold in such settings. To illustrate this point, assume that the real-valued process $X = (X(t) : t \geq 0)$ satisfies the stochastic differential equation

$$dX(t) = a(X(t))dt + b(X(t))dB(t) + d\Gamma(t),$$

where $B = (B(t) : t \geq 0)$ is a one-dimensional standard Brownian motion, and $\Gamma(\cdot)$ is the minimal non-decreasing process that increases only when X is at the origin and is such that the solution X is non-negative. If g is twice continuously differentiable, then

$$M(t) = g(X(t)) - \int_0^t (Lg)(X(s))ds - g'(0)\Gamma(t)$$

is a local martingale with respect to \mathbb{P}_x for each $x \geq 0$, where L is the differential operator defined in (14). Note that if $g'(0) \leq 0$, then

$$\tilde{M}(t) = g(X(t)) - \int_0^t (Lg)(X(s))ds$$

is a local supermartingale. The proof of Theorem 1 goes through without change in the local supermartingale setting. It follows that if f and g are non-negative functions with

$$(Lg)(x) \leq -f(x) + c$$

for $x \geq 0$ and with $g'(0) \leq 0$, then we may conclude that

$$\int_S \pi(dx) f(x) \leq c$$

for any stationary distribution π of X . This argument easily extends to other types of boundary behavior, as well as to higher-dimensional diffusions.

3. The Lower Bound

In this section, we turn to the question of when the lower bound (7) and its extensions to general Markov processes is valid. Such lower bounds would follow naturally from an inequality of the form

$$(18) \quad \int_S \pi(dx) (Ag)(x) \leq 0,$$

just as the upper bounds of Section 2 follow directly from Theorem 1.

Given our interest in obtaining bounds on the π -expectation of a non-negative function f and Section 2's discussion of the solution to Poisson's equation for such functions f , it is natural to restrict our attention to non-negative functions g for which

$$\sup_{s \in S} (Ag)(x) < \infty.$$

In view of Theorem 1, it is evident that (18) can hold only if we establish equality, namely, determining additional conditions on g ensuring that

$$(19) \quad \int_S \pi(dx) (Ag)(x) = 0.$$

In the setting of a discrete-time Markov chain, it is easily seen that the requirement that g be π -integrable suffices to guarantee (19).

Proposition 2. *Let π be a stationary distribution of the Markov chain $X = (X_n : n \geq 0)$ and suppose that $\pi|g| < \infty$. If $(Ag)(x) = \mathbb{E}_x g(X_1) - g(x)$, then $\pi Ag = 0$.*

Proof. Note that if $g = g^+ - g^-$ where $g^+ = \max\{g(x), 0\}$ and $g^- = (-g)^+$, then

$$\pi P g^+ = \pi g^+ < \infty,$$

so $\pi A g^+ = 0$. Similarly, $\pi A g^- = 0$, yielding the result. □

Corollary 5. *Suppose that f is non-negative. If there exist non-negative functions g_1 and g_2 and constants c_1 and c_2 for which*

$$\begin{aligned} (P g_1)(x) &\geq g_1(x) - f(x) + c_1, \\ (P g_2)(x) &\leq g_2(x) - g_1(x) + c_2, \end{aligned}$$

for all $x \in S$, then $\pi f \geq c_1$.

As our next example illustrates, π -integrability of g does not suffice to guarantee that $\pi Ag = 0$ in the setting of a continuous time Markov chain.

Example 2. Our counterexample is framed in the setting of a continuous time birth-death process X on $\mathbb{Z}_+ = \{0, 1, \dots\}$. Suppose that $\lambda(x) = \lambda r^x$ for $x \geq 0$ and $\mu(x) = \mu r^x$ for $x \geq 1$. Assume that $\mu > \lambda > 0$, and $r > 1$. Note that the embedded discrete-time Markov chain is a positive recurrent process (since $\lambda > \mu$). It follows that the jump process X is non-explosive. Let Q be the rate matrix of X and consider, as for Example 1, the solution g to the equation

$$Qg = \theta e$$

for some $\theta > 0$. Similar to Example 1, it is not difficult to verify that the solution g to the above is non-negative and satisfies

$$g(x) = \frac{\theta}{\lambda} \left(1 - (\mu r^2 / \lambda - r(1 + \mu/\lambda) + 1)^{-1}\right) (\lambda/(\mu - \lambda)(\mu/\lambda))^x \\ + \frac{\theta r}{\lambda} (\mu r^2 / \lambda - r(1 + \mu/\lambda) + 1)^{-1} x r^{-x}.$$

For this example, the stationary distribution π of X is given by

$$\pi(x) = (1 - \lambda/(r\mu))(\lambda/(r\mu))^x.$$

Note that if $r > 1 + \lambda/\mu$, then g is non-negative and π -integrable. However, $\pi Qg = \theta > 0$, thereby providing the required example.

Remark 5. Note that in the above example, both g and Qg are π -integrable, so evidently πQg can be positive, even if integrability of both g and Qg is imposed.

For Markov jump processes, X our next proposition provides a sufficient condition under which $\pi Qg = 0$.

Proposition 3. *Let X be a Markov jump process on discrete state space S with rate matrix Q and possessing a stationary distribution π . Suppose that g satisfies*

$$\sum_{x \in S} \pi(x) |Q(x, x)| |g(x)| < \infty.$$

Then, $\pi Qg = 0$.

Proof. Note that

$$\sum_{x, y} \pi(x) |Q(x, y)| |g(y)| = \sum_x \pi(x) \sum_{y \neq x} Q(x, y) |g(y)| + \sum_x \pi(x) |Q(x, x)| |g(x)| \\ = \sum_y |g(y)| \sum_{x \neq y} \pi(x) Q(x, y) + \sum_x \pi(x) |Q(x, x)| |g(x)| \\ = 2 \sum_y |g(y)| |Q(y, y)| \pi(y) < \infty.$$

It follows that

$$\sum_x \pi(x) (Qg)(x) = \sum_{x, y} \pi(x) Q(x, y) g(y) \\ = \sum_y g(y) \sum_x \pi(x) Q(x, y) = 0.$$

This concludes the proof. □

Corollary 6. *Suppose that f is non-negative. If there exist non-negative functions g_1 and g_2 , and constants c_1 and c_2 for which*

$$\begin{aligned} (Qg_1)(x) &\geq -f(x) + c_1, \\ (Qg_2)(x) &\leq -g_1(x)|Q(x, x)| + c_2, \end{aligned}$$

for all $x \in S$, then, $\pi f \geq c_1$.

To obtain lower bounds on SDEs and jump-diffusions, we offer the following result.

Theorem 2. *Suppose that $g \in D(A)$. Assume that the local martingale*

$$g(X(t)) - \int_0^t (Ag)(X(s))ds$$

is a martingale (adapted to the filtration of X) with respect to \mathbb{P}_x for each $x \in S$. If X has a stationary distribution π for which g is π -integrable and $\sup_{x \in S} (Ag)(x) < \infty$, then

$$\int_S \pi(dx)(Ag)(x) = 0.$$

Proof. By virtue of the martingale property,

$$\mathbb{E}_x \left[g(X(t)) - \int_0^t (Ag)(X(s))ds \right] = g(x)$$

for each $x \in S$. Note that $g(X(0))$ and $g(X(t))$ are both \mathbb{P}_π -integrable (since X is stationary under \mathbb{P}_π by definition). It follows that

$$(20) \quad \mathbb{E}_\pi \int_0^t (Ag)(X(s))ds = 0.$$

Because $\sup_{x \in S} (Ag)(x) < \infty$, either $\mathbb{E}_\pi (Ag)(X(s)) = \infty$, or $\mathbb{E}_\pi |(Ag)(X(s))| < \infty$ for each $s \geq 0$. In view of (20) we may conclude that $\mathbb{E}_\pi |(Ag)(X(s))| < \infty$ and hence (20) implies that

$$t \mathbb{E}_\pi (Ag)(X(0)) = 0,$$

which proves the result. □

The above result provides a mechanism for establishing lower bounds on stationary expectations for general Markov processes.

4. A Connection with Linear Programming

In this section, we explore connections between the bound (1) and linear programming characterizations of the stationary expectation $\alpha = \pi f$. We start by observing that when X is an irreducible finite-state discrete-time Markov chain, there always exists a solution g^* to Poisson's equation $g - Pg = f - \alpha$. Furthermore, because all functions are automatically π -integrable in this context, α can be characterized as the minimum of the following linear program (LP):

$$(21) \quad \begin{aligned} \min \quad & c \\ \text{s.t.} \quad & Pg \leq g - f + ce, \end{aligned}$$

where $e = (1, \dots, 1)^t$ is the column vector consisting entirely of 1s. A couple of observations are in order:

1. Note that if g is a solution of the inequality (21), we may always take the solution to be non-negative (without loss of generality). To see this, observe that $g + \beta e$ is then also a solution of (21) for any constant β . So, if g has a negative component, just choose $\beta = -\min\{g(x) : x \in S\}$. Hence, in finite state space, requiring g to be non-negative (as in Theorem 1) is no restriction on the class of “test functions” g .
2. It can be easily verified that the dual LP is

$$(22) \quad \begin{aligned} \max \quad & \nu f \\ \text{s.t.} \quad & \nu P = \nu \\ & \nu e = 1. \end{aligned}$$

Hence, the dual LP corresponds precisely to the standard equations that uniquely characterize the stationary distribution.

Of course, in infinite state space, a solution g to the linear inequalities system (21) may not be bounded from below, so that the non-negativity constraint on g in Theorem 1 could, in principle, limit the applicability of Theorem 1’s bound. In view of this, we offer the following result.

Theorem 3. *Suppose that f is a bounded non-negative function and that $X = (X_n : n \geq 0)$ is a uniformly ergodic Markov chain. Then, there exists a finite-valued non-negative function g and non-negative constant c that solve the linear inequality system*

$$(Pg)(x) \leq g(x) - f(x) + c, \quad x \in S.$$

Proof. We show that there exists a solution to Poisson’s equation $Pg = g - f + \alpha$ (where $\alpha = \pi f$) that is bounded below.

It is shown in Glynn and Meyn (1996) that one solution g to Poisson’s equation is

$$g^*(x) = \mathbb{E}_x \sum_{j=0}^{\tau-1} (f(X_j) - \alpha),$$

where τ is the regeneration time for the chain. In the uniformly ergodic case, the regeneration time τ has the property that

$$\mathbb{P}_x(\tau > n) = O(\gamma^n)$$

for some $\gamma \in (0, 1)$ that is uniform in x ; see Meyn and Tweedie (1993). Because f is non-negative,

$$g^*(x) \geq -\alpha \mathbb{E}_x \tau,$$

and it follows that g^* is bounded below. □

The following result suggests that for the types of functions f that arise in most applications, the non-negativity constraint on g is not a serious restriction.

Proposition 4. *Let $X = (X(t) : t \geq 0)$ be a positive recurrent irreducible Markov jump process on discrete state space S , with stationary distribution π . Suppose that $f : S \rightarrow \mathbb{R}_+$ is π -integrable and has the property that for each $c > 0$, $\{x : f(x) \leq c\}$ has finite cardinality. Then, there exists a finite-valued non-negative function g and a non-negative constant c that solve the linear inequality system:*

$$(Qg)(x) \leq -f(x) + x, \quad x \in S.$$

Proof. As in Theorem 3, we show that there exists an equality solution of the linear inequality system. In this setting, we choose z so that $f(z)$ is the minimum of f . A simple continuous-time adaptation of the reasoning of Glynn and Meyn (1996) establishes that one solution g to Poisson's equation is

$$g^*(x) = \mathbb{E}_x \int_0^{\tau(x)} (f(X(s)) - \alpha) ds,$$

where $\alpha = \pi f$. Let $K = \{x : f(x) \leq \alpha\}$, and note that if $T_K = \inf\{t \geq 0 : X(t) \in K\}$, it is evident that

$$g^*(x) = \mathbb{E}_x \int_0^{T_K} (f(X(s)) - \alpha) ds + \mathbb{E}_x g^*(X(T_K)).$$

Since $f(x) \geq \alpha$ for $x \in K^c$, it follows that

$$(23) \quad \mathbb{E}_x \int_0^{T_K} (f(X(s)) - \alpha) ds \geq 0$$

for all $x \in S$ (because (23) holds trivially for $x \in K$). The set K has finite cardinality so

$$\beta := \inf\{g^*(x) : x \in K\} > -\infty.$$

So, g^* is bounded below over S by β , and hence there exists a non-negative solution to Poisson's equation. \square

We conclude this section by showing how LP methods can be used to tighten the bound on $\alpha = \pi f$ relative to the constant c as determined by (6), where f is non-negative. In particular, suppose that one has found a non-negative Lyapunov function \tilde{g} satisfying (5) on the complement of some subset K . In order to obtain a finite-dimensional LP, suppose that K is a finite set for which

$$\sum_{y \in K^c} P(x, y) \tilde{g}(y)$$

can be computed for each $x \in K$. To tighten the bound on α relative to (6), consider the LP:

$$(24) \quad \begin{aligned} & \min c \\ & \text{s.t. } \sum_{y \in K} P(x, y) g(y) \leq - \sum_{y \in K^c} P(x, y) \tilde{g}(y) - f(x) + c, \quad \text{for all } x \in K \\ & 0 \leq g(x) \leq \tilde{g}(x) \quad \text{for all } x \in K. \end{aligned}$$

By setting $\hat{g} = g$ in K and $\hat{g} = \tilde{g}$ in K^c , we see that (\hat{g}, c^*) satisfies the hypotheses of Theorem 1, where c^* is the minimum of the LP (24). Hence, $\alpha \leq c^*$. Thus, the bounding method developed in this paper can be used in conjunction with LP ideas to create a numerical scheme for computing tight bounds on $\alpha = \pi f$.

5. Applications

This section presents various applications of the above results; the applications are grouped into two categories. The first set deals with application in discrete time, the focus being on reflected random walks. The second set of examples deals with applications in continuous-time: examples include analysis of a Markov jump process and a diffusion process with reflecting boundaries, both motivated by applications in queueing theory.

5.1. The single-server queue and random-walk processes

The single-server queue. Consider a queue that is fed by an arrival stream of jobs with i.i.d. *processing requirements* $V = (V_n : n \geq 0)$, and i.i.d. *inter-arrival times* $U = (U_n : n \geq 1)$ that are also independent of the processing requirements. We assume that $\mathbb{E}V_0 \leq \mathbb{E}U_1$, guaranteeing the stability of the system (i.e., $\rho := \mathbb{E}V_0/\mathbb{E}U_1 < 1$). We also assume that at time $t = 0$ the first job arrives at the queue and finds the system empty. Let $X = (X_n : n \geq 0)$ denote the *waiting time* process, where X_n is the time that the n th job spends in the system before receiving service. Taking $Z_n := V_n - U_{n+1}$, we can express the dynamics of the waiting time via the recursion

$$X_{n+1} = [X_n + Z_{n+1}]_+$$

where $[x]_+ := \max\{x, 0\}$. By construction, X is a discrete-time Markov chain taking values in $S = \mathbb{R}_+$. Given the negative drift condition $\mathbb{E}Z < 0$, basic stability theory for the $G/G/1$ queue ensures the existence (and uniqueness) of a stationary distribution π . Suppose we are interested in bounds on the first moment of this distribution, that is, suppose that $f(x) = x$. Put $g(x) = x^2$. Then, for any $x \in S$ we have

$$\begin{aligned} \mathbb{E}_x g(X_1) &= \mathbb{E}([x + Z_1]_+)^2 \\ &= \mathbb{E}(x + Z_1)^2 - \mathbb{E}((x + Z_1)^2; x + Z_1 < 0) \quad . \end{aligned}$$

So,

$$\begin{aligned} \mathbb{E}_x g(X_1) &\leq \mathbb{E}(x + Z_1)^2 \\ &= x^2 + 2x\mathbb{E}Z_1 + \mathbb{E}Z_1^2 \quad . \end{aligned}$$

Since $(Ag)(x) \leq \mathbb{E}Z_1^2$, then if Z has a finite second moment we can apply Theorem 1 (and Corollary 4), yielding the upper bound

$$\mathbb{E}_\pi X_1 \leq \frac{\mathbb{E}Z_1^2}{2|\mathbb{E}Z_1|} \quad ,$$

which is nothing but Kingman's bound; see Kingman (1962). If one considers the performance of this bound in heavy-traffic, i.e., along a sequence of systems $n = 1, 2, \dots$ in which $\rho^n \rightarrow 1$ in such a way that $\sqrt{n}(1 - \rho^n) \rightarrow \mu$ for some $\mu > 0$, then, denoting the corresponding sequence of waiting times $\{X^n : n \geq 1\}$ we have that $n^{-1/2}\mathbb{E}_\pi X_n \rightarrow \sigma^2/(2\mu)$, where $\sigma^2 = \mathbb{E}Z_1^2$. Hence the above inequality holds with equality in the aforementioned limiting sense.

To derive a lower bound on the mean, note that

$$\begin{aligned} -\mathbb{E}[(x + Z_1)^2; x + Z_1 < 0] &= \mathbb{E}\left[\int_{x+Z_1}^0 2udu; x + Z_1 < 0\right] \\ &= 2\mathbb{E}\left[\int_{-\infty}^0 u\mathbb{I}\{x + Z_1 \leq u, x + Z_1 \leq 0\}du\right] \\ &= 2\int_{-\infty}^0 u\mathbb{P}(Z_1 \leq u - x)du \\ &\geq 2\int_{-\infty}^0 u\mathbb{P}(Z_1 \leq u)du \\ &= -\mathbb{E}[Z_1^2; Z_1 < 0] \quad . \end{aligned}$$

If $\mathbb{E}|Z_1|^3 < \infty$ then $\pi g < \infty$ and hence Theorem 2 yields

$$\mathbb{E}_\pi X_1 \geq \frac{\mathbb{E} [Z_1^2; Z_1 \geq 0]}{2|\mathbb{E}Z_1|} .$$

Note that the upper bound only requires $\mathbb{E}Z_1^2 < \infty$.

5.2. Applications in continuous time

We consider two applications. The first derives bounds on mean queue lengths in a multi-class single server queue operating under the *longest queue first* (LQF) scheduling policy. We then derive bounds on moments of semi-martingale reflecting Brownian motion in the orthant.

Performance bounds for scheduling control in a single-server multi-class queue. Customer arrival are modeled as m mutually independent Poisson processes with rates $\lambda_1, \dots, \lambda_m$. The processing requirements of customers in each class follow an exponential distribution with mean $1/\mu_i$, $i = 1, \dots, d$ and are independent of each other and of the arrival processes. There is a single server which can serve customers at unit rate. Upon arrival, customers either get served immediately or are put into infinite capacity buffers, according to their class. Upon completion of service, a customer leaves the system. In any given class, at most one customer can be serviced and the sequencing within a class is according to a First-In-First-Out (FIFO) discipline. Customers that are not in service are said to be in the queue.

We will assume in what follows that

$$\rho := \sum_{i=1}^m \rho_i < 1 \quad ,$$

where $\rho_i := \lambda_i/\mu_i$ for $i = 1, \dots, m$. (The quantity ρ is referred to as the *traffic intensity* in the system.) It is well known that under the above condition every Markovian work-conserving policy is stable, in the sense that the associated CTMC is positive recurrent. (By work-conserving we mean that the server does not idle whenever there is work to be done.) Conversely, if $\rho > 1$ then any scheduling policy is unstable (i.e., there is no steady-state).

Denote the queue-length vector at time $t \geq 0$ by $X(t) = (X_1(t), \dots, X_m(t))$, and let $X = (X(t) : t \geq 0)$ denote the queue-length process. To illustrate the application of our Lyapunov inequality, we consider a simple state-dependent scheduling policy known as “serve the longest queue first,” denoted LQF for brevity. As the name suggests, this policy assigns the server to serve the class in which the queue length is the longest, and if no customers are in the system the server idles. We allow for service to be preempted if at any time instant the queue in one of the classes that is not being served increases beyond the length of the queue in the currently served class.

To formalize the verbal description of the scheduling policy, define a mapping $\delta : \mathbb{Z}_+^d \rightarrow \{0, \dots, m\}$, such that for any fixed vector of queue lengths $x \in \mathbb{Z}_+^d$, we have $a(x) \in \{e_0, e_1, \dots, e_d\}$ where e_i is the i th unit vector in \mathbb{R}^m , and e_0 is an m -dimensional zero vector. The *action* $a(x)$ specifies what customer class receives service when the system is in the given state x . Let $a_i(x)$ be the i th component of $a(x)$ for any state x . For the LQF discipline we have that

$$a(x) = \begin{cases} e_i & \text{if } x_i > \max\{x_j : j \neq i\} \\ e_0 & \text{otherwise} \quad , \end{cases}$$

and ties are broken arbitrarily by giving priority to the class with larger index. Let us denote by $i^*(x)$ the class that is granted priority under this scheduling rule. With this notation, the infinitesimal generator of the controlled CTMC is

$$A(x, y) = \begin{cases} \lambda_i & \text{if } y = x + e_i \\ \mu_i a_i(x) & \text{if } y = x - e_i \end{cases}$$

for any two states $x, y \in \mathbb{Z}_+^m$, such that $x \neq y$ (where this vector inequality is interpreted to hold if the two vectors differ at least in one coordinate), for $i = 1, \dots, m$. The diagonal entries in this matrix are defined by $A(x, x) = -\sum_{y \neq x} A(x, y)$.

Our objective is to obtain upper bounds on the steady-state queue lengths under the aforementioned LQF policy. Given our chosen scheduling rule, we are particularly interested in the behavior of the longest queue. Using previously established notation, put $f(x) = \|x\|_\infty := \max\{x_1, \dots, x_m\}$, and take the test function g to be

$$g(x) = \sum_{i=1}^m \frac{x_i^2}{\mu_i} + \sum_{i=1}^m \frac{x_i}{\mu_i} .$$

Using the definition of the infinitesimal generator, straightforward algebra yields that

$$(Ag)(x) = 2 \sum_{i=1}^m \rho_i x_i - 2x_{i^*} + 2\rho,$$

where $\rho_i = \lambda_i/\mu_i$, $\rho = \sum_i \rho_i$ and i^* is the index of the largest coordinate of x . Hence, we have that

$$(Ag)(x) \leq -2f(x)(1 - \rho) + 2\rho,$$

which serves as the basic inequality for the purposes of bounding the maximal queue length. In particular, by Theorem 1 and Corollary 1 we have

$$\mathbb{E}_\pi \|X(t)\|_\infty \leq \frac{\rho}{(1 - \rho)} .$$

A simple manipulation of the above bound gives us the following bound on the total workload in the system in steady-state

$$\mathbb{E}_\pi \left[\sum_{i=1}^m \frac{X_i(t)}{\mu_i} \right] \leq \frac{\rho^2}{\lambda_{\min}(1 - \rho)} ,$$

where $\lambda_{\min} = \min\{\lambda_1, \dots, \lambda_m\}$. The work reported on in Bertsimas et al. (2002) can be used to contrast this with a lower bound that holds for all stable Markovian policies, and is derived by other methods. In particular, Theorem 2 in Bertsimas et al. (2002) asserts that

$$\mathbb{E}_\pi \left[\sum_{i=1}^m \frac{1}{\mu_i} X_i(t) \right] \geq \frac{\rho^2}{4\lambda_{\max}(1 - \rho)} ,$$

where $\lambda_{\max} = \max\{\lambda_i : i = 1, \dots, d\}$. Hence our argument recovers the correct order of this lower bound. In particular, this implies that the performance of LQF scheduling is within a constant factor of the best possible scheduling rule, for all problem instances in which the ratio $\lambda_{\min}/\lambda_{\max}$ is held constant.

Reflected Brownian motion. A class of diffusion processes that play a central role in queueing theory are of the type most often referred to as *reflected Brownian*

motions (or RBMs); see, e.g., Harrison and Reiman (1981), Harrison and Williams (1987), Dupuis and Williams (1995) and references therein. (There are many more recent references on the topic but the latter are the most relevant to the examples presented below.) Specific instances of these processes have been shown to arise as diffusion limits of certain queueing networks that operate under so-called heavy traffic conditions. We next proceed with three concrete examples of such RBMs and illustrate how our basic inequalities can be used to obtain bounds on the tail of their stationary distribution.

Example 1: One dimensional RBM. The simplest instance of RBMs arises as a diffusion limit of the single server queue in heavy-traffic. The process can be defined as the unique strong solution of the following stochastic differential equation, which is a particular instance of the class of stochastic differential equations discussed in Section 3.2:

$$\begin{aligned} dX(t) &= -\mu dt + \sigma dB(t) + d\Gamma(t) \\ X(0) &= x_0, \end{aligned}$$

where $\mu, \sigma >$ are positive constants, $x_0 \geq 0$, and $\Gamma = (\Gamma(t) : t \geq 0)$ is the “pushing process” that keeps $X = (X(t) : t \geq 0)$ from going negative. It is well known that the stationary distribution of this process is exponential with mean $\sigma^2/(2\mu)$, that is, when $X(0)$ is drawn from this distribution the process X is stationary. Establishing this result is not difficult, but does require some modest amount of work and familiarity with properties of Brownian motion [see, e.g., Karatzas and Shreve (1991)]. We next illustrate how our basic inequality can be used to obtain rough bounds on the tail of the stationary distribution. Let $g(x) = \exp\{ax\} - ax$ for some positive constant a to be specified shortly. (It is possible to allow for two different constants to parameterize g , but this does not improve upon the bounds derived below.) With this definition we have $g'(0) = 0$, and

$$(Lg)(x) = -\mu ag(x) + \mu a + \frac{\sigma^2}{2} a^2 g(x).$$

Fix $\epsilon > 0$ and set $a := 2\mu(1 - \epsilon)/\sigma^2$. From this we get that $(Lg)(x) \leq -\epsilon(1 - \epsilon)\mu ag(x) + \mu a$. Hence, setting $f(x) = \exp\{ax\}$ and $c = (\epsilon(1 - \epsilon))^{-1}$, with the above choice of constant a we have by Corollary 2 (in particular, the discussion following Corollary 4) that $\pi f \leq c$. Now, using Markov’s inequality we have that $\mathbb{P}_\pi(X(0) \geq x) \leq (\epsilon(1 - \epsilon))^{-1} \exp\{-2\mu(1 - \epsilon)/\sigma^2 x\}$, for any $\epsilon > 0$.

Example 2: RBM in the orthant. Let $S = \mathbb{R}_+^d$ (the positive d -dimensional orthant). Let μ be a constant vector in \mathbb{R}^d , σ a $d \times d$ non-degenerate covariance matrix (symmetric and strictly positive definite), and R a $d \times d$ matrix. For each $x \in S$, a semimartingale reflecting Brownian motion (abbreviated as SRBM) associated with the data (S, μ, σ, R, x) is an \mathcal{F}_t -adapted, d -dimensional process $X = (X(t) : t \geq 0)$ defined on some filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, \mathbb{P})$ such that:

- (i) $X = W + R\Gamma$, \mathbb{P}_x -a.s.,
- (ii) \mathbb{P}_x -a.s., X has continuous paths and $X(t) \in S$ for all $t \geq 0$,
- (iii) W is a d -dimensional Brownian motion with drift vector μ , covariance matrix σ and $W(0) = x$. In addition, $M(t) = W(t) - \mu t$ is an \mathcal{F}_t -martingale.
- (iv) Γ is an \mathcal{F}_t -adapted d -dimensional process such that under \mathbb{P} it satisfies for each $j = 1, \dots, d$:
 - a.) $\Gamma_j(0) = 0$
 - b.) $(\Gamma_j(t) : t \geq 0)$ is continuous and non-decreasing

c.) $\Gamma_j(t)$ can increase only when X hits the face $F_j = \{x \in \mathbb{R}_+^d : x_j = 0\}$.

Loosely speaking, SRBM behaves like Brownian motion in the interior of the orthant, and is confined to the orthant by instantaneous “reflection” at the boundary faces, where the direction of reflection is dictated by the matrix R .

The most general condition currently known to ensure existence and uniqueness (in law) of SRBM in the orthant is that the matrix R is *completely S*. (That is, in the construction of the SRBM satisfying the properties detailed above is not done via a mapping from the initial condition and the Brownian motion, but rather as a weak solution.) The completely S condition is in fact necessary and sufficient; see Taylor and Williams (1993). [This property requires that for every principal sub-matrix of \tilde{R} of R there exists a vector v with strictly positive entries such that $\tilde{R}v$ is strictly positive.] For this class of SRBMs it is more challenging to characterize the existence of a stationary distribution. Dupuis and Williams (1995) prove that a sufficient condition for the existence and uniqueness of a stationary distribution is that all solutions of an associated deterministic Skorohod problem are attracted to the origin in finite time. (Their proof relies on a construction of a somewhat complicated piecewise linear Lyapunov function and uses the martingale structure of SRBM.) We next illustrate how a variation on that idea, using the basic inequalities developed in this paper, can be used to establish integrability of moments of an SRBM.

For the SRBM X , we require the following conditions to hold: (i) R is symmetric and positive definite; and (ii) $-\gamma := R^{-1}\mu < 0$ componentwise. Condition (ii) is necessary for the existence of a stationary distribution (see Dai and Harrison (2008)). As for the symmetry assumption, this is imposed primarily to facilitate the explicit construction of a simple test function g . Fix $a > 0$ and let $g(x) = \exp\{a\sqrt{1 + x^\top R^{-1}x}\}$. Straightforward algebra yields that

$$\begin{aligned}
 \nabla g \cdot \mu &= -a \frac{x \cdot \gamma}{\sqrt{1 + x^\top R^{-1}x}} g(x) \\
 \sum_{i,j=1,\dots,d} \sigma_{ij} \frac{\partial g^2(x)}{\partial x_i \partial x_j} &\leq a^2 c_1 g(x) + ac_2 (1 + x^\top R^{-1}x)^{-1/2} g(x) \\
 (\nabla g \cdot R)_i &= \frac{ax_i}{\sqrt{1 + x^\top R^{-1}x}} g(x),
 \end{aligned}
 \tag{25}$$

where c_1, c_2 are finite constants that depend only on the matrices R and σ , and can be computed explicitly in a straightforward manner (we omit such calculations for space considerations). Examining the third equality in (25), we may conclude that $\int \nabla g(X(t))^\top R d\Gamma(t) = 0$ for all t since $\int X(t) \cdot d\Gamma(t) = 0$ by definition of the SRBM. Examining the first and second inequalities in (25), it follows that for a suitable choice of $r > 0$, depending on γ and c_2 , we can ensure that $\nabla g \cdot \mu \leq -2ac_2g(x)$ for $x \notin \{x : 0 \leq x_i \leq r\}$. It then follows that by taking $a < c_2/(2c_1)$ we have that

$$(Lg)(x) \leq -(c_2^2/(2c_1))g(x) + c$$

where $c := \max\{|(Lg)(x)| : 0 \leq x_j \leq r, j = 1, \dots, d\}$. Put $K_n := \{x \in S : x_i \leq n, \text{ for all } i = 1, \dots, n\}$. Let $T_n = \inf\{t \geq 0 : X_i(t) \notin K_n\}$. By continuity of the paths of SRBM, we have that $T_n \rightarrow \infty$ a.s., as $n \rightarrow \infty$. We can now apply Itô’s differential rule to X , “localize” the martingale term using T_n and hence apply the same logic used in the proof of Theorem 1 and Corollary 2 (and the sketch provided for diffusion with reflecting boundaries) and arrive at the conclusion that $\pi f \leq c$.

This yields

$$\mathbb{E}_\pi \left[\exp \left\{ a \sqrt{1 + X(0)^\top R^{-1} X(0)} \right\} \right] \leq 2cc_1/c_2^2 .$$

Thus, for suitably small a we have exponential moments for the particular class of SRBMs satisfying assumptions (i) and (ii) above. Corresponding bounds on the tail of the stationary distribution follow immediately by using Markov's inequality. We should note that more precise characterization of the tail of SRBMs (without the need for condition (ii)) was recently derived by Budhiraja and Lee (2007).

Acknowledgements

We wish to thank an anonymous referee for his/her comments on an earlier version of the paper which led to an improved exposition, and highlighting the importance of Theorem 1. Many thanks are due to Stewart Ethier for giving us the opportunity to contribute this paper honoring Tom Kurtz on the occasion of his 65th birthday.

References

- [1] ASMUSSEN, S. (2003). *Applied Probability and Queues*. Springer, NY.
- [2] BERTSIMAS, D., GAMARNIK, D. and TSITSIKLIS, J. (2001). Performance of multiclass Markovian queueing networks via piecewise linear Lyapunov functions. *Ann. Appl. Probab.* **11** 1384–1428.
- [3] BERTSIMAS, D., PASCHALIDIS, I. and TSITSIKLIS, J. (1994). Optimization of multiclass queueing networks: Polyhedral and nonlinear characterizations of achievable performance. *Ann. Appl. Probab.* **4** 43–75.
- [4] BOROVKOV, A. A. (2000). *Ergodicity and Stability of Stochastic Processes*. John Wiley & Sons, New York.
- [5] BUDHIRAJA, A. and LEE, C. (2007). Long time asymptotics for constrained diffusions in polyhedral domains. *Stochastic Processes and Their Applications* **117** 1014–1036.
- [6] DAI, J. and HARRISON, J. M. (2008). Reflecting Brownian motion in the orthant: An illuminating example of instability. *Math. of Oper. Res.*, to appear.
- [7] DUPUIS, P. and WILLIAMS, R. J. (1994). Lyapunov functions for semimartingale reflecting Brownian motions. *Ann. Probab.* **22** 680–702.
- [8] FAYOLLE, G. (1989). On random walks arising in queueing systems: Ergodicity and transience via quadratic forms as Lyapunov functions. Part I. *Queueing Systems* **5** 167–184.
- [9] GAMARNIK, D. and ZEEVI, A. (2006). Validity of heavy-traffic steady-state approximations in generalized Jackson networks. *Ann. Appl. Probab.* **16** 56–90.
- [10] GLYNN, P. W. and MEYN, S. P. (1996). A Liapounov bound for solutions of the Poisson equation. *Ann. Probab.* **24** 916–931.
- [11] HAJEK, B. (1982). Hitting-time and occupation-time bounds implied by drift analysis with applications. *Ann. Appl. Probab.* **14** 502–525.
- [12] HAS'MINSKII, R. Z. (1980). *Stochastic Stability of Differential Equations*. Monographs and Textbooks on Mechanics of Solids and Fluids: Mechanics and Analysis **7**. Germantown, Md.
- [13] KARLIN, S. and TAYLOR, H. (1981). *A Second Course in Stochastic Processes*. Academic Press, San Diego.

- [14] KARATZAS, I. and SHREVE, S. E. (1991). *Brownian Motion and Stochastic Calculus*. Springer.
- [15] KINGMAN, J. F. C. (1962). Some inequalities for the queue $GI/G/1$. *Biometrika* **49** 315–324.
- [16] KUMAR, P. R. and MEYN, S. P. (1996). Duality and linear programs for stability and performance analysis of queueing networks and scheduling policies. *IEEE Trans. on Automatic Control* **41** 4–17.
- [17] KUMAR, S. and KUMAR, P. R. (1994). Performance bounds for queueing networks and scheduling policies. *IEEE Trans. on Automatic Control* **39** 1600–1611.
- [18] LASSERRE, J. (2002). Bounds on measures satisfying moment conditions. *Ann. Appl. Probab.* **12** 1114–1137.
- [19] LIONS, P.-L. and SZNITMAN, A.-S. (1984). Stochastic differential equations with reflecting boundary conditions. *Comm. Pure Appl. Math.* **37** 511–537.
- [20] MEYN, S. P. and TWEEDIE, R. L. (1993). *Markov Chains and Stochastic Stability*. Communications and Control Engineering Series. Springer-Verlag, London, 1993.
- [21] MEYN, S. P. and TWEEDIE, R. L. (1993). Stability of Markovian processes III: Foster–Lyapunov criteria for continuous time processes. *Adv. Appl. Probab.* **25** 518–548.
- [22] SIGMAN, K. and YAO, D. (1993). Finite moments for inventory processes. *Ann. Appl. Probab.* **3** 765–778.
- [23] TAYLOR, L. M. and WILLIAMS, R. J. (1993). Existence and uniqueness of semi-martingale reflecting Brownian motion in the orthant. *Probab. Theory Related Fields* **96** 283–317.
- [24] TWEEDIE, R. (1983). The existence of moments for stationary Markov chains. *J. Appl. Probab.* **20** 191–196.