A Crash Course on

Markov Chains and Stochastic Stability

Sean P. Meyn

Coordinated Science Laboratory
and the Dept. of Electrical and Computer Engineering
The University of Illinois
http://black.csl.uiuc.edu:80/~meyn

Contents

1 Introduction

These lecture notes are a self contained treatment of the solution of equilibria equations and ergodicity for "ψ-irreducible Markov chains". For a bit more depth, and in particular the treatment of general state spaces, see [12] (available at `black.csl.uiuc.edu/~meyn`), or see [11].

The Markov chains that we consider evolve on a countable state space, denoted X. The chain itself is denoted $\{X(t) : t \in \mathbb{Z}_+\}$, with transition law denoted P:

$$\mathsf{P}\{X(t+1) = x \mid X(t) = y\} = P(y, x), \qquad x, y \in \mathsf{X}.$$

2 Equilibria equations

There are two linear equations that we would like to solve:

(i) *Invariance equation*:

$$\pi P = \pi, \tag{1}$$

where π is seen as a row vector, P as a matrix. Hence

$$\sum_{x \in \mathsf{X}} \pi(x) P(x,y) = \pi(y), \quad y \in \mathsf{X}.$$

We hope that the solution is positive, $\pi(x) \geq 0$ for $x \in \mathsf{X}$, and finite so that $\pi(\mathsf{X}) = \sum \pi(x) < \infty$.

(ii) *Poisson's equation*: Given a π-integrable function $f : \mathsf{X} \to \mathbb{R}$,

$$P\hat{f} = \hat{f} - f + \eta, \tag{2}$$

where we write

$$\eta := \pi(f) := \sum_{x \in \mathsf{X}} \pi(x) f(x).$$

Here $\hat{f}$ is seen as a column vector

$$\sum_y P(x,y)\hat{f}(y) = \hat{f}(x) - f(x) + \eta \qquad x \in \mathsf{X}.$$

Letting $\Delta = P - I$ denote the difference operator, the two equations can be written,

$$\pi\Delta = 0 \qquad \text{and} \qquad \Delta\hat{f} = -f + \eta.$$

Solving either equation amounts to a form of inversion, but there are two difficulties. One is that the matrices to be inverted are not finite dimensional. The other is that these matrices are not invertable! For example, to solve Poisson's equation it appears that we must invert Δ. However, if there exists an invariant probablity π, then π is a member of the left null space of Δ which means that Δ cannot be invertible.

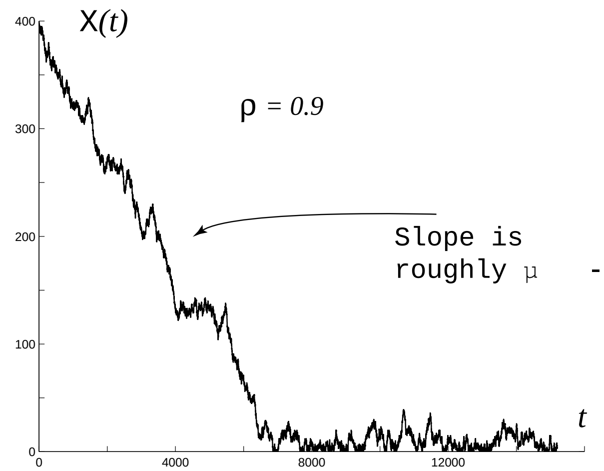

Figure 1: Sample path of a stable M/M/1 queue ($\rho < 1$).

2.1 Who cares?

The existence of π is equivalent to a form of *stability*, so that the sample path behavior looks like that shown in Figure 1, and this leads to ergodic theorems such as,

$$
\begin{aligned}
\frac{1}{N}\sum_{t=0}^{N-1} f(X(t)) &\to \eta, \qquad N \to \infty \\
\mathsf{E}[f(X(t))] &\to \eta, \qquad t \to \infty.
\end{aligned}
\tag{3}
$$

The existence of $\hat{f}$ leads to finer results:

(1) The solution to Poisson's equation is central to optimal control where f is a one-step cost function, and $\hat{f}$ is called the 'relative value function' [5, 6, 1, 16, 12].

(2) Approximate solutions to Poisson's equation lead to direct performance bounds (estimates of η) [8, 15, 7, 3, 4].

(3) The solution to Poisson's equation allows us to construct the useful martingale:

$$
M(t) = \Big(\sum_{i=0}^{t-1} f(X(i))\Big) + \hat{f}(X(t)) - \eta t
$$

This leads to the central limit theorem,

$$\frac{1}{\sqrt{N}}\sum_{0}^{N-1}(f(X(t))-\eta)\xrightarrow{\mathrm{d}}N(0,\gamma^2)$$

where

$$\gamma^2 = 2\pi(\hat{f}(f-\eta)) - \pi((f-\eta)^2).$$

Hence Poisson's equation provides tools for addressing performance of simulators [3].

Example $M/M/1$ queue When the arrival stream is renewal, and the service times are i.i.d., then the waiting time for a simple queue can be modeled as a Markov chain with state space $\mathsf{X} = \mathbb{R}_+$. The dynamics take the form of a one dimensional linear state space model, where the state space is constrained to the positive half line. The queue length process is itself a Markov process in the special case where the service times and interarrival times are exponentially distributed. By applying *uniformization* (i.e. sampling the process appropriately - see [10]), the queue length process $\mathbf{X}$ obeys the recursion

$$X(t+1) = X(t) + (1-I_{t+1})\pi(x_t) + I_{t+1}, \qquad t \in \mathbb{Z}_+,$$

where $\mathbf{I}$ is a Bernouli, i.i.d. random process: $\lambda = \mathsf{P}(I(t)=1)$ is the arrival rate, and $\mu = \mathsf{P}(I(t)=-1)$ is the service rate. Time has been normalized so that $\lambda+\mu=1$.

When $\rho = \lambda/\mu < 1$ then the invariant probability is given by

$$\pi(i) = (1-\rho)\rho^i, \qquad i \in \mathsf{X}.$$

A sample path of the process is shown in Figure 1 for this particulr model with $\rho = 0.9$. When f is the identity function, ($f(i) = i$ for $i \in \mathsf{X}$), then the solution to Poisson's equation is given by

$$\hat{f}(i) = \frac{1}{2}\frac{i^2+i}{(1-\rho)}.$$

We see that $\hat{f}$ is *quadratic*, which one could have guessed from looking at the equations: On iterating the formula $P\hat{f} = \hat{f} - f + \eta$,

$$P^2\hat{f} = \hat{f} - f - Pf + 2\eta \quad\Longrightarrow\quad P^3\hat{f} = \hat{f} - f - Pf - P^2 f + 3\eta \quad\Longrightarrow\quad \cdots,$$

one might expect the solution to take the form,

$$\hat{f} = \sum_0^\infty P^i (f - \eta)$$

This is true in the M/M/1 example, and one can also show that

$$\hat{f}(x) = \mathsf{E}_x \left[\sum_0^{\tau_\vartheta - 1} (f(X(t)) - \eta) \right]$$

where $\tau_\vartheta = \min(t \geq 1 : X(t) = \vartheta)$. The function f is linear, so its "integral" $\hat{f}$ is quadratic. This qualitative result carries over to general network models, and is a useful observation when developing the theory of network optimization.

We can similarly write,

$$\pi(x) = \frac{\mathsf{E}_\vartheta \left[\sum_{t=0}^{\tau_\vartheta - 1} \mathbb{1}(X(t) = x) \right]}{\mathsf{E}_\vartheta \left[\tau_\vartheta \right]}$$

provided $\mathsf{E}_\vartheta[\tau_\vartheta] < \infty$.

We next develop operator-theoretic formula which generalize easily to general state space processes.

2.2 The resolvent

These foundations are most easily formulated in terms of the resolvent,

$$R(x, y) = \sum_{t=0}^{\infty} 2^{-(t+1)} P^t(x, y), \qquad x, y \in \mathsf{X}. \tag{4}$$

The following *minorization condition* is basic:

Assumption A There is one state $\vartheta \in \mathsf{X}$ with

$$R(x, \vartheta) > 0 \qquad \text{for all } x \in \mathsf{X}.$$

That is, there is a single communicating class, reachable from any initial condition.

Assumption A will be satisfied in models on a general state space where there is an atom. For example, in a queueing network model, the state ϑ

might represent the network state where every buffer is empty. A rewriting of this condition brings Assumption A to a form which can be generalized further. Let ν denote the probability on X which is concentrated at ϑ. Assumption A implies that there exists a function $s\colon \mathsf{X} \to (0,1)$ with

$$R(x,y) \ge s(x)\nu(y) \qquad \text{all } x,y \in \mathsf{X}. \tag{5}$$

We can in fact take $s(x) := R(x,\vartheta)$, $x \in \mathsf{X}$. However, the inequality (5) is all that we require. When (5) holds for a positive s and a probabilityν then the chain is called *ψ-irreducible.*

The inequality (5) is a matrix in equality when R is viewed as a matrix. We can write this compactly as,

$$R \ge s \otimes \nu$$

where the right hand side is the outer product of the column vector s, and the row vector ν.

We have the relolvent equation

$$PR \;=\; RP \;=\; 2R - I. \tag{6}$$

From this it follows that any π is P-invariant if and only if it is R-invariant.

From the resolvent equation and the lower bound assumed on R we can now give a roadmap for solving the desired equations. As motivation, suppose that we already have an invariant probability π, so that

$$\pi R = \pi.$$

Then, on subtracting $s \otimes \nu$,

$$\begin{aligned} \pi(R - s\otimes\nu) &= \pi R - \pi[s\otimes\nu] \\ &= \pi - \pi(\mathsf{X})\cdot\nu = \pi - \nu \end{aligned}$$

Rewriting these equations gives

$$\pi[I - (R - s\otimes\nu)] = \nu.$$

We can now attempt an inversion. The point is, the operator $\Delta_R := I - R$ is not invertible, but by subtracting the outer product $s \otimes \nu$ there is some hope in attempting an inversion. Define the *potential kernel* as

$$G = \sum_{n=0}^{\infty} (R - s\otimes\nu)^n$$

Under certain conditions we do have $G = [I - (R - s \otimes \nu)]^{-1}$, and hence

$$\pi = \nu G$$

Now we can attempt the 'forward direction'. Given a pair s, ν satisfying the lower bound (5), we *define* $\mu := \nu G$. We must then answer two questions: (i) when is μ invariant? (ii) when is $\mu(X) < \infty$? If affirmative, then we do have an invariant probability measure, given by

$$\pi(A) = \frac{\mu(A)}{\mu(X)}, \qquad A \subset \mathsf{X}.$$

We will show that μ always exists, and that it is always *subinvariant*,

$$\mu(A) \geq \sum_{x \in \mathsf{X}} \mu(x) R(x, A)$$

Invariance and finiteness both require some form of *stability* for the process. The following is the key step in establishing subinvariance, and criteria for invariance. For $N \geq 0$, define $F_N : \mathsf{X} \to \mathbb{R}_+$ by

$$F_N = \sum_0^N (R - s \otimes \nu)^n s.$$

Lemma 2.1 *Under Assumption A, provided* (s, ν) *satisfy (5) we have*

(i) $\beta_N := \sup_x F_N(x) \leq 1$ *for all* N.

(ii) $F_N \uparrow Gs$, *as* $N \uparrow \infty$.

(iii) $\mu(s) \leq 1$.

Proof The proof of (i) is by induction: For $N = 0$, $f = s$, and $s \leq 1$ by assumption.

If true for N then

$$\begin{aligned} F_{N+1}(x) &= (R - s \otimes \nu) F_N(x) + s(x) \\ &\leq (R - s \otimes \nu) \mathbb{1}(x) \beta_N + s(x) \\ &= [R(x, \mathsf{X}) - s(x)\nu(\mathsf{X})]\beta_N + s(x) \\ &= \beta_N + (1 - \beta_N) s(x) \\ &\leq 1, \end{aligned}$$

where in the last equation we have again used the fact that $s \leq 1$.

The proof of (ii) is then the definintion of G, and (iii) is the definition of μ:

$$\mu(s) = \nu G s = \lim_{N\to\infty} \nu(F_N) \leq 1.$$

□

It is now easy to establish subinvarance:

Theorem 2.2 *For a ψ-irreducible Markov chain, the measure $\mu = \nu G$ is always subinvariant. Writing $\alpha = \nu G s$, we have*

(i) $\alpha \leq 1$*;*

(ii) *μ is invariant if and only if $\alpha = 1$.*

(iii) *μ is finite if and only if $\nu G(\mathsf{X}) < \infty$.*

PROOF Result (i) is already given in the lemma, and (iii) is just a restatement of the definition of μ. For (ii), write

$$\begin{aligned} \mu R &= \sum_0^\infty \nu(R - s\otimes\nu)^n R \\ &= \sum_0^\infty \nu(R - s\otimes\nu)^{n+1} + \sum_0^\infty \nu(R - s\otimes\nu)^i s\otimes\nu \\ &= \mu - \nu + \alpha\nu \leq \mu. \end{aligned}$$

□

It turns out that the case where $\alpha = 1$ amounts to a form of recurrence. In the countable state space case considered here, under Assumption A, one can show that $\alpha = 1$ if and only if $P(\tau_\vartheta < \infty \mid X(0) = \vartheta) = 1$.

Finiteness of μ is a positivity condition which, under the conditions imposed here, is equivalence to a finite mean return time to ϑ:

$$\mathsf{E}_\vartheta[\tau_\vartheta] < \infty.$$

This will be explored further in Section 3.

2.3 Solving Poisson's equation

To solve Poisson's equation (2) we will again use the resolvent. From (2) we have

$$RP\hat{f} = R\hat{f} - Rf + \eta$$

and on combining this with the resolvent equation (6) in the form $RP = 2R - I$ we arrive at

$$R\hat{f} = \hat{f} - Rf + \eta. \tag{7}$$

If we do manage to obtain a solution to (7), then this solution will also solve (2).

The solution $\hat{f}$ is not unique since we can always add a constant. This gives us some flexibility: *assume* that $\nu(\hat{f}) = 0$, so that $(R - s \otimes \nu)\hat{f} = R\hat{f} = \hat{f} - Rf + \eta$. This leads to a familiar looking identity,

$$[I - (R - s \otimes \nu)]\hat{f} = Rf - \eta$$

so that we can hope to write

$$\hat{f} = [I - (R - s \otimes \nu)]^{-1}(Rf - \eta),$$

which is equivalently expressed,

$$\hat{f} = G(Rf - \eta).$$

Provided the right-hand side is finite, this does solve (2).

3 Criteria for stability

For the $M/M/1$ queue, an invariant probability exists if and only if $\rho = \lambda/\mu < 1$. This means that the process drifts downward. In fact, $X(t \wedge \tau_\vartheta)$ is a supermartingale:

$$\mathsf{E}[X(t+1) \mid X_0^t] = X(t) - (\mu - \lambda), \qquad t < \tau_\vartheta.$$

We can also define a *Lyapunov function*,

$$V(x) = \frac{x}{\mu - \lambda}, \qquad x \in \mathsf{X},$$

which satisfies

$$PV\,(x) := \mathsf{E}[V(X(t+1)) \mid X_t = x] = V(x) - 1, \qquad x \neq \vartheta.$$

The existence of a Lyapunov function is known to be equivalent to positive recurrence. This is summarized in the following.

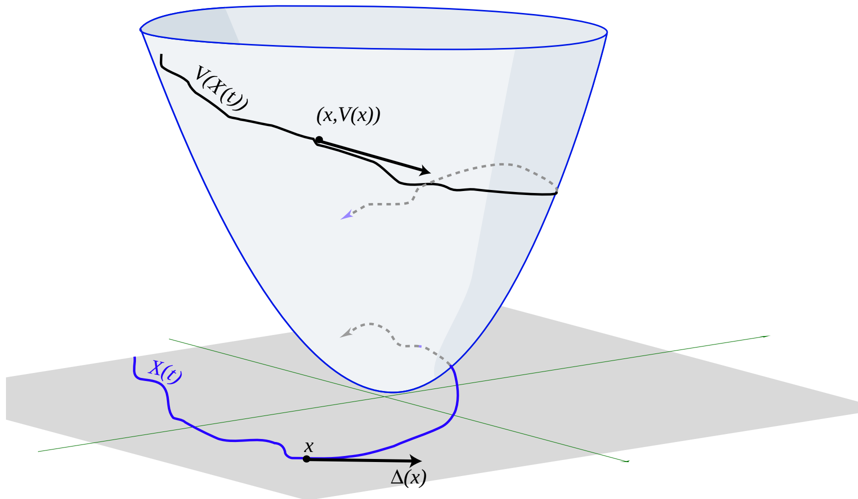

Figure 2: $V(X(t))$ is decreasing outside of the set S.

Theorem 3.1 (Fosters Criterion) *The following are equivalent for a ψ-irreducible Markov chain*

(i) *An invariant probability π exists*

(ii) $\mathsf{E}_\vartheta[\tau_\vartheta] < \infty$ *for one ϑ.*

(iii) *There exists $V : \mathsf{X} \to (0, \infty]$, finite for one $\vartheta \in \mathsf{X}$, a finite set $S \subset \mathsf{X}$ such that*

$$\mathsf{E}[V(X(t+1)) \mid X(t) = x] \leq V(x) - 1, \qquad x \in S^c.$$

If (3) holds then

$$\mathsf{E}_x[\tau_\vartheta] \leq V(x), \qquad x \in S^c.$$

Proof We just prove the implication (iii) → (i). We can write (iii) as

$$PV \leq V - f + b\mathbb{1}_S,$$

where $f \equiv 1$, and $b < \infty$. Hence,

$$RPV \leq RV - Rf + bR\mathbb{1}_S$$

By ψ-irreducibility (ϑ is reachable), there exists an integer n, and $\epsilon > 0$ such that

$$\sum_{i=1}^{n} P^i(x,\vartheta) \geq \epsilon \mathbb{1}_S(x), \qquad x \in \mathsf{X}.$$

Applying R to both sides then gives

$$\begin{aligned} R\mathbb{1}_S &\leq \epsilon^{-1} \sum_{i=1}^{n} RP^i(x,\vartheta) \\ &\leq \text{const.} R(x,\vartheta) \end{aligned}$$

If we set $s(x) = R(x,\vartheta)$ then we can thus find a $b_1 < \infty$ such that

$$RPV \leq RV - Rf + b_1 s.$$

From the resolvent equation (6) we then have

$$RV \leq V - Rf + b_1 s.$$

and hence also,

$$(R - s \otimes \nu)V \leq V - Rf + b_1 s.$$

Iterating this inequality gives

$$\begin{aligned} (R - s \otimes \nu)^2 V &\leq (R - s \otimes \nu)(V - Rf + b_1 s) \\ &\leq V - Rf + b_1 s \\ &\quad -(R - s \otimes \nu)Rf \\ &\quad +(R - s \otimes \nu)s. \end{aligned}$$

$$\begin{aligned} 0 \leq (R - s \otimes \nu)^n V &\leq V - \sum_0^{n-1} (R - s \otimes \nu)^i Rf \\ &\quad + b_1 \underbrace{\sum_0^{n-1} (R - s \otimes \nu)^i s}_{\leq 1 \text{ all } n.} \end{aligned}$$

Rearranging terms then gives,

$$\sum_0^{n-1} (R - s \otimes \nu)^i Rf \leq V + b_1,$$

and thus

$$GRf \le V + b_1.$$

Since $f \equiv 1$ we can then write $GRf\,(x) = G(x, \mathsf{X}) \le V + b_1$. Thus,

$$\mu(\mathsf{X}) = \nu G \mathbb{1} \le \nu(V) + b_1.$$

The minorization and the drift inequality (iii) give

$$(s \otimes \nu)(V) \le RV \le V - 1 + b_1 s,$$

which establishes finiteness of $\nu(V)$. □

In the same way we can prove,

Theorem 3.2 *Suppose that $f : \mathsf{X} \to [1, \infty)$, $V : \mathsf{X} \to (1, \infty)$ and that for a finite set S, and $b < \infty$,*

$$PV \le V - f + b\mathbb{1}_S.$$

Then X is positive recurrent, and there exists a solution to Poisson's equation

$$P\hat{f} = \hat{f} - f + \eta$$

where $\eta = \pi(f) < \infty$.

The solution satisfies $|\hat{f}(\mathsf{X})| \le V(x) + b_1$ for some $b_1 < \infty$.

If $\pi(V^2) := \sum_{x \in \mathsf{X}} V^2(x)\pi(x) < \infty$, then the central limit theorem holds for f,

$$\frac{1}{\sqrt{N}} \sum_{0}^{N-1} (f(X(t)) - \eta) \xrightarrow{\mathrm{d}} N(0, \gamma^2).$$

□

4 The mean ergodic theorem

We conclude by showing that the existence of a Lyapunov function leads to the ergodic theorems (3) which were cited as one of our goals in the introduction. A necessary condition for the mean ergodic theorem is *aperiodicity.* Under Assumption A, the Markov chain $\mathbf{X}$ is aperiodic provided $P^n(\vartheta, \vartheta) > 0$ for all n sufficient large.

4.1 Coupling

Coupling is a way of comparing the behavior of the process of interest $\mathbf{X}$ with another process $\mathbf{Y}$ which is already understood. To obtain ergodic theorems, the process $\mathbf{Y}$ can be taken as the stationary version of the process, with $Y(0) \sim \pi$. The ergodic theorems (3) obviously then hold for $\mathbf{Y}$.

Consider the bivariate chain

$$\Psi(t) = \begin{pmatrix} X(t) \\ Y(t) \end{pmatrix}, \qquad t \geq 0,$$

where X and Y are two copies of the chain with transition probability P, and different initial conditions. We define

$$T = \min(t : X(t) = Y(t) = \vartheta) = \min(t : \Psi(t) = \begin{pmatrix} \vartheta \\ \vartheta \end{pmatrix}).$$

To give a full statistical description of $\mathbf{\Psi}$ we need to explain how $\mathbf{X}$ and$\mathbf{Y}$ are related. We assume a form of conditional independence for $t < T$:

$$\begin{aligned} P(X(t+1) = x_1 \mid Y_0^t, X_0^t, X(t) = x_0) &= P(x_0, x_1) \\ P(Y(t+1) = y_1 \mid Y_0^t, X_0^t, Y(t) = y_0) &= P(y_0, y_1). \end{aligned}$$

For $t \geq T$ we assume that the chains coellesce so that $X(t) = Y(t)$ for $t > T$.

4.2 The coupling inequality

For any given $f : \mathsf{X} \to \mathbb{R}$ we have the *coupling inequality*,

$$|\mathsf{E}[f(X(t))] - \mathsf{E}[f(Y(t))]| \leq \mathsf{E}[|f(X(t)) \mid \mathbb{1}(T > t)] + E[f(Y(t))\mathbb{1}(T > t)].$$

One can show that $T < \infty$ a.s. if x is recurrent. Likewise, $\mathsf{E}[T] < \infty$ if $\mathbf{X}$ is positive recurrent. The latter can be proven by using the bivariate Lyapunov function $V_2(x, y) = max(V(x), V(y))$, $x, y \in \mathsf{X}$.

If we assume Y is stationary and f is bounded we obtain,

Theorem 4.1 *For any f, if $|f| \leq 1$ then,*

$$|\mathsf{E}_x[f(X(t))] - \pi(f)| \leq 2P_x(T > t) \to 0, \qquad t \to \infty.$$

□

4.3 Geometric ergodicity

Define the total variation norm between two measures μ and π by

$$\|\mu - \pi\| = \sup_{|f|\le 1} |\mu(f) - \pi(f)|.$$

Then Theorem 4.1 says

$$\|P^t(x,\,\cdot\,) - \pi(\,\cdot\,)\| \le 2\mathsf{P}(T > t \mid X(0) = x), \qquad x \in \mathsf{X},\ t \in \mathbb{Z}_+.$$

If we can control the tails of the coupling time T then we obtain a rate of convergence.

The chain is called *geometrically recurrent* if $\mathsf{E}[\exp(\epsilon T)] < \infty$ some $\epsilon > 0$. For such chains it is evident that the total variation norm above vanishes geometrically fast, in which case the chain is called *geometrically ergodic.*

The following extension of Theorem 4.1 shows that we can also move beyond bounded functions of $\mathbf{X}$ when the chain is geometrically recurrent. The V-total variation norm is given by

$$\|\mu - \pi\|_V = \sup_{|f|\le V} |\mu(f) - \pi(f)|.$$

Theorem 4.2 *The following are equivalent for an aperiodic, ψ-irreducible Markov chain*

(i) *The chain is geometrically recurrent.*

(ii) *There exists $V : \mathsf{X} \to [1,\infty]$ with $V(x) < \infty$ for some x, $\lambda < 1$, $b < \infty$, and a finite set S such that*

$$PV \le \lambda V + b\mathbb{1}_S$$

(iii) *For some $r > 1$ and some $M(x) < \infty$*

$$\sum_{n=0}^{\infty} \|P^n(x,\cdot) - \pi(\cdot)\| r^n \le M(x) \quad x \in \mathsf{X}.$$

If (i) holds then M can be taken as $M = cV$ for some c.

If any of the above conditions hold, then with V given in (ii), we can find $r > 1$ and $b < \infty$ such that

$$\|P^t(x,\,\cdot\,) - \pi(\,\cdot\,)\|_V \le br^{-t}V(x), \qquad x \in \mathsf{X},\ t \in \mathbb{Z}_+.$$

□

References

[1] V. S. Borkar. *Topics in controlled Markov chains.* Pitman Research Notes in Mathematics Series # 240, Longman Scientific & Technical, UK, 1991.

[2] P. W. Glynn and S. P. Meyn. A Lyapunov bound for solutions of Poisson's equation. *Ann. Probab.*, 24, April 1996.

[3] Henderson, S. G. *Variance Reduction Via an Approximating Markov Process.* Ph.D. thesis. Department of Operations Research, Stanford University. Stanford, California, USA, 1997.

[4] S.G. Henderson and S.P. Meyn. Variance reduction for simulation in multiclass queueing networks. *submitted to the IIE Transactions on Operations Engineering: special issue honoring Alan Pritsker on simulation in industrial engineering*, 1999.

[5] A. Hordijk. *Dynamic Programming and Markov Potential Theory.* 1977.

[6] R. A. Howard. *Dynamic Programming and Markov Processes.* John Wiley and Sons/MIT Press, New York, NY, 1960.

[7] P.R. Kumar and S.P. Meyn. Duality and linear programs for stability and performance analysis queueing networks and scheduling policies. *IEEE Transactions on Automatic Control*, 41(1):4–17, January 1996.

[8] S. Kumar and P. R. Kumar. Performance bounds for queueing networks and scheduling policies. *IEEE Trans. Automat. Control*, AC-39:1600–1611, August 1994.

[9] H. Kwakernaak and R. Sivan. *Linear Optimal Control Systems.* Wiley-Interscience, New York, NY, 1972.

[10] S. Lippman. Applying a new device in the optimization of exponential queueing systems. *Operations Research*, 23:687–710, 1975.

[11] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability.* Springer-Verlag, London, 1993.

[12] S.P. Meyn. Stability, Performance Evaluation, and Optimization To appear as a chapter in *MARKOV DECISION PROCESSES: Models, Methods, Directions, and Open Problems*, 2000.

[13] E. Nummelin. *General Irreducible Markov Chains and Non-Negative Operators.* Cambridge University Press, Cambridge, 1984.

[14] E. Nummelin. On the Poisson equation in the potential theory of a single kernel. *Math. Scand.*, 68:59–82, 1991.

[15] I. Ch. Paschalidis D. Bertsimas and J. N. Tsitsiklis. Scheduling of multiclass queueing networks: Bounds on achievable performance. In *Workshop on Hierarchical Control for Real–Time Scheduling of Manufacturing Systems*, Lincoln, New Hampshire, October 16–18, 1992.

[16] M. L. Puterman. *Markov Decision Processes.* Wiley, New York, 1994.